ETH zürich

Chair of Photogrammetry and Remote Sensing

Report for Interdisciplinary Project - Spring 2022

Building Density Estimation from Low Resolution Satellite Data

Rushan Wang

Supervised by
Prof. Dr. Konrad Schindler
Dr. Rodrigo Caye Daudt
Nando Metzger

Acknowledgements

I would like to thank Prof. Dr. Konrad Schindler for giving me this great opportunity to work in this project. I also want to thank Nando Metzger and Dr. Rodrigo Caye Daudt for constantly giving me support and good advice during the project.

Abstract

Building density is a key ingredient for population estimation. This project defines building density as the number of buildings in each pixel and aims to achieve building density estimation using low-resolution satellite data. A deep learning method based on the fully convolutional network is implemented to build the estimation model. The model uses the architecture of U-Net, which is the state-of-art method in the field of fully convolutional networks. Two layers are output from the U-Net, performing the regression task and the segmentation task. Weighted binary cross-entropy is used to deal with unbalanced data. The result shows that our approach can give a decent output, especially on the classification task. However, the regression task still needs further improvement to predict more accurate building density. The model preforms better in lower building density area.

Keywords

Remote sensing; Building density; CNNs; Segmentation; Regression

Contents

1 Introduction

Africa is the world's second-largest and second-most-populous continent [1]. Its population has increased fast over the last decades and is consequently relatively young. With relatively high population growth and relatively old census data, it is desirable to be able to rapidly estimate the local population density in Africa. Such estimates are essential for health planning, refugee support, epidemiological modeling, and allocation of public resources and services[2]. Building density is an effective metric to estimate such population density. In addition, building density also plays an important part in many other aspects such as city planning, land management, and environment protection[3].

Previous methods use high-resolution satellite imagery to calculate building density. They refer to the building density as the ratio of the coverage of the buildings[4].

Nevertheless, high-resolution satellite data is not available all the time. On the one hand, high-resolution satellite data are not publicly accessible, and you have to pay for each data update. On the other hand, their spatial coverage is incomplete, which means that some land strips will have no image.

Google has recently released a dataset of building footprints of Africa, named Open Buildings dataset, where buildings were derived from high-resolution imagery. Even though this dataset has 516M building footprints data, it still has data missing in some important regions, for instance, North Kivu showed in Figure 1.

To fill in the gaps in areas with data missing, this project aims to develop a method to estimate building density from low-resolution satellite data. By collecting data from regions with known building footprints, a neural network can be trained to predict building density in each pixel based on low-resolution satellite imagery. Finally, we can apply the model to areas that lack of building data.

Nevertheless, the land surface of Africa has a wide diversity of terrain and building types, which can lead to many challenges[5]. For instance, the range of geological or vegetation features can be confused with built structures, also many contiguous buildings do not have clear delineations. Another challenging problem is the sparseness of buildings in the majority of regions of Africa, which cause unbalanced training data, with an extremely low ratio of pixels containing building.

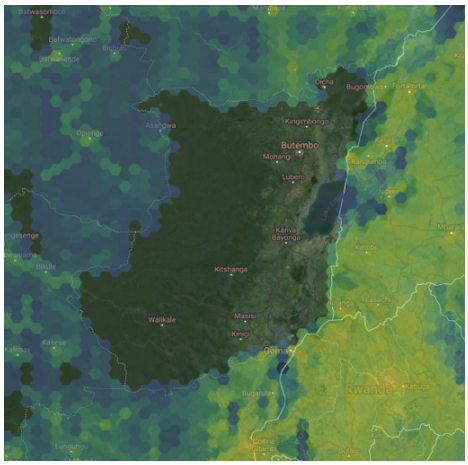

Figure 1: Data missing in North Kivu (The yellow color represents more building data, the blue color represents less building data, and no color covering the satellite imagery means no building data)

Source: https://sites.research.google/open-buildings/

In this report, the development of the building density estimation pipeline for Africa will be described.

The report is structured as follows. Section 2 presents the previous studies with convolutional neural networks in remote sensing and the previous method to estimate building density. Section 3 introduces the Open Building dataset and Sentinel-2 data used in this project. Section 4 draws the basic structure of the pipeline used, including data preprocessing, training network, and assessing metrics. Section 5 gives the results of this project. Section 6 summarizes the contributions of this project and gives outlooks for future study.

2 Related Work

2.1 Deep Convolutional Neural Networks

Deep convolutional neural networks(CNNs) have become the critical approach in remote sensing since they can automatically learn powerful representations from input imagery[6]. With multiple convolutional layers and pooling layers connected, deep CNNs can efficiently extract multi-level features from the spatial and spectral information of satellite imagery[7].

Various CNN-based approaches have been proposed for the pixel-wise semantic segmentation of satellite imagery. They take arbitrary-sized inputs and predict labels for each pixel. And they are generally based on fully convolutional networks(FCNs)[2]. The fully convolutional network classifies each pixel on the image.

This project is based on the U-Net model, which is a fully convolutional neural network that is designed to learn from fewer training samples. U-Net is an architecture developed for biomedical image segmentation in the beginning[8]. Figure 2 shows the architecture of the U-Net model.

