Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

IGP
Institut für Geodäsie und Photogrammetrie

Leonard Haas
20-926-044

Mapping Settlement Structures with the help of Crowd-source and Open-source Labels for Humanitarian Purposes with Deep Learning

Bachelor Thesis

Bachelor's Degree Programme in Geospatial Engineering
Swiss Federal Institute of Technology (ETH) Zurich

Supervision

Prof. Dr. Schindler, Nando Metzger, Olivier Dietrich
June 16, 2023

Abstract

In this thesis, the two building data sets Microsoft Building Footprints and Google Open Buildings are compared. For one, a direct comparison between the two is made. Furthermore, it is tested which data set can be better used as training labels for building prediction models. The direct comparison shows that both data sets have a lot of inaccuracies. But the building areas of Microsoft Building Footprints are more accurate. Additionally, Microsoft's building counts are better. As training labels, Microsoft's data only did slightly better than Googles. The resulting building prediction models catch the general form of most buildings. Some buildings and holes aren't detected by the model, the limited resolution being a key issue. Overall, using Microsoft Building Labels over Google Open Buildings for building detection is recommended. However, both options can still be tested with different comparison approaches.

Acknowledgements

I would like to thank Prof. Dr. Konrad Schindler, Nando Metzger and Olivier Dietrich for supervising this thesis and always supporting me throughout the semester. I would like to thank Nando Metzger for providing some code as a starting point. I would also like to thank Daniel Pfister for proof reading.

Contents

List of Figures

Chapter 1

Introduction

A Comparison between two building data sets, Microsoft Building Footprints and Google Open Buildings is made in this thesis. Their potential as training data for building detection models is tested.

1.1 Motivation

This bachelor thesis is part of a project that aims to provide accurate, high-resolution and up-to-date population mapping for the whole world using satellite imagery. In a first step of the project, context variables including buildings, distance to road and topography are computed from satellite imagery. After that, population counts are predicted from those variables. Those population counts are valuable for the International Committee of the Red Cross when providing humanitarian aid. For example, in vaccine distribution, help of war refugees and response to natural disasters, high resolution population maps are really useful.

1.2 Research Questions

The main goal is to compare the suitability of two existing building data sets, Microsoft Building Footprints and Google Open Buildings, as training labels for the network. On one hand, the data sets themselves are compared to a ground truth. On the other hand, it is compared which one helps creating better models with a deep learning network.
A comparison of the quality of the building predictions for different locations is done as well. It is shown what kinds of buildings the resulting model predicts well.

Additionally, a goal is to compare if images from Sentinel 1 or Sentinel 2 work better for predicting building footprints. It is also determined if a combination of the two works achieves better results than just one.

Chapter 2

Data

2.1 Image Data

Image data is needed as input to train and run the deep learning model. Imagery from the european satellite missions Sentinel 1 and Sentinel 2 is used. Sentinel 1 provides radar images (ESA (2023a)). 2 bands are available in a 10x10m resolution. Only imagery from ascending satellites is used because it is available for all locations. Mixing ascending and descending satellite imagery is avoided, since the data is different due to different geometries. Sentinel 1 has the major advantage that it is usable regardless of weather and daylight. Meanwhile, Sentinel 2 delivers 13 optical bands (ESA (2023b)). However, only four of those bands are available in a 10x10m resolution. Since the resolution is critical when detecting buildings, only those bands are used. They contain the bands 2, 3 and 4, the three RGB channels and band 8, the NIR (near-infrared) band.
For both Sentinel 1 and Sentinel 2, the median of all images for was computed for each season.

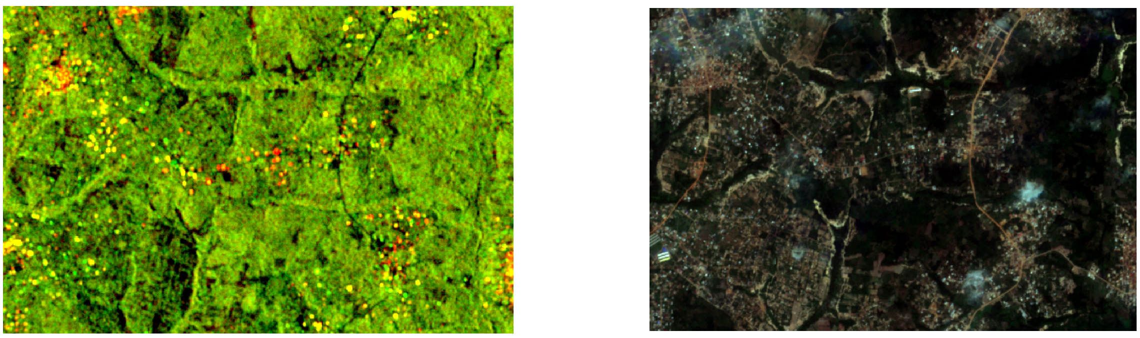

Figure 2.1: Sentinel 1 (left)(ESA (2023a)) and Sentinel 2 (right)(ESA (2023b)) example

2.2 Building Data

Microsoft Building Footprints (Microsoft (2023)) and Google Open Buildings (Google (2023)) are the two data sets that are compared in this thesis. Both cover large areas in a number of regions. This potentially makes them useful as training and validation labels for a building prediction model. Microsoft's Buildings are available for most continents, while Google's data covers Africa and South East Asia (Figure 2.2).

