INVESTIGATING THE POTENTIAL OF SUPER-RESOLUTION FOR ROAD SEGMENTATION IN SENTINEL-2 IMAGES

Sandeep Kumar Jangir, Corentin Henry, Nina Merkle
Remote Sensing Technology Institute, German Aerospace Center (DLR), Oberpfaffenhofen, Germany
sandeep.jangir@dlr.de, corentin.henry@dlr.de, nina.merkle@dlr.de

***Abstract*—Road segmentation from Sentinel-2 imagery is challenging due to its coarse 10 m spatial resolution, yet its global coverage makes it valuable for applications like disaster relief and infrastructure monitoring. Traditional segmentation methods rely on high-resolution data, but recent approaches have explored super-resolution to enhance the spatial resolution of Sentinel-2 images. This study investigates the potential of super-resolution to perform road segmentation at 62.5 cm resolution from single-image Sentinel-2 RGB data, bridging the resolution domain gap. Both the super-resolution and the segmentation models are trained on high-resolution data, making the task more difficult. We demonstrate that these models can generalize to low-resolution data and deliver usable results for various applications, particularly in regions lacking up-to-date high-resolution imagery.**

***Index Terms*—Road segmentation, Sentinel-2, Super-resolution, Deep learning, Image enhancement**

I. Introduction

Using Sentinel-2 imagery to perform road segmentation presents unique challenges and opportunities due to its low spatial resolution but short revisit period and multispectral capabilities, making it well suited for large-scale monitoring and mapping applications. Road segmentation from these images can provide valuable insights into infrastructure development, particularly in remote or underdeveloped areas where high-resolution data might not be available. While Sentinel-2 does not capture fine-grained details, its global coverage and spectral richness make it an essential tool for large-scale mapping, disaster response, or environmental monitoring.

The recent advancements in deep learning models [1], particularly in terms of image quality, have made deep learning-based approaches the preferred choice for various image enhancement tasks, such as Super-Resolution (SR). Deep learning-based SR methods can be categorized into two types: Single-Image Super-Resolution (SISR), which uses information from a single image for upsampling, and Multi-Image Super-Resolution (MISR), which aggregates and merges data from multiple images to produce a higher-quality result. However, the low-resolution data of satellites such as Sentinel-2 makes resolving fine features challenging. Many methods [2–4] use time series to upsample the 20 m and 60 m bands to 10 m. In contrast, SISR methods rely on a single observations to generate SR images. For instance, Lanaras et al. [5] and Salgueiro et al. [6] upsample the 60 m and 20 m bands to 10 m using information from the original 10 m band and single observations. In Galar et al. [7] a SR model is trained on high-resolution satellite data, such as WorldView, and applied to super-resolve Sentinel-2's 10 m bands. However, the primary challenge of achieving a resolution finer than 10 m lies in the absence of Sentinel-2 bands with a lower GSD. To address this, one approach is to train SR networks on higher-resolution satellite images which share overlapping image profiles, similar to Sentinel-2's.

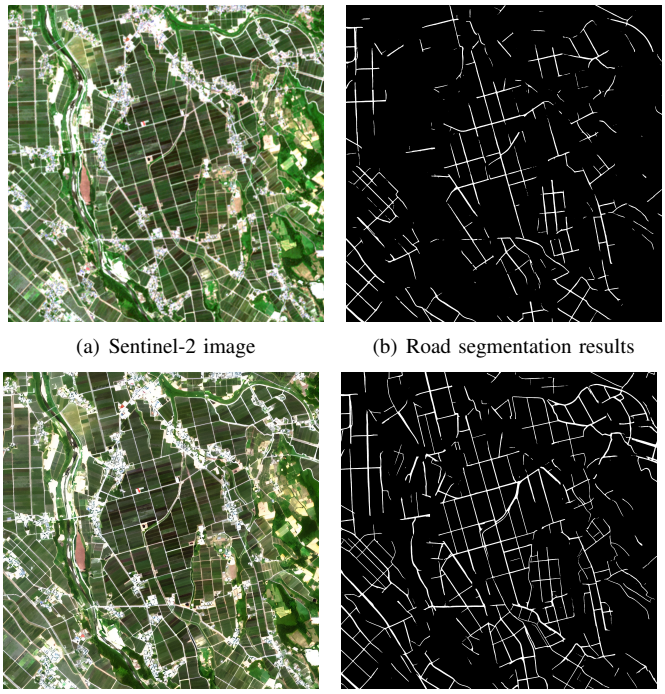

(a) Sentinel-2 image (b) Road segmentation results

(c) Enhanced Sentinel-2 image (d) Road segmentation results

Fig. 1. Road segmentation with and without super-resolution over Japan.