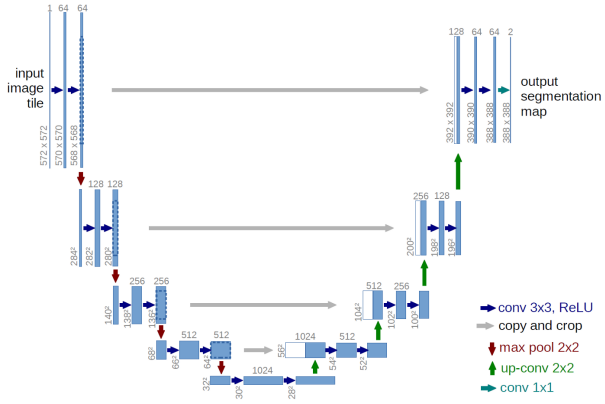

Figure 2: U-Net Architecture[8]

U-Net is a U-shaped encoder-decoder network architecture. It consists of four encoder blocks and four decoder blocks, which are connected through bridges. The encoder network extracts features and learns an abstract representation of the input image through a sequence of the encoder blocks[9]. At each encoder block, the network has half the spatial dimensions and double the number of feature channels. The decoder network takes the abstract representation and generates a semantic segmentation mask. At each decoder block, the network doubles the spatial dimensions and half the number of feature channels. The bridge is used to connect the encoder and the decoder network and to complete the flow of information.

Through this kind of architecture, U-Net can transfer the high frequency signal of the input to

the output, and learn features from fewer training data.

2.2 Building Estimation

Building detection has always been a popular research area[4]. Many scholars have used different methods to solve this problem.

Yu[10] used an object-based method to extract building objects automatically and then they computed building density with the LiDAR dataset. Wu[3] estimated urban building density by using high-resolution Synthetic Aperture Radar (SAR) images. They applied the empirical threshold method and ffmax-filter algorithm for the detection. Zhou[11] applied the CART algorithm and integrated SAR and optical data to estimate building density. CART is a classification and regression tree-based approach, which can handle both classification and regression tasks. In the work of Zhang[11], building density was estimated by using multiple features and support vector regression on optical very high-resolution satellite imagery.

However, none of the studies has proposed an approach using low-resolution satellite data, meaning those methods can not be applied to a large-scale area with incomplete data, which is exactly the current research gap. Therefore, in this project, we will conduct the building density estimation using low-resolution satellite data.

3 Data

3.1 Study Sites

We focus on overall six specific regions, five of which are in Africa and one region is in Iraq, Asia. Those regions were selected due to a lack of data and the fact that they all raise great interest in humanitarian proposals. Figure 3 shows the locations of those regions.

Tal Afar is a city in the Nineveh Governorate of northwestern Iraq with no official census data exists. North Kivu is a province bordering the eastern Democratic Republic of the Congo. The region is politically unstable and since 1998 has been one of the flashpoints of the military conflicts in the region. Maiduguri is the capital and the largest city of Borno State in north-eastern Nigeria, which has a lot of violence. Mokolo is the departmental capital and largest city of the Mayo-Tsanaga department, in the Far North Province of Cameroon. Juba is the capital and largest city of South Sudan. Montepuez is the second-largest city in the province of Cabo Delgado in Mozambique, after the provincial capital of Pemba. All six regions are politically unstable and lack accurate official demographic Data, indicating the significance of population density estimations in these regions.

Figure 3: The study sites

3.2 Open Buildings Dataset

The Open Buildings dataset[1] covers 516M building footprints across Africa. The building footprint data was detected by deep learning models on high-resolution aerial imagery at a continental scale. However, some area of Africa is not covered by this kind of high-resolution satellite data, leading to gaps in the dataset.

Moreover, the data quality of the Open Buildings dataset is not homogeneous, as shown in Figure 4. Therefore, it is essential to choose the training area with relatively high quality.

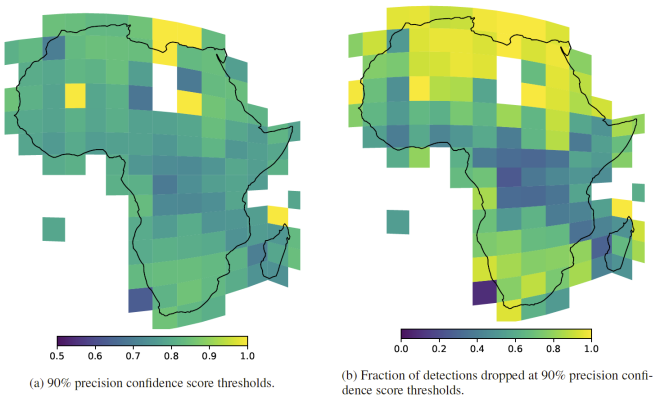

(a) 90% precision confidence score thresholds.

(b) Fraction of detections dropped at 90% precision confidence score thresholds.

Figure 4: Spatial variance of data quality in Open Buildings dataset

Source: https://sites.research.google/open-buildings/

3.3 Satellite Imagery

Sentinel-2 is an Earth observation mission from the Copernicus Program. It systematically acquires optical imagery at high spatial resolution (10 m to 60 m) over various terrain surface[12]. It has multi-spectral data with 13 bands in the visible, near-infrared, and short wave infrared parts of the spectrum.

