

# Mapping the world population one building at a time

Tobias G. Tiecke<sup>1†</sup>, Xianming Liu<sup>1</sup>, Amy Zhang<sup>1</sup>, Andreas Gros<sup>1</sup>,  
Nan Li<sup>1</sup>, Gregory Yetman<sup>2</sup>, Talip Kilic<sup>3</sup>, Siobhan Murray<sup>4</sup>,  
Brian Blankespoor<sup>4</sup>, Espen B. Prydz<sup>3</sup>, Hai-Anh H. Dang<sup>4</sup>

<sup>1</sup>Facebook Inc., 1 Hacker Way, Menlo Park, CA 94025, USA

<sup>2</sup>Center for International Earth Science Information Network (CIESIN)  
Columbia University, Palisades, NY 10964, USA

<sup>3</sup>Development Data Group, The World Bank, Rome, Italy

<sup>4</sup>Development Data Group, The World Bank, Washington DC, USA

<sup>†</sup>To whom correspondence should be addressed. Email: [ttiecke@fb.com](mailto:ttiecke@fb.com)

**High resolution datasets of population density which accurately map sparsely-distributed human populations do not exist at a global scale. Typically, population data is obtained using censuses and statistical modeling. More recently, methods using remotely-sensed data have emerged, capable of effectively identifying urbanized areas. Obtaining high accuracy in estimation of population distribution in rural areas remains a very challenging task due to the simultaneous requirements of sufficient sensitivity and resolution to detect very sparse populations through remote sensing as well as reliable performance at a global scale. Here, we present a computer vision method based on machine learning to create population maps from satellite imagery at a global scale, with a spatial sensitivity corresponding to individual buildings and suitable for global deployment. By combining this settlement data with census data, we create population maps with  $\sim 30$  meter resolution for 18 countries. We validate our**

**method, and find that the building identification has an average precision and recall of 0.95 and 0.91, respectively and that the population estimates have a standard error of a factor  $\sim 2$  or less. Based on our data, we analyze 29 percent of the world population, and show that 99 percent lives within 36 km of the nearest urban cluster. The resulting high-resolution population datasets have applications in infrastructure planning, vaccination campaign planning, disaster response efforts and risk analysis such as high accuracy flood risk analysis.**

Accurate information on global population distribution is crucial to many disciplines. A population and housing census is the traditional tool for deriving small-area detailed statistics on population and its spatial distribution[1]. However, censuses are time-consuming, and the spatial resolution is naturally set by the census enumeration areas (EA), which lack fine-grained information about the aggregation of population. The sizes of the EAs vary by many orders of magnitude from country to country, ranging from hundreds of square meters in urban areas to tens of thousands of square kilometers in low population areas, resulting in an average spatial resolution[2] of a census unit of 33 km at a global scale. Recently, multiple higher resolution maps of human-made built up areas have emerged [15, 16], most notably the Global Human Settlement Layer (GHSL)[3, 4], the Global Urban Footprint (GUF)[5, 6], the WorldPop project[7], Landsat[8] and Missing Maps project[9]. However, none provide a scalable solution with high accuracy in rural areas.

Over the past decade high-resolution (sub-meter) satellite imagery has become widely available, enabling the global collection of recent and cloud-free earth imagery. Additionally, in the past years, the surge in research on computer vision and in particular convolutional neural networks (CNN) have enabled bulk processing of imagery in a rapid manner [10, 11, 12]. The combination of these methods enables the global analysis of high-resolution imagery as a promising

method for detecting individual buildings; combining building estimates with available census data to produce updated and higher resolution population maps; and offering alternative, state-of-the-art population estimates in the absence of census data. Various approaches using machine learning have been demonstrated on small areas [13, 14], yet a method which allows global mapping has remained elusive.

Here, we present a method to generate high-resolution population density maps at a global scale (see figure 1) (referred to as *High Resolution Settlement Layer (HRSL)*). By using various CNN architectures, we identify human-made buildings in high-resolution satellite imagery, for a wide range of terrains, seasons and climates, even under poor image-quality conditions. Under the assumption that buildings act as a proxy for where people live, we obtain population estimates on a country-wide level, with 1x1 arcsecond resolution ( $\sim 30 \times 30$  m at the equator) and sensitivity to individual buildings, enabling accurate studies of population aggregation in rural areas.