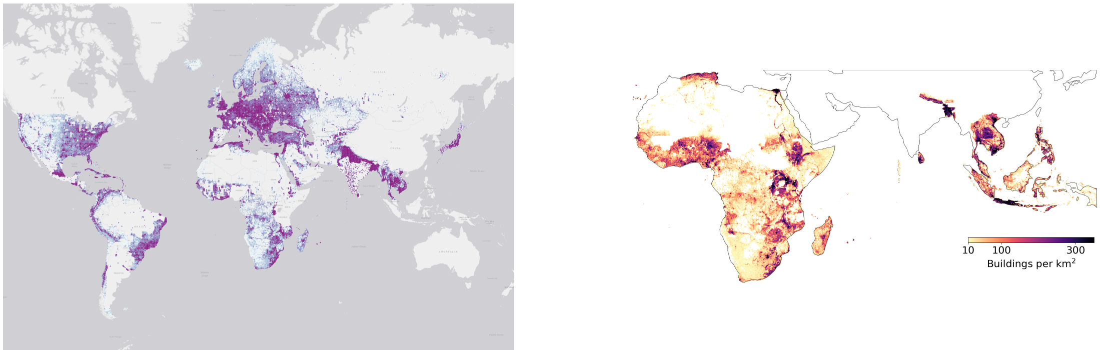

Figure 2.2: Microsoft Building Footprints (left)(Microsoft (2023)) and Google Open Buildings (right)(Google (2023)) coverage

However, both of those building footprints are not always accurate. They were also determined from a deep learning model with satellite imagery as input ((Microsoft (2023)),(Google (2023)). This automated approach results in a lot of errors.

As shown in the example in Figure 2.3, buildings are often mapped in incorrect shapes (middle and right image). Additionally, some containers are mistaken for buildings in both data sets.

Figure 2.3: SpaceNet7 (left)((SpaceNet (2020)), Google Open Buildings (center)(Google (2023)), Microsoft Building Footprint (right)((Microsoft (2023)) example from Senegal

Those two building data sets are downloaded via Google Earth Engine. The Microsoft Building Data was created from image data between 2014 and 2021. No more precise dates are known. The Bangladesh Google tile was created in August 2022, with the other Google tiles dating back to April 2021.

Another source of building footprints is the SpaceNet7 dataset (SpaceNet (2020)). The SpaceNet buildings are the most accurate ones that are found. The buildings were manually labeled which guarantees higher quality. For example in Figure 2.3, it is visible that buildings are accurately mapped by the SpaceNet data set. No other objects are mistaken for buildings either.

Since only 60 4x4km tiles are available, as we can see in Figure 2.4 (SpaceNet (2020)), SpaceNet7 does not provide enough data to train a versatile deep learning model. However SpaceNet7 data can still be used as verification.

Figure 2.4: SpaceNet7 Tiles

2.2.1 Data preprocessing

To compare the Microsoft and Google building data sets as training data for building detection, the building data needs to be rasterized. This step is necessary for all Microsoft and Google training and validation labels. Furthermore, the SpaceNet7 data needs to be rasterized to test the models.
The vector building tiles are rasterized so that they have the same amount of pixels in x and y dimension as the corresponding imagery. That means the building labels have the same resolution of 10x10 meters. All pixels whose center is covered by a building polygon are marked as building area in this rasterization.

2.3 Data Allocation

2.3.1 Direct Comparison

For the direct comparison of Microsoft Building Footprints and Google Open Buildings using SpaceNet7 as footprints, only areas that are covered by all three data sets can be used. This limits the selection to just 8 4x4 km tiles (Figure 2.5). Most of them are located in Africa, with one in Bangladesh. Three tiles near the city of Cairo in Egypt are used.

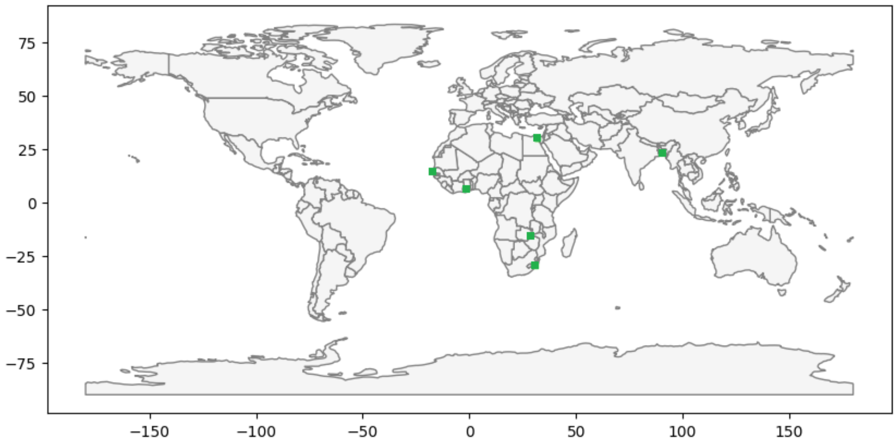

Figure 2.5: Tiles for Direct Comparison and Model Testing

The countries of all tiles are listed here:

Country
Bangladesh
Egypt (3 tiles)
Ghana
Senegal
South Africa
Zambia

Table 2.1: List of Testing Tile Countries

2.3.2 Comparison of Model Predictions

The same 16 square kilometre tiles displayed in Figure 2.5 are used to test the building prediction models that are generated. Other than in the direct comparison, the rasterized building labels are used.

A lot of the training and validation tiles are selected to be nearby the testing tiles. The idea is to train on tiles with similar geography for optimal learning. Different tiles are used to ensure that correlation between training and testing data is not too strong. Some tiles in other African countries are used as well (Figure 2.6).