Road segmentation is usually performed on imagery acquired at 1 m GSD or below, but rarely above, as streets are no longer distinguishable in cities at coarser resolutions. Most datasets focus on images acquired between 1 m and 10 cm GSD, e.g. SpaceNet 3 [8] at 30 cm and DeepGlobe18 at 50 cm [9], where the roads width makes them easier to detect for machine learning models. Nevertheless, Ayala et al. [10] extract roads at 2.5 m GSD from native 10 m GSD

Sentinel-2 RGB images, using a U-Net-like architecture with a ResNet34 encoder backbone. They obtain the final 2.5 m output resolution via 4× bicubic upsampling of the input images. Trier et al. [11] also present a U-Net-based method to extract forest roads using 10 spectral bands of Sentinel-2 images at 10 m and 20 m GSD, upsampling the latter to 10 m via bilinear interpolation. They report getting a high rate of false negatives and false positives when using single images, and suggest using time series instead. Looking similarly for unofficial roads in the Amazonian for forest monitoring and therefore with a focus on non-urban areas, Botelho et al. [12] purposefully use the SWIR1, NIR, and red bands from 10 m Sentinel-2 images and optical Green-Red-NIR Landsat images. Using a modified U-Net model and a post-processing algorithm for road reconnection and vectorization. In contrast, Jia et al. [13] do not rely on hyperspectral bands and instead use settlement footprints to inform their model as to the likely proximity of roads. They train their model on Sentinel-2 RGB-NIR images at 10 m GSD. A final upsampling module learns to classify the roads at a 2.5 m GSD. Aiming at a 50 cm GSD, Sirko et al. [14] train a teacher HRNet model fed with single frames of 50 cm GSD satellite imagery and distill its knowledge into a student modified HRNet model fed with 32-layer multi-temporal 10 m GSD Sentinel-2 13-band imagery. The latter directly predicts the roads at a 50 cm GSD from the 10 m input image, achieving qualitatively high results.

In this study, we investigate the potential of super-resolution techniques to improve the performance of road segmentation models applied to Sentinel-2 RGB imagery. Neither the applied SR model nor the road segmentation network have been trained on low-resolution data, which presents a significant challenge in achieving optimal results. Our aim is to investigate whether these models can generalize to low-resolution Sentinel-2 imagery to support applications such as humanitarian aid with up-to-date road maps in areas where current high-resolution imagery is not available in time. To this end, we test the ability of the segmentation model to generalize to four areas worldwide and discuss the impact of SR on road segmentation model performance.

II. Methods

To extract the road network from Sentinel-2 imagery, we follow a two-step approach. In the first step, we use a SR framework to improve the spatial resolution of the Sentinel-2 images (see Section II-A). As a second step, we extract the road network from enhanced Sentinel-2 images using a deep-learning based approach (see Section II-B).