To enable global analysis we develop a building detection model that is country-agnostic, fast, and works with easily obtainable binary labeled data. Our pipeline consists of several steps; we extract 64x64 pixel images (*'patches'*) around detected straight lines using a conventional edge detector, which reduces the amount of data for classification by  $\sim 4$  times. A portion of those candidates are sampled across all countries and labeled as training and evaluation data for the CNNs (see Supplemental Information (SI)). The computer vision models are trained on a single machine with four GPUs, whereas the classification runs on Facebook's infrastructure on a CPU cluster. This approach allows us to process  $0.5 \text{ Mkm}^2$  in  $\sim 24$  hours.

We use three different types of CNN (see figure 2): a classification model based on the SegNet [18] architecture in order to perform binary classification (building/no-building); a feedback neural network (FeedbackNet) performing weakly-supervised segmentation of the satellite images enabling us to obtain building footprints [19], and a denoising network which is capable

of improving the quality of the source data by removing high-frequency noise from the satellite imagery (see SI). For all models the tradeoff between good performance at a global scale (larger networks) and the run time of the networks to perform global analysis (smaller networks) determined the limit in performance of the models.

The encoder-decoder style SegNet, is customized to perform the classification at the level of a *patch*. The encoder (a convolutional sub-network) is used to extract abstract information about the input, and the decoder (a deconvolutional sub-network) is trained to upsample the output of the encoder into a spatially meaningful probability map representing the possibility of house existence in the input. The probabilities generated by the decoder are averaged over all spatial locations within the patch to derive the final classification. This yields high accuracy and a reduced false positive rate on a global scale compared to other methods we explored, however, at the expense of loss of spatial details such as the boundaries and shapes of individual buildings.

To recover the building shapes in addition to classification we use a separate CNN to perform image semantic segmentation [20, 21]. Traditional methods, such as region proposal [22] based CNNs [23], are not able to handle small-sized foreground objects. Moreover, these methods require a large amount of pixel-wise labeled training samples following a supervised learning setting. Obtaining such a large volume of labeled data is currently intractable to achieve at a global scale because of it being time-consuming to acquire, as well as it having a large possibility to accumulate errors in supervision. To facilitate a generalized and scalable model, we employ the weakly-supervised learning that takes the abundant and easy-to-get image level categorical supervision (binary labeling) into training, and perform pixel-level prediction during deployment [19]. The methodology is motivated by the feedback mechanism in human cognition [24] and recent advances of computational models in Feedback Neural Networks [25], which deactivates the non-relevant neurons within hidden layers of neural networks and achieve pixel-wise semantic segmentation. Both models are trained on 150,000 binary labeled (build-

ing/no building) *patches*, randomly sampled from all countries and seasons, covering both rural areas and urban areas.

We have ran our model on 18 countries across 3 continents, spanning 29% of the world’s population. For nearly all countries we analyze over 95% of the landmass, of which 1.9% contains a building at a  $1 \times 1$  arcsecond resolution. The coverage is limited by the availability of cloud-free, high-resolution satellite imagery (see SI) and no significant differences in performance between countries is observed, unless the source imagery is of poor quality.

A well characterized accuracy of the dataset is crucial for its further use. We consider the accuracy of the building identification and the accuracy of the population redistribution over buildings separately, with a focus on the former, which is the main topic of our work; the latter will be addressed briefly. Since a global ground truth dataset is not available four independent analyses are performed, each testing different potential errors in our dataset. Statistical errors and systematic errors are treated separately, the former having random origin and the latter originate e.g. from repetitive errors, such as large rocks, boats or mountain ridges (see Figure 3) being falsely classified as buildings. Systematic errors are more challenging to quantify but are of particular concern since they can result in clusters of false positives potentially misinterpreted as a settlement. To quantify statistical errors, we first study the accuracy of the computer vision model by comparing our classification results with a reference (human labeled) dataset. Averaged over 18 countries the SegNet achieves a precision and recall of  $Pr = (TP/(TP + FP)) = 0.947$ ,  $Re = (TP/(TP + FN)) = 0.913$  ( $TP$ ,  $FP$ ,  $FN$  are true positive, false positive and false negative counts respectively) on a global dataset where only 7% of the randomly sampled patches contains a building. This analysis tests the performance of our model, however, it doesn’t capture errors due to clouds, missed straight edge detection or missing source imagery.

To address this, a second analysis is performed comparing the identified buildings to the

GPS coordinates of the households interviewed for the Malawi Third Integrated Household Survey (IHS3). This dataset is obtained by interviewing a nationally-representative sample of households, and thus, is independent of artifacts of remote sensing. The incidence of finding a household is measured within a distance of at most 100m from the nearest populated pixel in our building dataset, thereby accounting for the limited accuracy of the survey location data (see SI). At the national-level, 98.3 percent of the 11,957 IHS3 households coincide with a populated pixel in our building dataset.