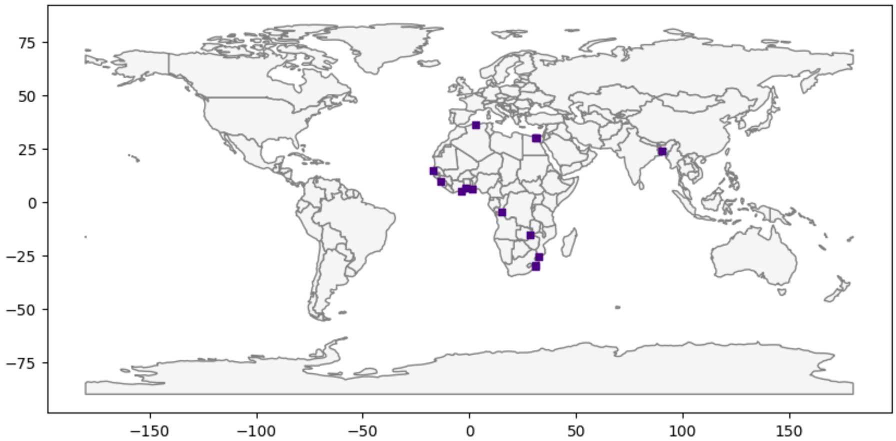

Figure 2.6: Tiles for Training and Validation of Model

Training and Validation tiles from the following countries are used:

Country
Algeria
Bangladesh
Egypt (2 tiles)
Ghana
Guinea
Ivory Coast
Mozambique
South Africa (2 tiles)
Togo
Zambia

Table 2.2: List of Testing Tile Countries

The raster data that is used, consisting of imagery and building labels, is split into training, validation and testing data. For training and validation, data from the same downloaded tiles is used. The left 80 % of each training+validation tile is used for training. This means training+validation tiles contain 1782 x 2228 pixels, which corresponds to around 17.8km x 22.3 km. An area of about 400 square kilometres is covered. The 20 % on the right hand side of each tile is used for validation (Figure 2.7). The validation part makes up 446 x 2228 pixels or 4.5 km x 22.3 km. The validation area for each tile is 100 square kilometers big.

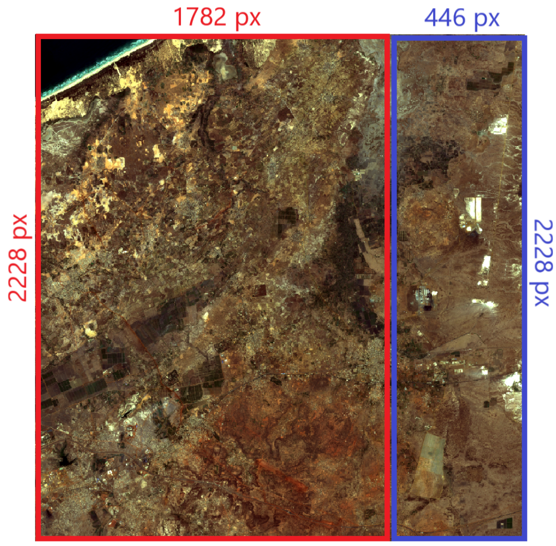

Figure 2.7: Training Area (red) and Validation Area (blue) of Senegal, Sentinel 2 image

Chapter 3

Methods

To compare Microsoft Building Footprints and Google Open Buildings, two approaches are chosen. On one hand, the building data itself is compared to the ground truth. On the other hand, it is tested which dataset produces a better building prediction model when used as training labels. In the end, the results from the two different approaches are compared and discussed.

3.1 Direct Comparison of Building Data

The Microsoft and Google building data sets are directly compared the SpaceNet7 data set as a ground truth. Multiple metrics are computed to assess the quality of both data sets. This also helps to determine the best possible results that can be expected from the deep learning model.

3.1.1 Comparison Metrics

It order to assess and compare the quality of Microsoft and Google Buildings, both are compared to the SpaceNet7 labels individually. The original vector data sets are used to compute multiple metrics. For each test tile, the confusion matrix is computed. It consists of the sizes of the true positive (TP), true negative (TN), false positive (FP) and false negative (FN) areas computed by intersecting the building areas.

TP are the tile sections where both data sets, the tested Microsoft or Google buildings and the SpaceNet7 ground truth, predict building area. FP sections are only marked as building area in the tested data, FN sections only in the ground truth. TN sections aren't marked as building area in either data set, which means the negative prediction of the tested data is correct. The following metrics are computed with those formulas:

$$IoU\ (Intersection\,over\,Union) = \frac{TP}{TP + FP + FN}$$

$$Precision = \frac{TP}{TP + FP}$$

$$Accuracy = \frac{TP + FN}{TP + FP + TN + FN}$$

$$Recall = \frac{TP}{TP + FN}$$

$$F1 - Score = \frac{2}{\left(\frac{1}{\text{Precision}} + \frac{1}{\text{Recall}}\right)}$$

Each metric is then compared for Microsoft Building Footprints and Google Open Buildings.

3.2 Comparison of Model Predictions

In the second approach, both Microsoft and Google are used as building labels when training a deep learning model. It is tested which building labels lead to the best model. Imagery is put into the model as input and trained with the building labels. A U-Net is used as the deep learning network. For testing, the SpaceNet7 buildings are once again used as a ground truth. The SpaceNet data is used as building labels in testing to examine the quality of different generated models. Evaluation metrics are computed and Microsoft and Google models are compared.

3.2.1 Neural Network

The neural network is used for a random 128x128 pixel patch each time. A random season of the imagery is used.
The U-Net is used as the architecture of the deep learning network (O'Sullivan (2023)). It is a commonly used convolutional neural network for image segmentation. Image Segmentation aims to assign a class to each pixel of an image. In our case, the classes are building area and non building area. A U-Net contains an encoder and a decoder part. First, the input image goes through the encoder part. The image is downsampled here. It runs through multiple convolutional and pooling layers. In convolution, a kernel is applied to every pixel. The pooling layers downsample the image. Afterwards, the image runs through the decoder part. The image is upsampled until the dimensions are the same as in the beginning. Deconvolution layers increase the image dimensions. Skip connections between encoder and decoder layers enable the flow of both high-level and low-level feature information to improve segmentation performance.