Google Earth Engine (GEE) has powerful capabilities in accessing and processing massive volumes of multi-source, multi-temporal, multi-scale Earth Observation data through a cloud platform [13]. Available datasets in the GEE catalog include satellite imagery, geophysical data, climate and weather data, and demographic data.

In this project, 4 bands (Band2,3,4,8) of the Sentinel 2 Multispectral Instrument were used, all with 10m spatial resolution, as shown in Table 1. The data were available at the top of the atmosphere reflectance catalog in GEE. A cloud-free composite image was created from the image

[1] https://sites.research.google/open-buildings/

collection by making use of the provided cloud mask. To correspond to the time when the Open Buildings dataset was created, we captured the imagery in the time between 1 January and 31 December 2021. Therefore, the temporal effect of the change in buildings can be reduced.

For each study region, the north, south, west, and east sides of the neighbor area were chosen as training areas. For each satellite image, it contains approximately 50km $\times$ 50km area. Since the Open Buildings dataset does not cover Iraq, Egypt will be chosen as the training area for Tal Afar.

Bands	Central wavelength(nm)	Bandwidth(nm)	Spatial resolution(m)
Band 2 – Blue	492.4	66	10
Band 3 – Green	559.8	36	10
Band 4 – Red	664.6	31	10
Band 8 – NIR	832.8	106	10

Table 1: Spectral bands for the Sentinel 2.

4 Methodology

4.1 Preprocessing

The building footprints are aggregated to the same raster grid of 10m resolution as the Sentinel-2 imagery.

First, a raster file of the same size as the satellite image is created, which also has a 10m resolution. Then, the center of the building polygon is extracted and assigned to the corresponding pixel. Finally, the number of buildings for each pixel is calculated, which is the building density that would be used in this project. Figure 5 shows the ground truth we produced. Each pixel represents a 10m $\times$ 10m area, and the color represent different building density.

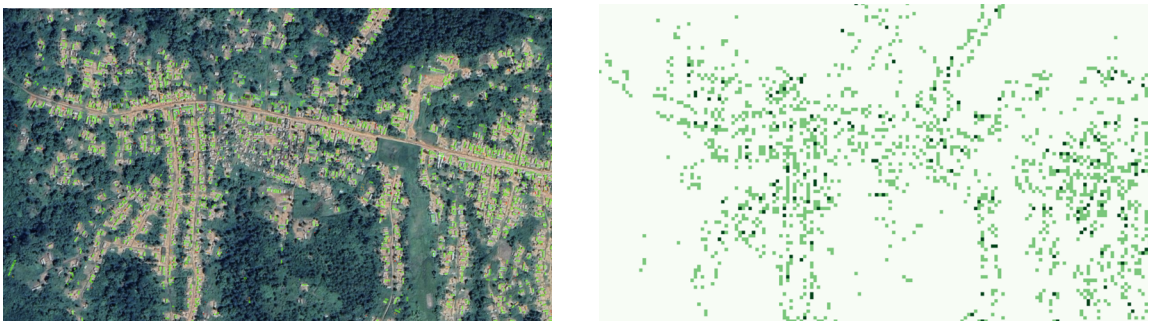

(a) Open Building data (b) Building density

Figure 5: Data Preprocessing

4.2 Overview of the proposed framework

Figure 6 shows the structure of the pipeline for this project. First, the 4 channels of satellite imagery will be fed into the U-Net network and output two layers. One output layer is to carry out the regression task, which gives each pixel the number of how many buildings it has. The other output layer is to carry out semantic segmentation to classify each pixel in the image as containing building or non containing building. Then a binary build-up mask can be created and added to the regression layer to preserve only the pixels that have buildings and set all the pixels without building as zero.

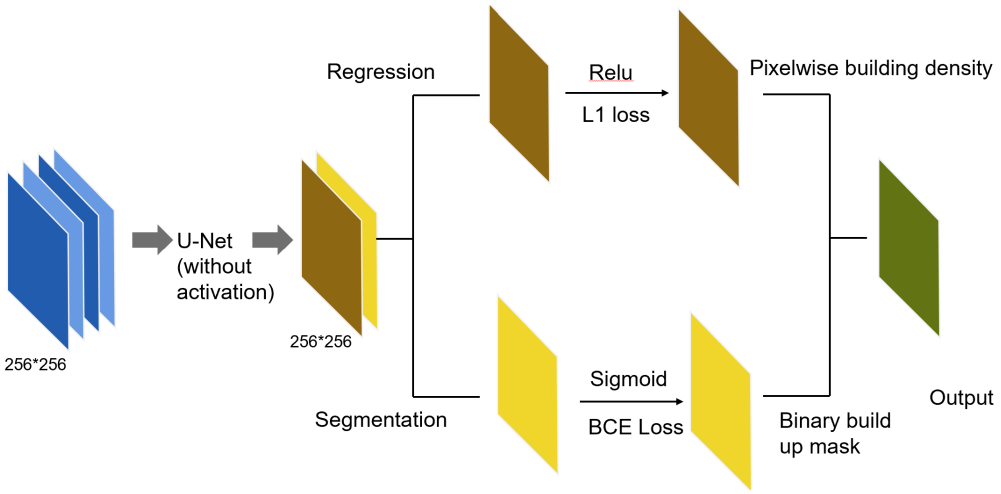

Figure 6: Training Pipeline

By using two output layers and conducting classification and regression tasks separately, it can prevent the model to predict every pixel as zero and give a more accurate estimation.