A. Single Sentinel-2 image super-resolution

We use U2D2 [15], a generalized framework for enhancing aerial and satellite imagery. U2D2 comprises three components: the Degradation Module (DM), which simulates degradations like noise, blurring, and artifacts to generate Low-Quality (LQ) images for training; the DL-Upsampler (DLU), a lightweight CNN for 2× SR which upscales and enhances images; and DiffRef, an adapted diffusion model based on SR3 [16], refining 2× upsampled images for sharper, more detailed outputs. U2D2's modular structure decomposes the enhancement task into distinct stages, allowing the DLU and DiffRef networks to operate independently and produce enhanced images. During training, the DM processes datasets by simulating degradations online at varying GSDs (15 cm–1.5 m) using bicubic downsampling. The DLU network removes these degradations while performing 2× upscaling. Its output, combined with the bicubic-upsampled input, is refined by DiffRef to produce enhanced SR images with sharper details and fewer artifacts. U2D2 effectively generates enhanced images (2× SR, denoising, and detail restoration) for aerial and satellite data outside its training distribution, achieving consistent performance across GSDs below 1.5 m. This consistently benefits downstream applications compared to both original data and other SR methods. In this study, we evaluate U2D2 on Sentinel-2 data, enhancing the 10 m RGB bands to a lower GSD without retraining. Previously, we showed that 2× SR balances preserving image details and generating new features. To evaluate enhanced Sentinel-2 images on a road segmentation network requiring high-resolution inputs, we target a GSD of 1.25 m. We explore an approach combining SR and bicubic interpolation to enhance Sentinel-2 images to 1.25 m GSD step by step: 10 m→5 m→2.5 m→1.25 m (as summarized in Table I). This strategy preserves original information while achieving the resolution required for downstream applications.

TABLE I
Overview of the upsampling strategies to reach 0.625 m, from the native 10 m GSD of Sentinel-2.

Exp.	10→ 5 m	5→ 2.5 m	2.5→ 1.25 m	1.25→ 0.625 m
#1	Bicubic	Bicubic	Bicubic	Bicubic
#2	Bicubic	Bicubic	SR	Bicubic
#3	Bicubic	SR	Bicubic	Bicubic
#4	Bicubic	SR	SR	Bicubic
#5	SR	SR	Bicubic	Bicubic
#6	SR	SR	SR	Bicubic

B. Road segmentation

We use a OneFormer segmentation model [17], a model architecture that enables us to benefit from the learning capabilities of vision transformer networks which were empirically shown to converge faster and reach a performance equivalent or higher to convolutional models on large-scale training sets [18]. It is made of a Swin-Tranformer Tiny encoder backbone [19] and a Multi-Scale Deformable transformer-based Attention module (MSDeformAttn) [20], followed by another multi-scale transformer-based decoder learning to predict more fine-grained spatial and semantic details. The model outputs a road segmentation mask at a $\frac{1}{4}$ resolution compared to the input image, necessitating a bilinear resampling to obtain the final full-resolution map. This architecture, when trained for road segmentation, has been shown to generalize well to complex areas throughout the world. For comprehensive details on the architecture, please refer to our original study [21].

III. Results & discussion

In the following, we provide an overview of our test areas and the corresponding annotations (Subsection III-B), the training parameters of the SR framework, the road segmentation model (Subsection III-A), and our results (Subsection III-C & Subsection III-D).

A. Test areas and ground truth data

We selected Sentinel-2 L2A image patches from four countries as our test images and use the 8-bit RGB bands:

- T30TVK_20241223T110623 (Spain)
- T32UPU_20241101T102059 (Germany)
- T37MBU_20240805T075902 (Kenya)
- 54SUH_20240620T013529 (Japan)

To provide a quantitative evaluation of the performance of our approach, we manually labeled all four images with vector road centerlines, using existing OpenStreetMap (OSM) data as an indication of the presence of roads. Roads missing in OSM but visible in the Sentinel-2 images are added to our annotations. In the case where roads are not visible, or are identified as dirt paths by looking at high-resolution satellite imagery from Google, they are not included in our annotations.

Since the road segmentation model has been trained on images acquired with a GSD mostly between 1 m and 30 cm, we apply a $2\times$ bicubic upsampling to 62.5cm to the SR-enhanced images before generating the prediction, as shown in Table I. After resampling, the evaluation effectively uses a patch size of 1024×1024 px with an overlap of 102 px. The logits for the overlapping regions are merged via an extrema operator before the probability calculation, so as to take the strongest predictions into account in the final road map.

B. Training parameters

The U2D2 model [15] was trained on aerial images from the open-source EAGLE dataset [22], with GSDs from 5 cm to 45 cm, plus two WorldView-4 (WV4) RGB images with a 30 cm GSD. To simulate varying levels of degradation, the training images were downsampled to generate LQ images (i.e. with GSDs from 10 cm to 1.5 m). The DLU network was trained over 500k iterations, starting with a learning rate of $2e^{-5}$, reduced to $2e^{-6}$ for the final 50k iterations. The DiffRef model uses a modified U-Net architecture for denoising during the diffusion process and was trained for 4.5M iterations. A linear noise schedule is applied across 500 time steps, with noise levels increasing gradually from $1e^{-4}$ to $1e^{-2}$, ensuring a smooth and stable progression through the diffusion process. For further details, please refer to the U2D2 paper [15].