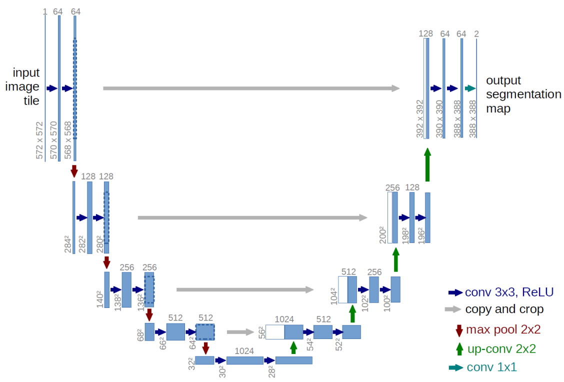

Figure 3.1: U-Net Architecture (Freiburg (2023))

The output of the model is compared to the building labels. The Binary Cross-Entropy (BCE) Loss is computed to measure the measure the dissimilarity of the predictions and the building labels (Godoy (2018)).
The BCE-Loss is computed as follows:

$$\text{Average BCE Loss} = \frac{1}{N}\sum_{i=1}^{N} - \left[y_i \cdot \log(p_i) + (1 - y_i) \cdot \log(1 - p_i)\right]$$

The loss is computed for each pixel i. The p is the given pseudoprobability of the model prediction that the pixel contains a building. The y is a binary value. The value is 1 if the pixel actually contains building area according to the building label.
BCE Loss is minimized in order to improve the model.

3.2.2 Training Arguments

The following data settings are used:

Argument	**Entry**
Number of workers	4
Batch Size	20
Crop Size	128
Scaling	8

Table 3.1: Data Settings

The used hyper parameters mostly remain the same for the different models. The following training settings are used:

Argument	**Entry**
Downsampling Steps	3
Number of Epochs	250
Optimizer	Adam
Learning Rate	0.0001
Momentum	0.9
W-decay	1e-5
Learning Rate Scheduler	Step
Learning Rate Step	10
Learning Rate Gamma	0.9
Gradient Clip	0.01

Table 3.2: Training Settings

3.2.3 Evaluation Metrics

Multiple evaluation metrics are computed during training, validation and testing. Different to the BCE Loss, they are not used by the U-net to improve the model. They provide additional feedback about the quality of a model. The goal is to compare different models. This way, the influence of using different building labels as training can be detected. Also, the model quality for varying imagery and the success on different testing locations is estimated.
The following metrics are computed:

$$Precision = \frac{TP}{TP + FP}$$

$$Recall = \frac{TP}{TP + FN}$$

$$F1 - Score = \frac{2}{\left(\frac{1}{\text{Precision}} + \frac{1}{\text{Recall}}\right)}$$

The positive or negative values from the predictions aren't binary values, but an estimate of the probability between 0 and 1. In order to compute those metrics, they need to be transformed to binary numbers using a threshold. The threshold specifies above what probability a 1 is assigned to a pixel. The threshold has a big impact on the metrics and can be optimized.
If the average Recall is really low for a model, not enough buildings are being detected leading to a high amount of false negatives. In that case, the threshold needs to be decreased so more pixels are declared building areas in the predictions. Meanwhile, it can also happen that too much building area is predicted. In that case, the average Precision is low due to the false positives. The threshold would have to be increased, only considering pixels where pseudoprobability is higher.
After the threshold is optimized on average, different predictions can be compared using the evaluation metrics. The F1-Score contains both Precision and Recall as weights. That makes it a good measurement for the quality of a prediction. That is why the F1-Score is used to compare training labels, imagery and locations.
Even though some of the same metrics are computed, those metrics can not directly be compared to the ones from the direct building data comparison in chapter 3.1. Here, raster data is used instead of vector data, which might influence the results.

Chapter 4

Results and Discussion

4.1 Direct Comparison

4.1.1 Building Areas

Results

The metrics in Figure 4.1 show the difference between SpaceNet and the other 2 building datasets is quite substantial in general. The Microsoft data is a little closer to SpaceNet, with all metrics being higher.

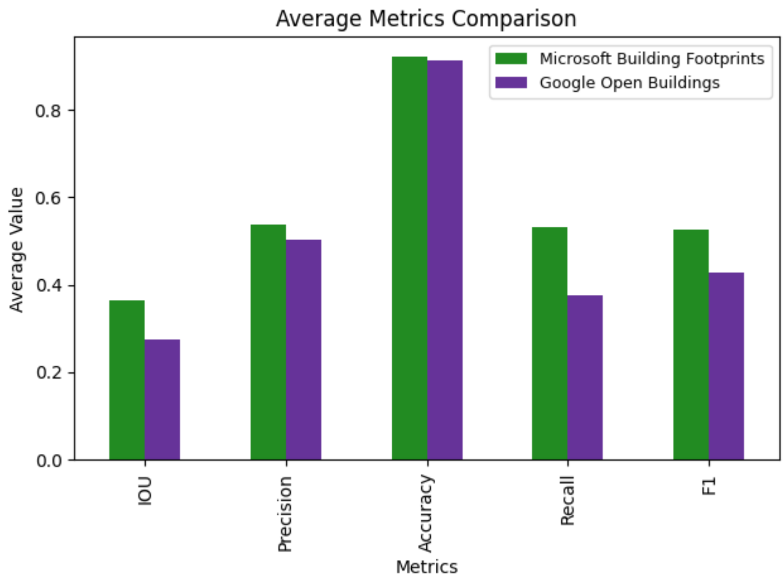

Figure 4.1: Metrics of Microsoft and Google with SpaceNet as ground truth

Discussion

Microsoft's metrics are better than Google's. That suggests that the quality of the Microsoft data set is higher. Especially the recall is much better. However, both data sets are far off from the SpaceNet ground truth. F1-Scores lie at around 0.5. The low Recall values suggest that a lot of buildings area is missed. Meanwhile, the low precision values show that a lot of false building area is found in does data sets.