4.3 Loss functions

We calculate the overall loss as the combination of the two losses from the regression task and segmentation task. The equation of the overall loss function is shown as (1). λ is one of the parameters we will discuss in the following section.

$$\mathcal{L}_{overall} = \mathcal{L}_{L1} + \lambda \mathcal{L}_{BCE} \tag{1}$$

L1 loss is chosen as the loss function of the regression task. It is used to minimize the error which is the sum of all the absolute differences between the true value and the predicted value. The equation of L1 loss is shown in (2).

$$\mathcal{L}_{L1} = \sum_{i=1}^{n} |y_{true} - y_{predicted}| \tag{2}$$

For the loss function of the segmentation task, binary cross-entropy loss is used. Entropy is a measure of the uncertainty associated with a given distribution. The binary cross-entropy loss function calculates the loss by computing the average shown in (3).

$$\mathcal{L}_{BCE} = -\frac{1}{N} \sum_{i=1}^{N} y_i \cdot \log\left(p(y_i)\right) + (1 - y_i) \cdot \log\left(1 - p(y_i)\right) \tag{3}$$

4.4 Accuracy assessment

Since the model has two layers as output, which include regression and classification, various metrics are used to assess the results. For the segmentation task, overall accuracy, precision, recall, and F1 score would be used as the metrics. For the regression output, root mean square error(RMSE) would be used as the metric. It is worth noting that when assessing the result of the regression task, a binary building mask generated from the ground truth should be implemented in the layer at first.

5 Results

In this section, the experimental results of this project are presented.

5.1 Experimental setup

Each study region has four patches of satellite images selected, two of them are used for training, one patch is used for validation, and the other one is used for testing. For all the experiments, the Adam optimizer is used, and the number of training epochs is set to 40. The learning rate is set to 0.0001 if not mentioned. When reading all the data and conducting the training process, the patch size is set to 256 $\times$ 256 pixels.

5.2 Loss function parameter

Different value of λ is tested to see the effect on the result, as shown in Table 2. The model performs better when the loss function has λ as 100, which also matches the fact that the value of $\mathcal{L}_{\mathcal{L}\infty}$ is about 100 times larger than $\mathcal{L}_{BCE}$. When equalizing these two losses in the loss function, the network can be more effective in improving both of the performances.

λ	1	50	100	200
Accuracy(%)	99.10	99.95	99.96	99.95
Recall(%)	71.33	70.92	71.82	72.73
F1(%)	1.72	23.14	26.45	24.87
Precision(%)	0.63	9.81	11.65	10.91
RMSE	0.0951	0.0232	0.0213	0.0223

Table 2: Effect of different λ on the loss.

Besides λ, another setting for the binary cross-entropy loss is also tested. Since the cross-entropy loss evaluates the class predictions for each pixel separately and then averages the losses across all pixels, it is essentially learning equally for each pixel in the image. IF multiple classes are unevenly distributed in the image, this may lead to the training process being dominated by the class with a high number of pixels. That is to say, the model can primarily learn the features of the class with a high number of pixels and the learning model will be more biased in predicting the pixels of that class. In our case, the number of pixels without a building is much more than the number of pixels with buildings, which can easily lead to bias. To solve that unbalanced classes problem, we can add weight to the calculation of loss, to make the model learn more features of the class with a low number of pixels. The weight can be defined as the ratio between the number of two classes. And the weighted binary cross-entropy loss function can turn into the

following equation.

$$posweight = \frac{negnum}{posnum} \tag{4}$$

$$loss = -posweight \times y_{true} log(y_{pred}) - (1 - y_{true}) log(1 - y_{pred}) \tag{5}$$

Therefore, we try to test whether it is better to have weighted on the binary cross-entropy loss. Table 3shows the result of weighted and unweighted binary cross-entropy loss. We can see from the table that the recall of unweighted loss is much worse than the weighted loss.

Weighted	No	Yes
Accuracy(%)	99.99	99.94
Recall(%)	0.01	69.22
F1(%)	0.03	23.14
Precision(%)	0.02	9.62
RMSE	1.19	2.29

Table 3: Effect of weighted and unweighted binary cross entropy loss.

5.3 Different Learning rate

The learning rate is a tuning parameter in an optimization algorithm that determines the step size at each iteration while moving toward a minimum of a loss function. In setting a learning rate, there is a trade-off between the rate of convergence and overshooting.

Learning rate	0.001	0.0001	0.00001
Accuracy(%)	99.99	99.95	99.92
Recall(%)	1.51	68.05	68.67
F1(%)	2.87	23.21	16.30
Precision(%)	0.49	9.53	6.35
RMSE	0.0120	0.0227	0.0284

Table 4: Effect of different learning rate.

5.4 Model Comparison

After choosing the best parameter, we now need to test all the study areas. We trained several individual models with respect to different area and also trained one model for the whole area.

5.4.1 Egypt

Table 5 shows the result of models training with all data and only data near Egypt. The Accuracy and RMSE of model training only from data of Egypt are significantly better than model training from all data.

