

Figure 1: Examples of synthetic satellite images generated with diffusion models conditioned on Open Street Maps (OSM). The real satellite images are provided as reference (2nd column) but they are not used at inference. We cover diverse landscapes (urbanised and rural areas).

Analysis



Figure 2: The model encodes class variances such as seasonality changes, lighting conditions, and human activity. We show examples on water bodies of multiple sizes and challenging urban areas.

Overview

- We use public sources [1,2] to create two large datasets of map-satellite image pairs over the regions of Mainland Scotland and the Central Belt
- We train a ControlNet [3] (with stable-diffusion-v1-5 as a frozen backbone) to generate realistic satellite images conditioned on cartographic data
- We present a qualitative evaluation of results and observe that we obtain both high quality and high fidelity images
- We explore historical map reconstruction
- We discuss the opportunities and challenges in using these models for remote sensing

Datasets



Figure 3: Sampling regions used for the dataset construction. We visualise some example pairs (map-satellite image). Mainland Scotland is largely rural, whereas the central belt has build up cities including Edinburgh and Glasgow.

Future work

- Diffusion models could be used for data augmentation; this is useful when data collection is expensive
- They can be used to generate private high-resolution datasets without compromising confidentiality
- They can be used to make adversarial datasets to train models for the identification of manipulated satellite imagery

References

- [1] Openstreetmap. 2017.
- [2] ArcGIS World Imagery (Clarity). February 2023.
- [3] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. ICCV 2023




Figure 4: Satellite images don't tend to match georeferenced historical maps for the same region due to a significant time difference (>100 years). However, after training the model with a large semi-paired dataset we observe consistent generations that closely follow the map layout.