Our OneFormer model is pre-trained on the ADE20K dataset [23] and trained on 6 aerial and satellite image datasets for road extraction with a worldwide coverage [21]: Massachusetts Roads [24], DeepGlobe18 [9], SpaceNet 3 & 5 [8, 25], RoadNet [26], and City-Scale [27]. We use an AdamW optimizer with an initial learning rate of $5e^{-5}$ and an exponential decay factor of 0.89 after each of the 20 training epochs. We use bootstrapped cross-entropy, binary cross-entropy and dice losses [28] with weights of 2, 5, and 5, respectively. We compensate OneFormer's non-deterministic optimization by training each epoch 4 times separately, proceeding with the checkpoint with the highest validation IoU.

TABLE II
Effect of upsampling strategies on road extraction results.

Exp.	Ups. Strategy	Quality ↑	Completeness ↑	Correctness ↑
#1	B-B-B-B	46.53 %	51.45 %	**86.94 %**
#2	B-B-SR-B	53.30 %	67.89 %	71.84 %
#3	B-SR-B-B	59.91 %	**77.45 %**	72.77 %
#4	B-SR-SR-B	56.30 %	69.15 %	76.20 %
#5	SR-SR-B-B	**60.97 %**	76.46 %	75.48 %
#6	SR-SR-SR-B	55.12 %	69.38 %	73.26 %

C. Super-resolution enhancement of Sentinel-2 images

Figure 2 presents the results of SR enhancement for Sentinel-2 images, improving the GSD from 10 m to 1.25 m using the U2D2 framework with the SR-SR-B-B strategy. This $\times 8$ resolution enhancement introduces significant new information in the final image. While smaller objects and structures that were indistinguishable in the original observation are rendered as blobs in the SR-enhanced image, larger isolated structures such as roads, buildings, and fields, which are vaguely distinguishable in the original 10 m image, are resolved with clear boundary definitions, though without detailed texture. Aliasing artifacts on roads are eliminated, resulting in improved roads definition that facilitates their detection by road segmentation algorithms. As the U2D2 framework is trained to remove artifacts commonly found in satellite images, RGB artifacts caused by stacking individual RGB bands are also eliminated. This improvement is visible on linear structures such as roads in Figure 2. However, in some cases, the framework may unintentionally remove color information, such as on rooftops, as the network might perceive them as artifacts, which is particularly noticeable in row two of Figure 2. Despite these minor limitations, it is worth noting that the U2D2 network was not specifically trained on Sentinel-2 or lower-resolution data. Nevertheless, this SR enhancement of low-resolution Sentinel-2 images has the potential to support downstream applications which can then provide a preliminary extraction of larger objects without requiring any model fine-tuning. For the road segmentation task, we aimed at finding the optimal strategy for upsampling native Sentinel-2 RGB image bands, ensuring that SR enhancement not only improves image quality but also increases the segmentation performance. We provide the corresponding analysis in the part below.

D. Road segmentation results

We report our results in terms of road quality, completeness, and correctness [29], metrics similar to IoU, recall, and precision, respectively, but adapted to road centerline labels. They convert the predictions into centerlines via morphological thinning, then estimate the relative proximity of the predictions and the labels within a controllable buffer distance. Here we choose a buffer diameter of 30 m, hence 3 pixels at the original