Training Data	All Data	Egypt
Accuracy(%)	72.63	88.13
Recall(%)	77.07	31.60
F1(%)	39.51	38.19
Precision(%)	23.14	23.19
RMSE	0.6356	0.5383

Table 5: Effect of different training dataset for Egypt.

Figure 7shows the result from the two different models and the ground truth and satellite image of the test area. Figure(a)(b)(c) shows the building density by color. The darker the red is, the more buildings the pixel contains. From Figure(a), the model training with all data tends to predict more area as contain buildings, and tend to predict all pixels with similar building density value. From Figure(b), the model trained with only Egypt data has a better estimation.

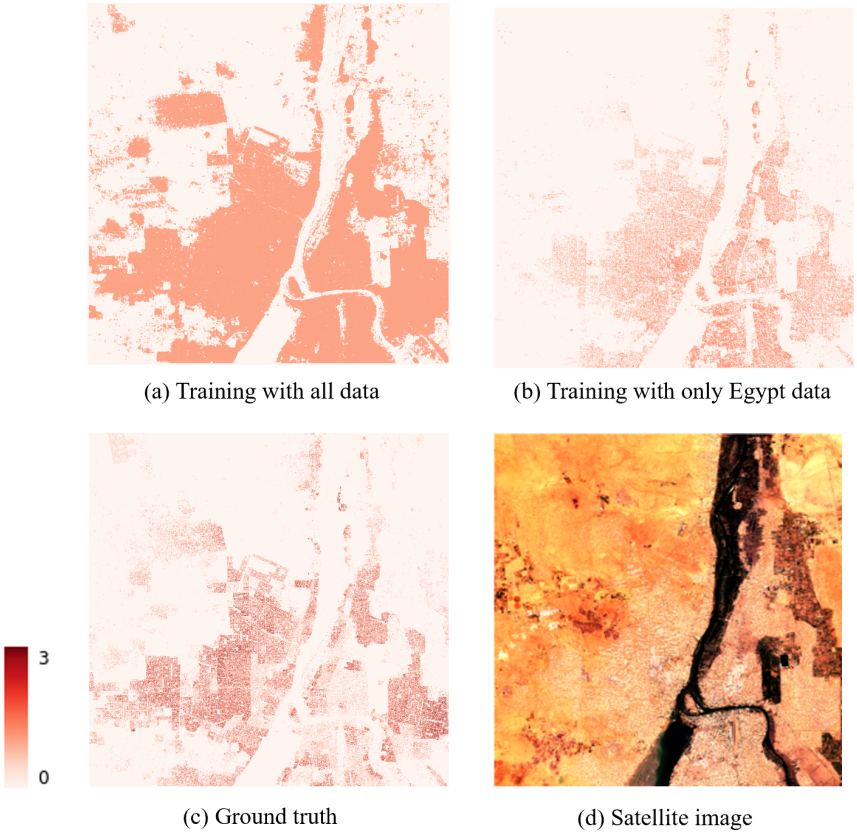

(a) Training with all data (b) Training with only Egypt data

(c) Ground truth (d) Satellite image

Figure 7: Result of Egypt

5.4.2 North Kivu

Table 6 shows the result of the two different models for North Kivu. The model training on only data from Kivu performs better. Figure 8shows the result from the two different models and the

Training Data	All Data	North Kivu
Accuracy(%)	99.95	99.95
Recall(%)	68.05	76.66
F1(%)	23.21	24.28
Precision(%)	9.53	11.06
RMSE	0.02276	0.0233

Table 6: Effect of different training dataset for North Kivu.

ground truth and satellite image of the test area. Figure(a)(b)(c) shows the building density by color. The darker the red is, the more buildings the pixel contains.