4.1.2 Building Counts

Results

In terms of building counts per tile, Microsoft and Google have recorded too many buildings compared to the SpaceNet ground truth. However, the Microsoft Building counts are much closer.

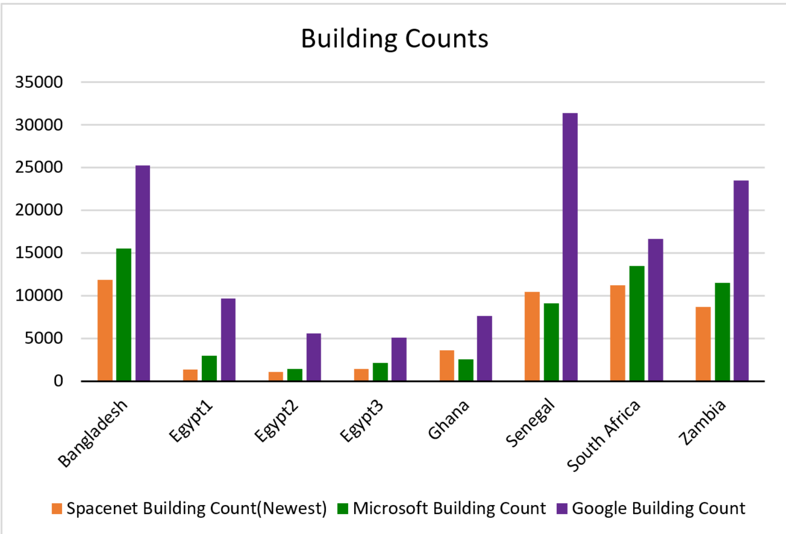

Figure 4.2: Building Counts of Testing Tiles

Discussion

The results suggest that the Microsoft data set is better. Building count is an important metric with the general goal of population density mapping in mind. The amount of buildings is used to calculate inhabitants of an area.

4.2 Comparison of Model Predictions

4.2.1 Comparison of Training Labels

Results

The best model created with Microsoft Building Labels is compared to the best model from Google Building Labels. The test metrics are slightly better for the Microsoft model.

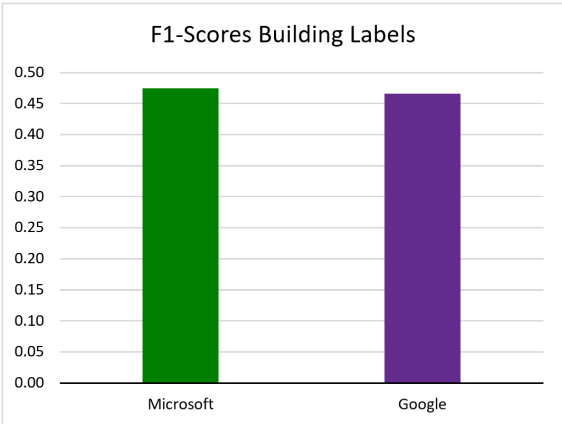

Figure 4.3: F1-Scores Training Labels

Discussion

Microsoft buildings are more similar to SpaceNet than the Google buildings according to section 4.1 . That's why it is expected that the models trained with Microsoft labels produce a better F1-Score. However, the difference in F1-Score is really small at -0.0083. Only 8 testing tiles are used and make for a small sample size. That's why it can't be clearly stated that the models from Microsoft training worked better than the ones from Google. For both data sets, the predictions have a lot of room for improvement.

4.2.2 Best Model

Description

The best model according to the evaluation metrics in the test tiles is obtained when using Microsoft Building Footprints over Google Open Buildings as training labels. Furthermore, only Sentinel 2 imagery was used for the best model.

Output

As seen in Figure 4.4, the overall building structures are captured by the network. However, in some dense areas, the buildings are to thick. Small gaps between buildings go unnoticed and not every building is recognised individually. At the same time, some buildings are missed in areas with few, isolated buildings.

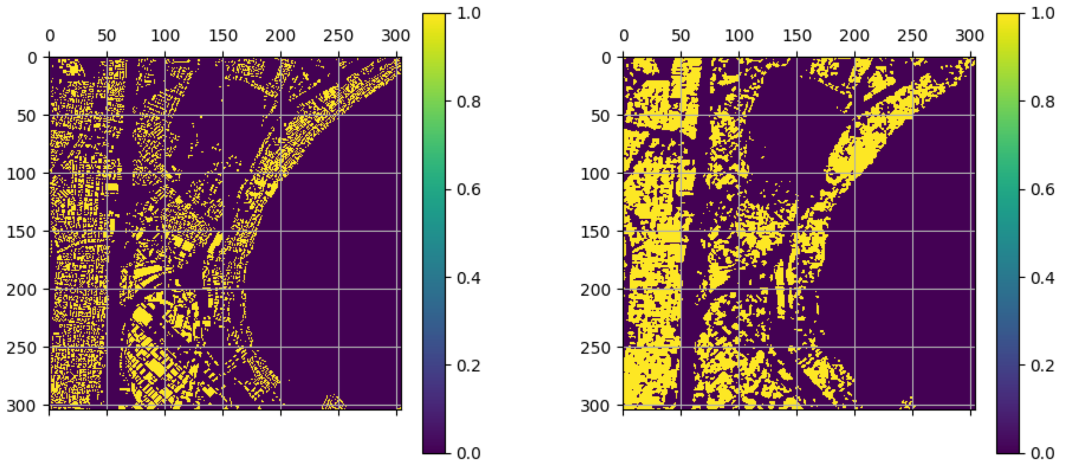

Figure 4.4: SpaceNet labels (left)(SpaceNet (2020)) and Model predictions (right) for Senegal tile

On average, recall is slightly above 0.5 for the best model. That means only half of the building areas are found. Meanwhile, precision lies below 0.5. More then half of the predicted building areas do not actually contain buildings.