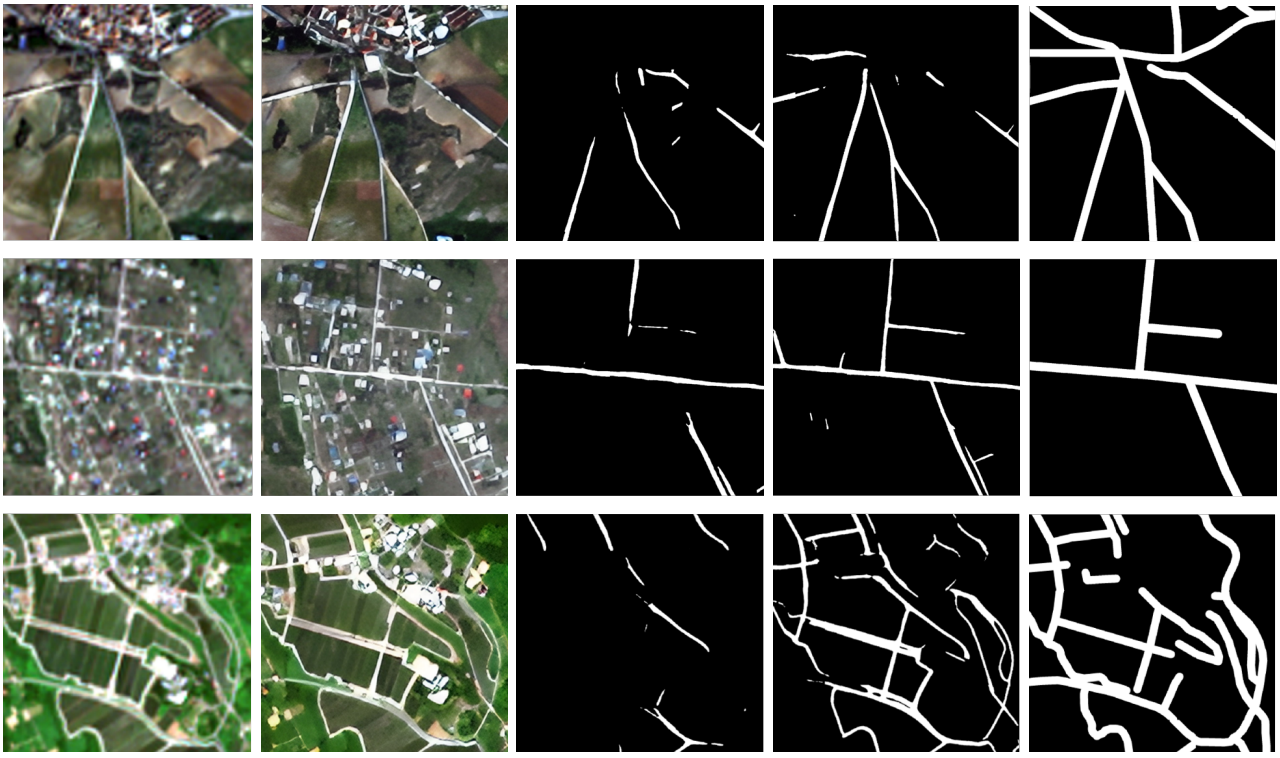

(a) Sentinel-2 (b) enhanced Sentinel-2 (c) pred. on Sen2 (d) pred. on enhanced Sen2 (e) road buffer labels

Fig. 2. Qualitative comparison of bicubic upsampling (a) and SR as in Exp #5 (b) of Sentinel-2 optical RGB images, with corresponding road segmentation labels and results (d-e) at 62.5cm GSD. The selected image patches were acquired over Spain (row 1), Kenya (row 2), and Japan (row 3).

10 m GSD of Sentinel-2 images, to account for the fact that the roads may be located up to 10 m away from their location as visible in the images. Moreover, to accommodate the relatively lower confidence of the model on the target data, we apply a 35% probability threshold on the road predictions. Using a total of 16× bicubic upsampling on the original Sentinel-2 images, we obtain a baseline quality of 46% with a medium, conservative completeness of 51% and a correspondingly high correctness of 87% (Exp. #1 in Table II). Replacing any of those 2× upsampling operations by our SR model yields better results in quality and completeness, though correctness, while still high, logically decreases as our segmentation model is able to predict more roads but includes some false positives. We observe that applying a 2× bicubic upsampling as a first step leads to the least performance increase overall (Exp. #2-4), lagging 1-7.5% behind the best performing model in terms of quality (Exp. #5). The latter uses two 2× SR upsampling at the first steps, followed by successive 2× bicubic upsampling until reaching a 62.5 cm GSD, and achieves a 61% quality, a 76% completeness, and a 75% precision. While other strategies helped the model get slightly better completeness or correctness (around 1% higher), they significantly underperform in the complementary metric (around 3-7% lower), which effectively leads to a lower quality score. On the other hand, adding too many 2× SR operations in a row creates artifacts, which confuse the model, leading to poor performance improvements with only 55% quality (Exp. #6). Figure 2 shows results for pure bicubic resampling of Sentinel-2 images and the approach of Exp. #5, with a clear increase in the level of details and road extraction comprehensiveness for all test areas: main roads, streets and services roads were successfully detected in the enhanced images whereas they were only partially detected, if at all, in the original images upsampled via bicubic interpolation.