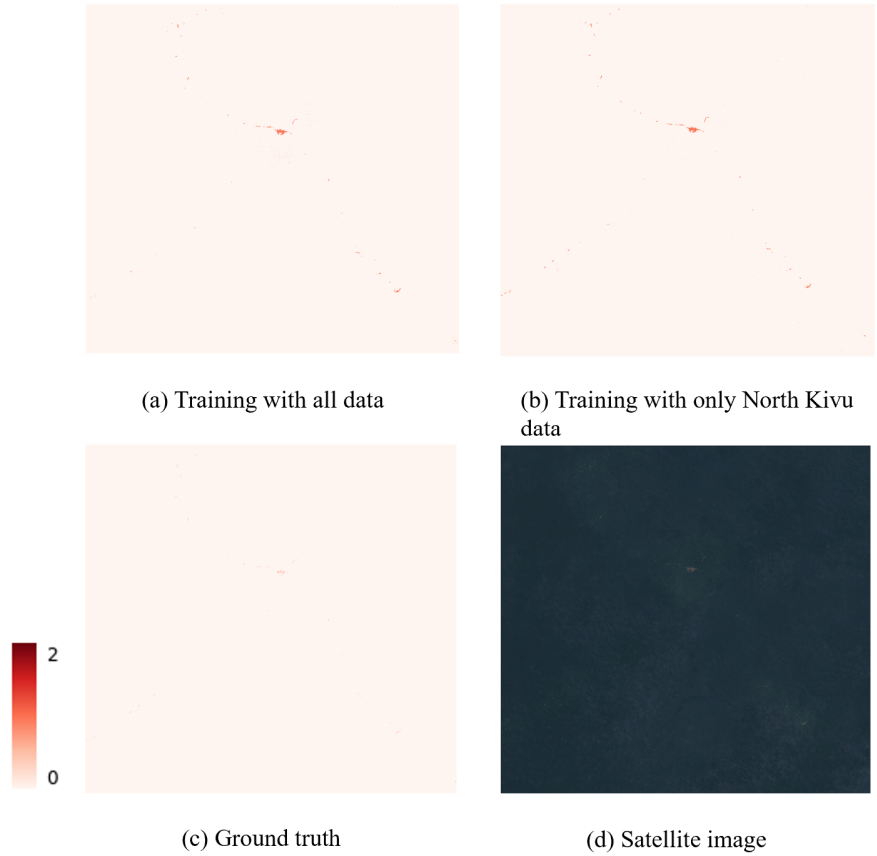

(a) Training with all data

(b) Training with only North Kivu data

(c) Ground truth

(d) Satellite image

Figure 8: Result of North Kivu

5.4.3 Maiduguri and Mokolo

Since Maiduguri and Mokolo are very close to each other, they share similar building style and landscape. We combine this two regions together as one study region.

Table 7 shows the result of the two different models. The model training on only data from Maiduguri and Mokolo performs better. Figure 9shows the result from the two different models and the ground truth and satellite image of the test area. Figure(a)(b)(c) shows the building

Training Data	All Data	Maiduguri and Mokolo
Accuracy(%)	93.92	94.97
Recall(%)	88.50	88.39
F1(%)	25.54	29.28
Precision(%)	13.34	15.64
RMSE	0.2541	0.2309

Table 7: Effect of different training dataset for Maiduguri and Mokolo.

density by color. The darker the red is, the more buildings the pixel contains.

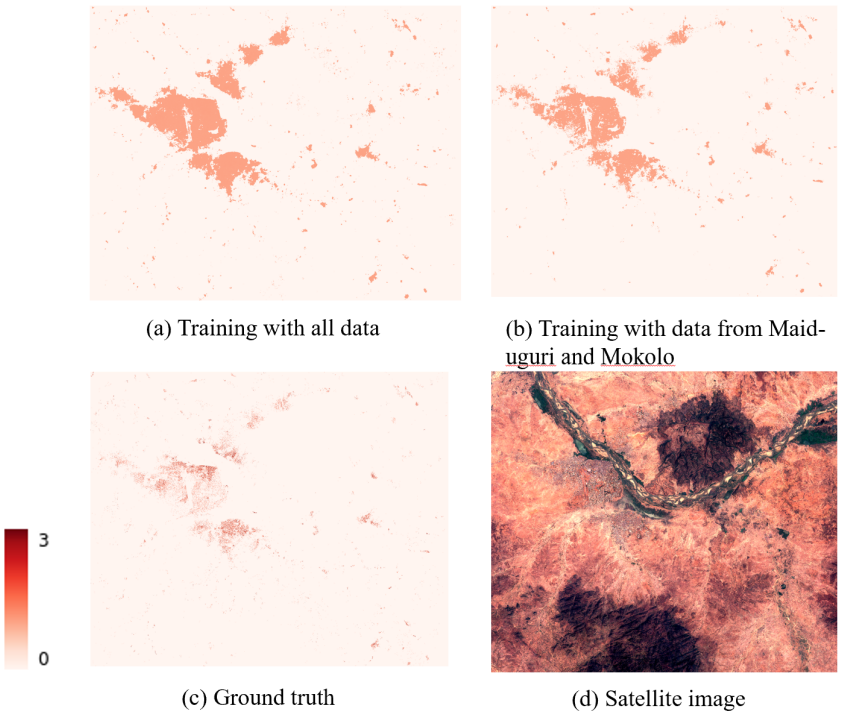

(a) Training with all data

(b) Training with data from Maiduguri and Mokolo

(c) Ground truth

(d) Satellite image

Figure 9: Result of Maiduguri and Mokolo

5.4.4 Montepuez

Table 8 shows the result of the two different models for Montepuez. The model training on only data from Montepuez performs better. Figure 10shows the result from the two different models

Training Data	All Data	Montepuez
Accuracy(%)	91.80	96.07
Recall(%)	75.78	71.86
F1(%)	8.76	15.96
Precision(%)	3.65	6.60
RMSE	0.2872	0.2003

Table 8: Effect of different training dataset for Montepuez.

and the ground truth and satellite image of the test area. Figure(a)(b)(c) shows the building

density by color. The darker the red is, the more buildings the pixel contains.