In order to quantify the systematic errors we compare our data against two datasets (GUF and GHSL) generated using independent data sources and methodologies. The three datasets are compared at the 1 arc-second resolution of the binary classification to create all possible combinations of absence or presence of a settled area (see figure 3c). The areas where our classification disagrees with both the GUF and GHSL classifications represented as contiguous vectors are visually inspected using the satellite imagery to identify false positive and negative units. Country totals of agreement and disagreement as a proportion of country area are shown in Table 1 (see SI for additional countries). Four test cases of Malawi, Ghana, Vietnam and Haiti are presented since they cover a wide range of geographies and image quality.

We find that the visually confirmed systematic errors contribute to less than 1% of the total landmass and only a few percent of the country built up area. Additionally, in general areas of systematic false negatives occur more often than areas of systematic false positives. In order to assess the impact of larger clusters of disagreement which could yield more severe systematic areas we consider the  $\sim 500$  largest areas of disagreement. These errors contribute to less than 1% of the built up area for all 4 countries. The GUF and GHSL often miss buildings which are discovered in ours.

Finally, we perform an analysis comparing the three datasets (HRSL, GUF, GHSL) against a single region near Blantyre, Malawi which has been completely manually labeled, at a reso-

lution of 2x2 arcsecond to allow for misalignment between the datasets (see figure 3c)[9]. The northern urban area (above the upper dotted line) is captured fairly well in all three datasets expressed in the recall values of 0.83/0.82/0.99 for the GHSL/GUF and HRSL respectively. However, for the southern rural area (below the lower dotted line) the recall values are 0.04/0.06/0.84 for the GHSL/GUF and HRSL respectively, demonstrating the superior accuracy of our dataset in rural areas. In conclusion, all errors are estimated to be a few percent or less, indicating a high degree of confidence in the building dataset. Additionally, our results are superior over existing datasets in rural areas.

After assessing the accuracy of the buildings, we turn to the population redistribution over the buildings. Population estimates are obtained using a minimally modeled approach of proportional allocation [26, 27]. We take two approaches, first, the population is distributed equally to settled areas in the binary classification that fall inside census units and second, the population is distributed proportionally to the fraction of built up area within a 1 arcsecond gridcell (see SI for details)[28]. The population allocations of the second method are expected to be more accurate, however, since currently no method exists to validate the population data globally at sub-arcsecond lengthscales we limit our analysis to the first method. The largest uncertainty of the population estimates originates from census data obtained at a too coarse administrative unit. The impact is estimated by comparing the population estimates performed on a coarse administrative unit to known population counts of finer administrative units (see SI). The error (standard deviation) in population estimate is a factor of 2 or less depending on the available country specific census data. Finally, the errors are slightly larger in urban areas (2.5 or less) compared to rural areas. Our data is well suited for a more sophisticated population allocation, which is an active area of research[7], however, this is beyond the scope of this work.

Our results represent a large improvement over the state-of-the art in globally robust building identification via satellite imagery, and thus population estimation, particularly for rural areas.

As mentioned in the introduction, population estimates are a crucial interdisciplinary concern, and a cursory evaluation of our data hints at the nature of insights to come. For example, the improved rural accuracy enables the study the population distribution with respect to the nearest urbanized center. Definitions of urban versus rural vary widely, and are an active research area. Using the definition of Ref. [29] for *urban clusters*; clusters with a population density of at least  $300 / \text{km}^2$  and a minimum population of 5000 we calculate the distribution of the rural population as a function of the distance to the nearest urban gridcell. Figure 4 shows the distribution of the 18 countries we have analyzed, and shows that 90%/95%/99% of this population lives within 7/14/36 km of the nearest urban cluster. This proximity of rural population to urban clusters provides guidelines for telecommunication infrastructure development.

In conclusion, our model achieves high accuracy and is applicable at a global scale. The errors in building identification are estimated to be at the few percent level on a nation wide scale, while maintaining an unprecedented sensitivity to individual buildings. The uncertainty in population estimates originates from the underlying census data and the method of proportional allocation which is a topic of future research.