IV. Conclusion & future works

In this study, we demonstrated the effectiveness of combining SR and road segmentation for enhancing Sentinel-2 imagery, improving spatial resolution from 10 m to 1.25 m. Despite not being trained on low-resolution data, the SR-enhanced images were shown to support road segmentation tasks, yielding better performance compared to conventional bicubic interpolation. Our experiments showed that a balanced combination of SR and bicubic interpolation led to the best segmentation quality and completeness. The proposed approach provides a valuable solution for applications like humanitarian relief and infrastructure monitoring, where high-resolution imagery is often unavailable. Future works should focus on refining the SR model for low-resolution data for 10 m and 20 m bands and exploring its application for the improvement of other downstream applications on diverse regions and domains with limited access to timely high-resolution data.

REFERENCES

[1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[2] Y. Zhu, C. Geiß, and E. So, "Image super-resolution with dense-sampling residual channel-spatial attention networks for multi-temporal remote sensing image classification," *International Journal of Applied Earth Observation and Geoinformation*, vol. 104, p. 102543, 2021.

[3] T. Tarasiewicz, J. Nalepa, R. A. Farrugia, G. Valentino, M. Chen, J. A. Briffa, and M. Kawulok, "Multitemporal and multispectral data fusion for super-resolution of Sentinel-2 images," *IEEE Transactions on Geoscience and Remote Sensing*, 2023.

[4] M. T. Razzak, G. Mateo-García, G. Lecuyer, L. Gómez-Chova, Y. Gal, and F. Kalaitzis, "Multi-spectral multi-image super-resolution of Sentinel-2 with radiometric consistency losses and its effect on building delineation," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 195, pp. 1–13, 2023.

[5] C. Lanaras, J. Bioucas-Dias, S. Galliani, E. Baltsavias, and K. Schindler, "Super-resolution of Sentinel-2 images: Learning a globally applicable deep neural network," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 146, pp. 305–319, 2018.

[6] L. Salgueiro, J. Marcello, and V. Vilaplana, "Single-image super-resolution of Sentinel-2 low resolution bands with residual dense convolutional neural networks," *Remote Sensing*, vol. 13, no. 24, p. 5007, 2021.

[7] M. Galar, R. Sesma, C. Ayala, L. Albizua, and C. Aranda, "Super-resolution of Sentinel-2 images using convolutional neural networks and real ground truth data," *Remote Sensing*, vol. 12, no. 18, p. 2941, 2020.

[8] A. V. Etten, D. Lindenbaum, and T. M. Bacastow, "Spacenet: A remote sensing dataset and challenge series," *CoRR*, vol. abs/1807.01232, 2018.

[9] I. Demir, K. Koperski, D. Lindenbaum, G. Pang, J. Huang, S. Basu, F. Hughes, D. Tuia, and R. Raskar, "DeepGlobe 2018: A challenge to parse the Earth through satellite images," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. Salt Lake City, UT: IEEE, Jun. 2018, pp. 172–17 209.

[10] C. Ayala, C. Aranda, and M. Galar, "Towards fine-grained road maps extraction using Sentinel-2 imagery," *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. V-3-2021, pp. 9–14, 2021. [Online]. Available: https://isprs-annals.copernicus.org/articles/V-3-2021/9/2021/

[11] Ø. Trier, A.-B. Salberg, R. Larsen, and O. Nyvoll, "Detection of forest roads in Sentinel-2 images using U-Net," in *Proceedings of the Northern Lights Deep Learning Workshop*, vol. 3, 03 2022.