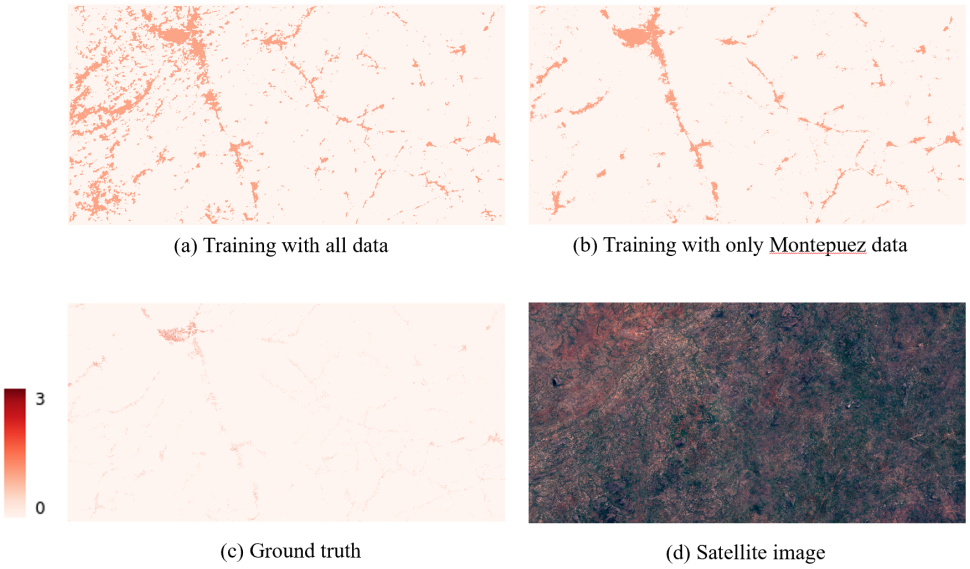

(a) Training with all data (b) Training with only Montepuez data

(c) Ground truth (d) Satellite image

Figure 10: Result of Montepuez

5.4.5 Juba

Table 9 shows the result of the two different models for Juba. The model training on only data from Juba performs better. Figure 11shows the result from the two different models and the

Training Data	All Data	Juba
Accuracy(%)	90.18	93.77
Recall(%)	93.89	86.23
F1(%)	16.84	22.67
Precision(%)	8.75	11.40
RMSE	0.3172	0.2521

Table 9: Effect of different training dataset for Juba.

ground truth and satellite image of the test area. Figure(a)(b)(c) shows the building density by color. The darker the red is, the more buildings the pixel contains.

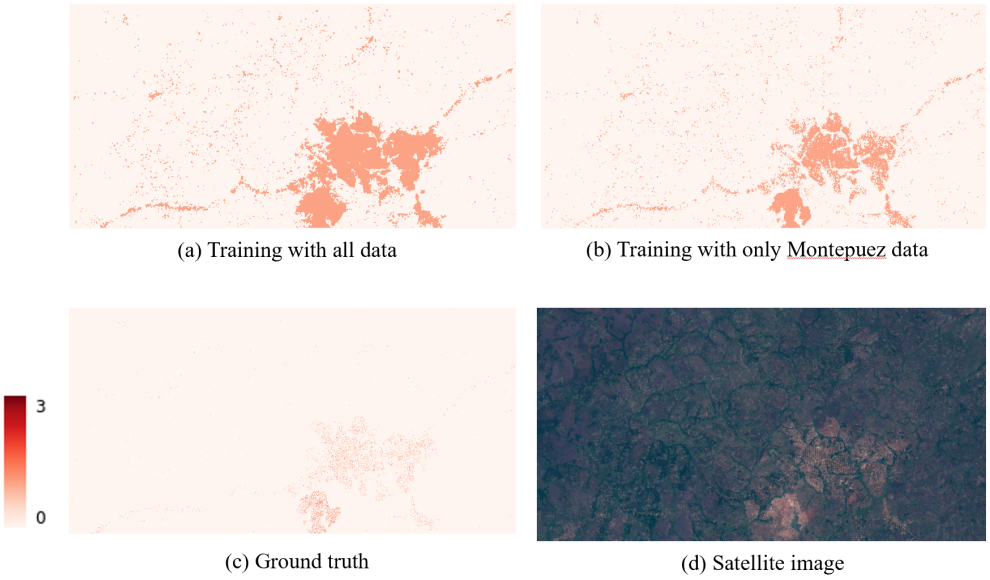

(a) Training with all data

(b) Training with only Montepuez data

(c) Ground truth

(d) Satellite image

Figure 11: Result of Juba

6 Discussion

The advantage of this method is its success in estimating building density from low-resolution satellite data, which none of the previous works have implemented.

The downside is that it tends to over-predict the region containing buildings, which may be due to the fact that the building footprints fell into more than one pixel, but we only chose the pixel where the building polygon's center is located. As a result, they forecast an outcome that differs from the ground truth we created. Also, when the pixel has a larger building number, it can't forecast as correctly.

If we compare the results of the same models performed in different regions, we can observe that the model performed better in jungle areas like North Kivu than in desert areas like Egypt. One of the reasons can be that the training area has more jungle than deserts, leading to the model learning more features about the jungle. This can account for the lower accuracy. Another reason might be the model tends to predict pixels with buildings with a building number of 1, which is also caused by the fact that among all those pixels with a building, the building number of 1 is much more than other building numbers. Since the building density in Egypt is denser, the model failed to predict a higher number. And this can be the account of higher RMSE.

7 Conclusion

7.1 Summary

The main contribution of this project is that it puts forward a new approach to estimating building density where lack of high-resolution satellite imagery.