Failure cases

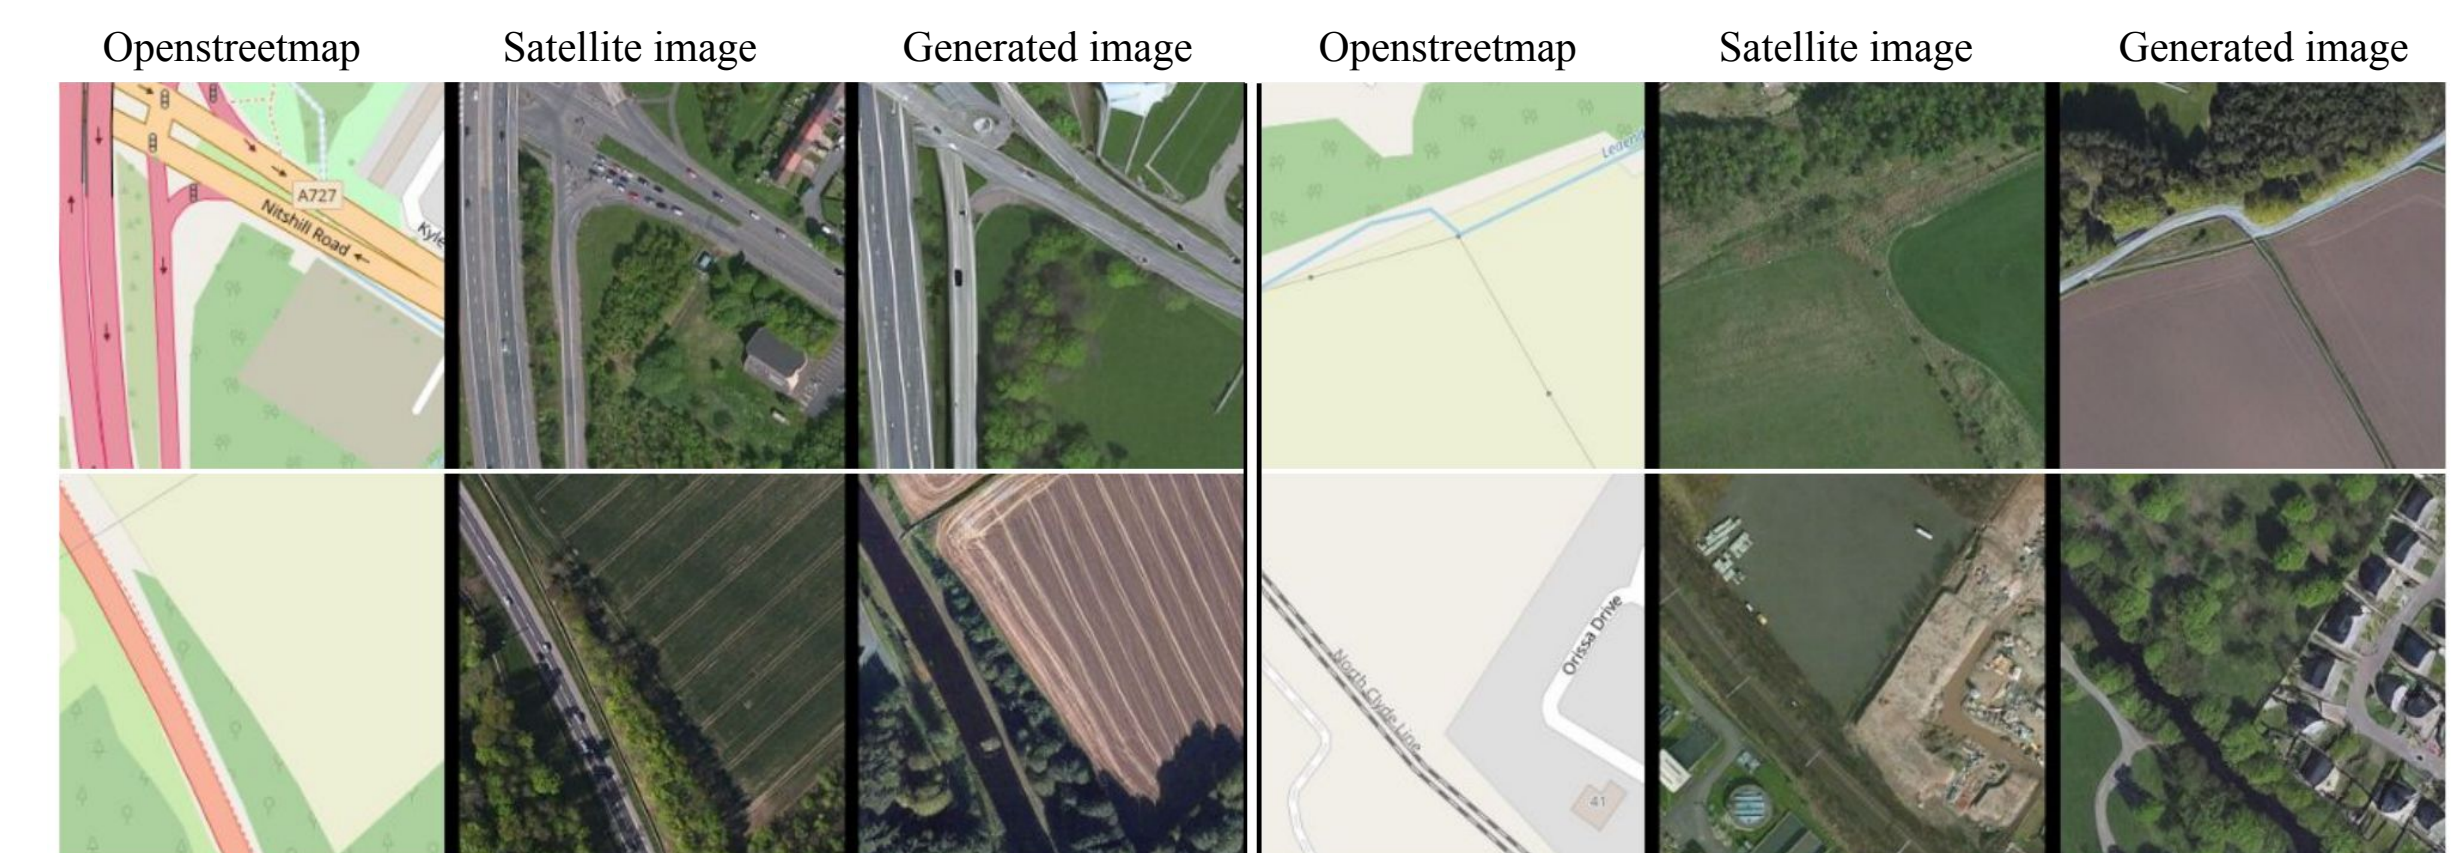


Figure 5: Our model finds it challenging to generate large roads with multiple lanes. It also struggles with intersections and overpasses. Rivers are sometimes mistaken for roads. Generating railroads is challenging; we attribute this to underrepresentation in our datasets.