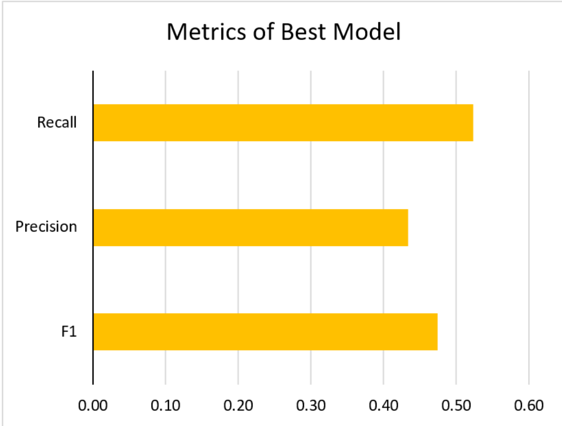

Figure 4.5: Metrics of Best Model

Discussion

The deep learning model often fails to catch small gaps between buildings and small isolated buildings. The limited resolution is probably a main contributor to those issues. A lot of smaller building geometries can't be derived from 10x10 meter satellite images. The rasterization of the vector building data to a 10x10 meter raster is causing issues too. The rasterized building labels are less representative of the actual buildings, which makes training more difficult.

Considering no higher resolution image data is available, those issues are hard to avoid. The resolution sets a limit for the evaluation metrics that can be achieved. For the building rasterization, a different rasterization method might improve the results. Instead of only considering the center of the pixel, pixels that intersect with a building polygon anywhere could be considered building areas. This way, buildings smaller than 10x10 meter would always be considered. At the same time, the models might overpredict the building area this way. Another approach could be to consider what percentage of each pixel is covered by building area and using that in training.

4.2.3 Different Locations with Best Model

Results

When the F1-Scores of the different testing locations are compared, big differences come up. The gap between the highest F1-Score (0.693, Egypt1) and the lowest F1-Score (0.255, Ghana) is immense.

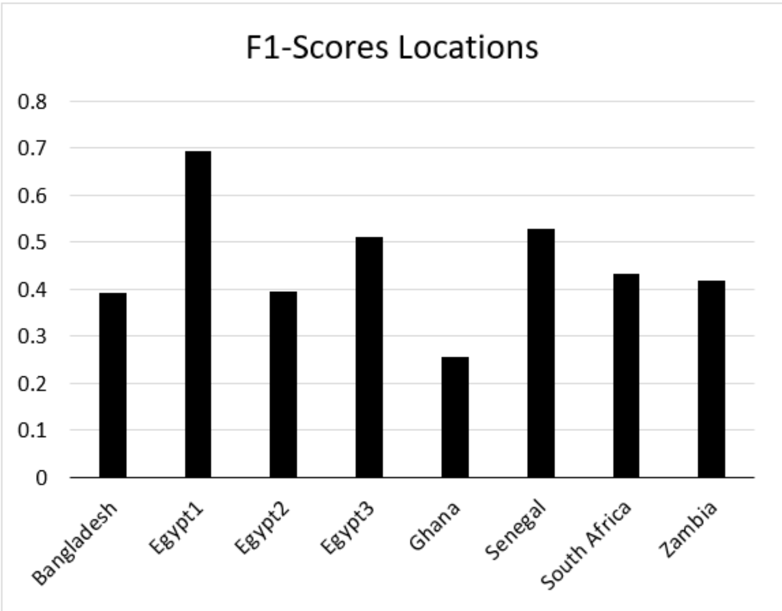

Figure 4.6: Different locations with Best Model

Discussion

A comparison between the location with the highest F1-Score (0.693, Egypt1 tile) and the one with the lowest F1-Score (0.255, Ghana) makes sense. As seen in Figure 4.7, the individual buildings in the Egypt1 tile are much bigger than the ones in the Ghana tile. This explains the major difference in F1-Score. The limited resolution is much less of an issue for the prediction in Egypt1, since all buildings are big enough to be spotted.

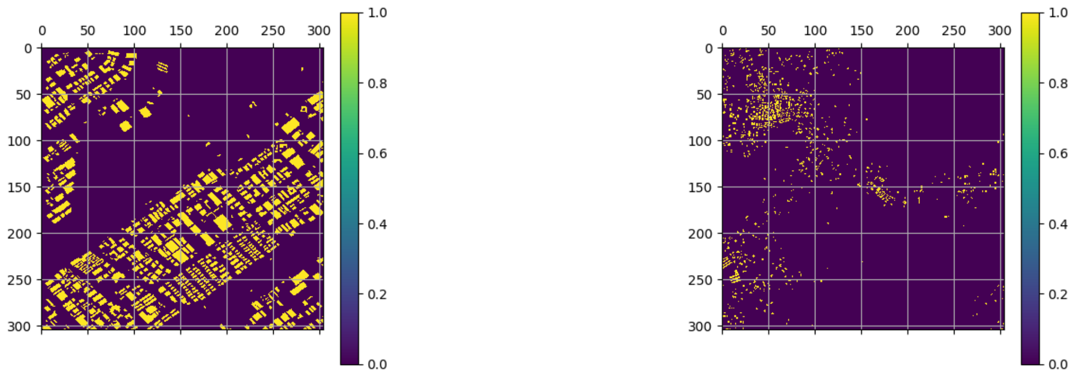

Figure 4.7: Rasterized SpaceNet Labels - Egypt1 (left) and Ghana (right)

For Ghana, the model struggles to predicts the various small buildings. Figure 4.8 shows the model predictions for the two tiles.