In this project, a pipeline was developed to estimate pixelwise building density from low-resolution satellite imagery. Particularly, It fills the gap in estimating building density in Africa. Using a Fully Convolutional Neural Network, the accuracy reaches over 90 percent, and RMSE can be as low as 0.2.

Since the building in Africa is very sparse, only a low ratio of pixels contain building. In other words, the classes are unevenly distributed. The number of pixels without buildings is far more than the number of pixels with buildings, which can lead to the model learning more features from pixels without buildings and predicting almost all the pixels as zero. Therefore, weighted binary cross-entropy loss was used to deal with the unbalanced data. In addition, by combining segmentation and regression tasks, the project success in preventing the model to predict all-zero values.

We can also conclude from the result of the experiment that training individual models for different terrain areas can perform better than training only one model for the whole area, which can due to the significant variance of terrain landscape in Africa and the different buildings.

7.2 Outlook

For future studies, there are still some improvements that can be implemented.

First of all, more training data should be used. Since the Open Buildings dataset has much more data than we only used in this project, the future study can use most of them to train a more generic and robust model. Using more bands of the Sentinel-2 data as input can also help adding features.

Second, different network models can be tried out in the future. In this project we only used U-Net architecture, it is worth trying other networks.

Third, it should be considered that some building footprints might fall into different pixels. There needs to have a better way to generate ground truth.

In addition, the mean building area as well as other geography background about Africa can be considered in the future to make it closer to the stage to estimate population density.

References

[1] A. Sayre, *Africa*, ser. 7 continents. Twenty-First Century Books, 1999. [Online]. Available: https://books.google.ch/books?id=V9ziwQP26uwC

[2] E. Maggiori, Y. Tarabalka, G. Charpiat, and P. Alliez, "Convolutional neural networks for large-scale remote-sensing image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 55, no. 2, pp. 645–657, 2017.

[3] Q. Wu, R. Chen, H. Sun, and Y. Cao, "Urban building density detection using high resolution sar imagery," in *2011 Joint Urban Remote Sensing Event*, 2011, pp. 45–48.

[4] N. T. Süberk and H. F. Ateş, "Deep learning for building density estimation in remotely sensed imagery," in *2019 4th International Conference on Computer Science and Engineering (UBMK)*, 2019, pp. 423–428.

[5] W. Sirko, S. Kashubin, M. Ritter, A. Annkah, Y. S. E. Bouchareb, Y. Dauphin, D. Keysers, M. Neumann, M. Cissé, and J. Quinn, "Continental-scale building detection from high resolution satellite imagery," *ArXiv*, vol. abs/2107.12283, 2021.

[6] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, "Cnn features off-the-shelf: An astounding baseline for recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2014.

[7] Y. Wang, D. Zhang, and G. Dai, "Classification of high resolution satellite images using improved u–net," *International Journal of Applied Mathematics and Computer Science*, vol. 30, no. 3, pp. 399–413, 2020. [Online]. Available: https://doi.org/10.34768/amcs-2020-0030

[8] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," vol. 9351, 10 2015, pp. 234–241.

[9] N. Tomar, "What is unet?" Jan 2021. [Online]. Available: https://medium.com/analytics-vidhya/what-is-unet-157314c87634

[10] B. Yu, H. Liu, J. Wu, Y. Hu, and L. Zhang, "Automated derivation of urban building density information using airborne lidar data and object-based method," *Landscape and Urban Planning*, vol. 98, no. 3, pp. 210–219, 2010, climate Change and Spatial Planning. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S016920461000191X

[11] Y. Zhou, C. Lin, S. Wang, W. Liu, and Y. Tian, "Estimation of building density with the integrated use of gf-1 pms and radarsat-2 data," *Remote Sensing*, vol. 8, no. 11, 2016. [Online]. Available: https://www.mdpi.com/2072-4292/8/11/969

[12] [Online]. Available: https://en.wikipedia.org/wiki/Sentinel-2

[13] N. Gorelick, M. Hancher, M. Dixon, S. Ilyushchenko, D. Thau, and R. Moore, "Google earth engine: Planetary-scale geospatial analysis for everyone," *Remote Sensing of Environment*, vol. 202, pp. 18–27, 2017, big Remotely Sensed Data: tools, applications and experiences. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0034425717302900

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Declaration of originality

The signed declaration of originality is a component of every semester paper, Bachelor's thesis, Master's thesis and any other degree paper undertaken during the course of studies, including the respective electronic versions.

Lecturers may also require a declaration of originality for other written papers compiled for their courses.

I hereby confirm that I am the sole author of the written work here enclosed and that I have compiled it in my own words. Parts excepted are corrections of form and content by the supervisor.

Title of work (in block letters):

Building Density Estimation from Low Resolution Satellite Data

Authored by (in block letters):

For papers written by groups the names of all authors are required.

Name(s):	**First name(s):**
Rushan	Wang

With my signature I confirm that

- I have committed none of the forms of plagiarism described in the 'Citation etiquette' information sheet.
- I have documented all methods, data and processes truthfully.
- I have not manipulated any data.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for plagiarism.

Place, date	**Signature(s)**
Zurich, 9,6,2022	Rushan Wang

For papers written by groups the names of all authors are required. Their signatures collectively guarantee the entire content of the written paper.