[12] J. Botelho, S. C. P. Costa, J. G. Ribeiro, and C. M. Souza, "Mapping roads in the Brazilian Amazon with artificial intelligence and Sentinel-2," *Remote Sensing*, vol. 14, no. 15, 2022. [Online]. Available: https://www.mdpi.com/2072-4292/14/15/3625

[13] Y. Jia, X. Zhang, R. Xiang, and Y. Ge, "Super-resolution rural road extraction from Sentinel-2 imagery using a spatial relationship-informed network," *Remote Sensing*, vol. 15, no. 17, 2023. [Online]. Available: https://www.mdpi.com/2072-4292/15/17/4193

[14] W. Sirko, E. A. Brempong, J. T. C. Marcos, A. Annkah, A. Korme, M. A. Hassen, K. Sapkota, T. Shekel, A. Diack, S. Nevo, J. Hickey, and J. Quinn, "High-resolution building and road detection from Sentinel-2," 2024. [Online]. Available: https://arxiv.org/abs/2310.11622

[15] S. K. Jangir, M. N. Mühlhaus, N. Merkle, C. Henry, and R. Bahmanyar, "U2D2: A blind super resolution and enhancement framework for aerial and satellite images," 2025, submitted to IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing.

[16] C. Saharia, J. Ho, W. Chan, T. Salimans, D. J. Fleet, and M. Norouzi, "Image super-resolution via iterative refinement," *IEEE transactions on pattern analysis and machine intelligence*, vol. 45, no. 4, pp. 4713–4726, 2022.

[17] J. Jain, J. Li, M. Chiu, A. Hassani, N. Orlov, and H. Shi, "OneFormer: One transformer to rule universal image segmentation," in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 2989–2998.

[18] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2021.

[19] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 10 012–10 022.

[20] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, "Deformable {detr}: Deformable transformers for end-to-end object detection," in *International Conference on Learning Representations*, 2021.

[21] C. Henry and F. Fraundorfer, "Worldwide high-fidelity road extraction from aerial and satellite imagery enabled by low-fidelity openstreetmap labels," in *Pattern Recognition*, D. Cremers, Z. Lähner, M. Moeller, M. Nießner, B. Ommer, and R. Triebel, Eds. Cham: Springer Nature Switzerland, 2025, pp. 302–316.

[22] S. M. Azimi, R. Bahmanyar, C. Henry, and F. Kurz, "EAGLE: Large-scale vehicle detection dataset in real-world scenarios using aerial imagery," in *2020 25th international conference on pattern recognition (ICPR)*. IEEE, 2021, pp. 6920–6927.

[23] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, "Scene parsing through ADE20K dataset," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 5122–5130.

[24] V. Mnih, "Machine learning for aerial image labeling," Ph.D. dissertation, University of Toronto, 2013.

[25] The SpaceNet Partners, "SpaceNet 5: Automated road network extraction and route travel time estimation from satellite imagery," https://spacenet.ai/sn5-challenge/, 2020.

[26] Z. Zhang, Q. Liu, and Y. Wang, "Road extraction by deep residual U-Net," *IEEE Geoscience and Remote Sensing Letters*, vol. 15, no. 5, pp. 749–753, May 2018.

[27] S. He, F. Bastani, S. Jagwani, M. Alizadeh, H. Balakrishnan, S. Chawla, M. M. Elshrif, S. Madden, and M. A. Sadeghi, "Sat2Graph: Road graph extraction through graph-tensor encoding," in *European Conference on Computer Vision (ECCV)*, Aug. 2020.

[28] C. Henry, F. Fraundorfer, and E. Vig, "Aerial road segmentation in the presence of topological label noise," in *2020 25th International Conference on Pattern Recognition (ICPR)*, 2021, pp. 2336–2343.

[29] C. Wiedemann, C. Heipke, H. Mayer, and O. Jamet, "Empirical evaluation of automatically extracted road axes," in *Empirical Evaluation Techniques in Computer Vision*. IEEE Computer Society Press, 1998, pp. 172–187.