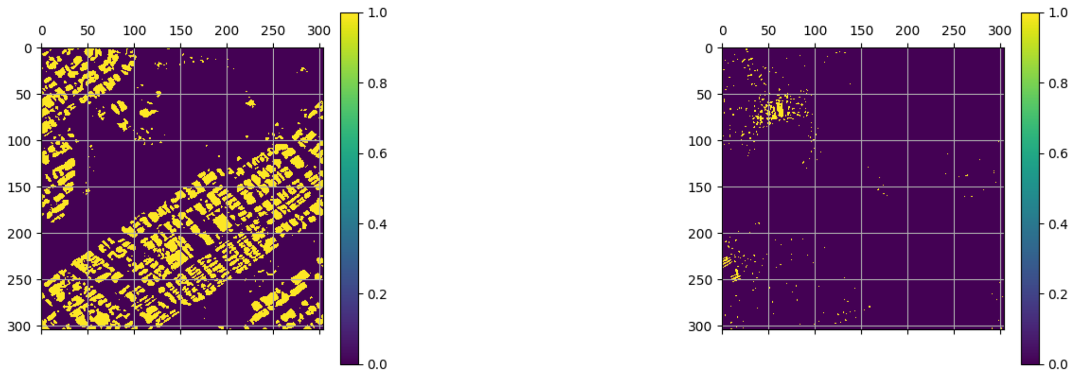

Figure 4.8: Predictions - Egypt1 (left) and Ghana (right)

A closer look at part of the Ghana tile helps explaining the low F1-Score. Looking at the Sentinel 2 image in 10x10 meter resolution (Figure 4.9, left), it is hard to spot the individual buildings even by eye. The model doesn't predict the buildings accurately (4.9, right).
However, this image section shows that the model succeeded at differentiating buildings and some other objects with similar pixel colors. Empty spaces (top left corner) and roads weren't confused for buildings.

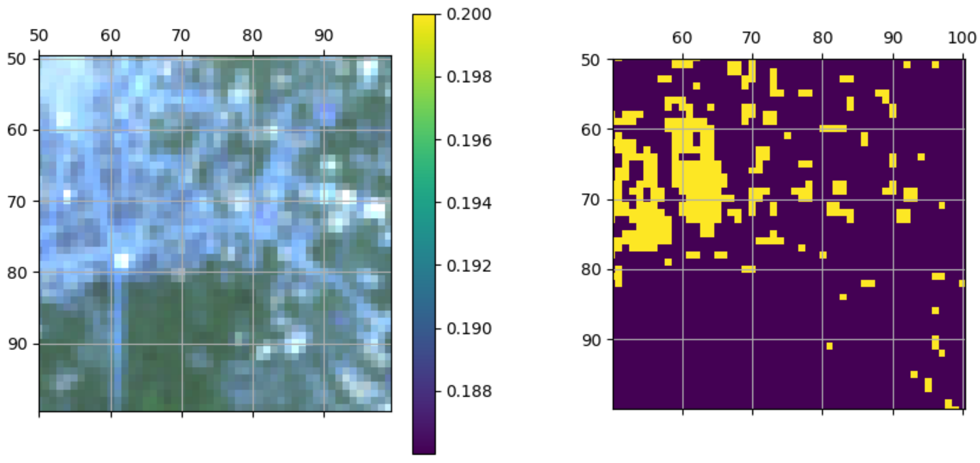

Figure 4.9: Sentinel 2 image (left)((ESA (2023b)) and model prediction (right) of example in Ghana

Figure 4.10 shows the actual ground truth building labels.

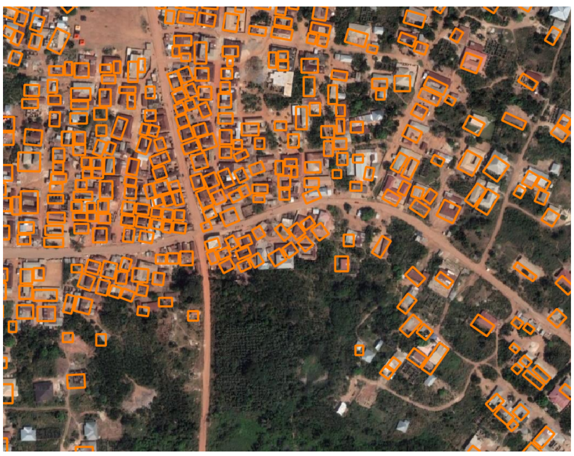

Figure 4.10: SpaceNet7 labels of example in Ghana

4.2.4 Imagery Ablations

Results

Predicting buildings using the radar imagery of Sentinel 1 instead of Sentinel 2 yields the worst results. When the optical Sentinel 2 imagery is used, the F1-Score is much higher. Using both Sentinel 1 and Sentinel 2 combined doesn't improve the results any further.

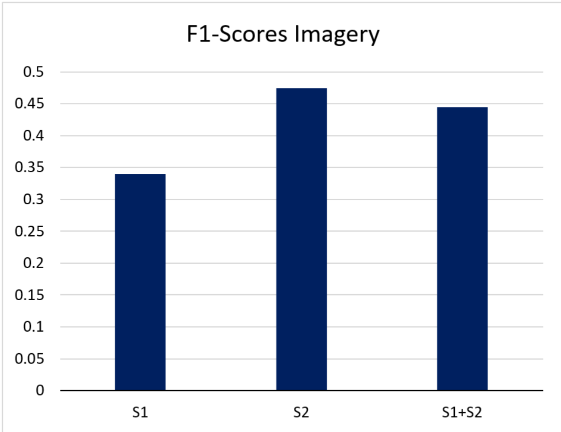

Figure 4.11: F1-Scores Imagery

Discussion

Other research suggests that the more accurate prediction with Sentinel 2 compared to Sentinel 1 is to be expected. However, a fusion of Sentinel 1 and Sentinel 2 should bring a slight improvement over only using Sentinel 2 (Sandhini Putri et al. (2022)). The key take away is that Sentinel 2 works much better than Sentinel 1 as input data. It can't be explained yet why using Sentinel 1 and Sentinel 2 at the same time made the test metrics worse. Different machine learning approaches might justify using both.

Chapter 5

Conclusion

5.1 Main Findings

Using Microsoft Building Footprints over Google Open Buildings as training data is generally recommended. It's helpful that this data set is widely spread. Building Footprints of North America, South America, Europe, Asia and Africa are available and would allow for more diverse training tiles (Figure 2.4).

However, footprints are also missing or unavailable in some regions. In that case, other building labels such as Google Open Buildings can be used.
Another approach might be to combine the two data sets for building labels. The intersection or union area of both could be used. Coordinates that contain buildings in both data sets would have a higher certainty. In a union of the two data sets, less buildings would be completely missed, since they would often appear in at least one data set.

5.2 Recommendations

Adding more training data would help to create a more versatile deep learning model. Training, validation and testing is mostly done for African locations so far and urban and suburban regions are used. If the produced models were tested on different continents and regions, the evaluation metrics would probably be worse. Depending on a region, building architectures and material colors are different. Different challenges are faced when detecting those buildings. In order to predict worldwide building counts accurately, training needs to be done in more diverse locations.
In the direct comparison of Microsoft Building Footprints and Google Open Buildings, some different metrics could be computed. Instead of focusing on the areas,

the amount of buildings that overlap between the tested data set and the SpaceNet7 ground truth could be determined. Less focus would be on the area sizes and shapes, with more focus on the existence of the correct buildings.
It would make sense to ensure that the images and the training labels from the same dates. Building areas changing over time might also be an issue as of now.
For evaluation of the building predictions, some different metrics could be computed. So far, the focus was on the amount of building area that overlaps. It would be interesting to compute building counts and building sizes from the building predictions and comparing those to the ground truth building counts. With regard to the goal of the overall project, predicting population densities, those values need to be accurate. It could be compared if Microsoft Building Footprints or Google Open Buildings produce the better building counts and sizes.
Generally, is it recommended to keep using both Microsoft Building Footprints and Google Open Buildings as building labels when creating new models or calculating new metrics. The difference in performance was quite small so far. With new approaches, Google's data set performing better than Microsoft's is still possible.

Bibliography

E. S. A. ESA, "Sentinel-1," 2023. [Online]. Available: https://sentinels.copernicus.eu/web/sentinel/missions/sentinel-1

——, "Sentinel 2 mission guide," 2023. [Online]. Available: https://sentinel.esa.int/web/sentinel/missions/sentinel-2

U. Freiburg, "U-net: Convolutional networks for biomedical image segmentation," 2023. [Online]. Available: https://lmb.informatik.uni-freiburg.de/people/ronneber/u-net/

D. Godoy, "Understanding binary cross-entropy / log loss: a visual explanation," 2018, https://towardsdatascience.com/understanding-binary-cross-entropy-log-loss-a-visual-explanation-a3ac6025181a.

Google, "Open buildings - a dataset of building footprints to support social good applications." 2023. [Online]. Available: https://sites.research.google/open-buildings/

Microsoft, "Microsoft building footprints," 2023. [Online]. Available: https://github.com/microsoft/GlobalMLBuildingFootprints

C. O'Sullivan, "U-net explained: Understanding its image segmentation architecture," 2023. [Online]. Available: https://towardsdatascience.com/u-net-explained-understanding-its-image-segmentation-architecture-56e4842e313a

A. F. Sandhini Putri, W. Widyatmanti, and D. A. Umarhadi, "Sentinel-1 and sentinel-2 data fusion to distinguish building damage level of the 2018 lombok earthquake," *Remote Sensing Applications: Society and Environment*, vol. 26, p. 100724, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352938522000325

SpaceNet, "Sn7: Multi-temporal urban development challenge," 2020. [Online]. Available: https://spacenet.ai/sn7-challenge/

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Eigenständigkeitserklärung

Die unterzeichnete Eigenständigkeitserklärung ist Bestandteil jeder während des Studiums verfassten Semester-, Bachelor- und Master-Arbeit oder anderen Abschlussarbeit (auch der jeweils elektronischen Version).

Die Dozentinnen und Dozenten können auch für andere bei ihnen verfasste schriftliche Arbeiten eine Eigenständigkeitserklärung verlangen.

Ich bestätige, die vorliegende Arbeit selbständig und in eigenen Worten verfasst zu haben. Davon ausgenommen sind sprachliche und inhaltliche Korrekturvorschläge durch die Betreuer und Betreuerinnen der Arbeit.

Titel der Arbeit (in Druckschrift):

Mapping Settlement Structures with the help of Crowd-source and Open-source Labels for Humanitarian Purposes with Deep Learning

Verfasst von (in Druckschrift):
Bei Gruppenarbeiten sind die Namen aller Verfasserinnen und Verfasser erforderlich.

Name(n):	Vorname(n):
Haas	Leonard

Ich bestätige mit meiner Unterschrift:
- Ich habe keine im Merkblatt „Zitier-Knigge" beschriebene Form des Plagiats begangen.
- Ich habe alle Methoden, Daten und Arbeitsabläufe wahrheitsgetreu dokumentiert.
- Ich habe keine Daten manipuliert.
- Ich habe alle Personen erwähnt, welche die Arbeit wesentlich unterstützt haben.

Ich nehme zur Kenntnis, dass die Arbeit mit elektronischen Hilfsmitteln auf Plagiate überprüft werden kann.

Ort, Datum	Unterschrift(en)
Zürich, 16.06.23	L Haas

Bei Gruppenarbeiten sind die Namen aller Verfasserinnen und Verfasser erforderlich. Durch die Unterschriften bürgen sie gemeinsam für den gesamten Inhalt dieser schriftlichen Arbeit.