## **Acknowledgement**

We thank D. Liu and T. Huang for insightful discussions on the denoising and feedback networks, and C. Deuskar and B. Stewart for insightful discussions on urbanization. All satellite

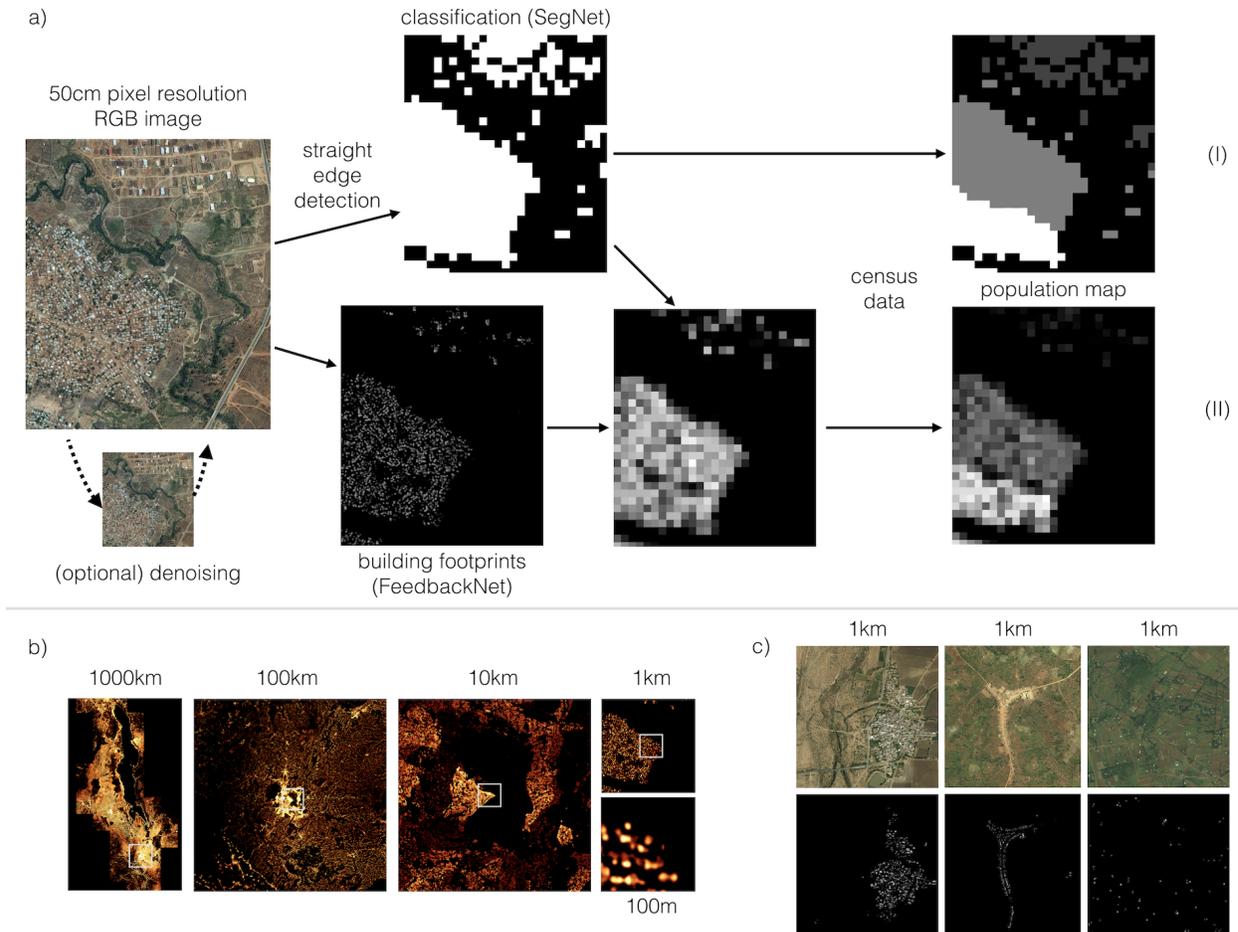


Figure 1: **Method to create population maps from high resolution satellite imagery.** (a) Our method starts with satellite imagery to which a denoising network is applied depending on the image quality. Subsequently, the image is analyzed using straight edge detection and two independent CNNs, SegNet and FeedbackNet, resulting in accurate classification and determination of the building footprint. Two independent approaches are studied to create a built up area map (I) uses classification results and for (II) cascades the CNN outputs where each gridcell corresponds to the fraction of built up area. By combining this with census data a population estimate at 1x1 arcsecond resolution. (b) Examples of the results for Malawi covering 5 orders of magnitude in length scales. (c) Examples of the model output for (left-right) India, Mozambique and Kenya.

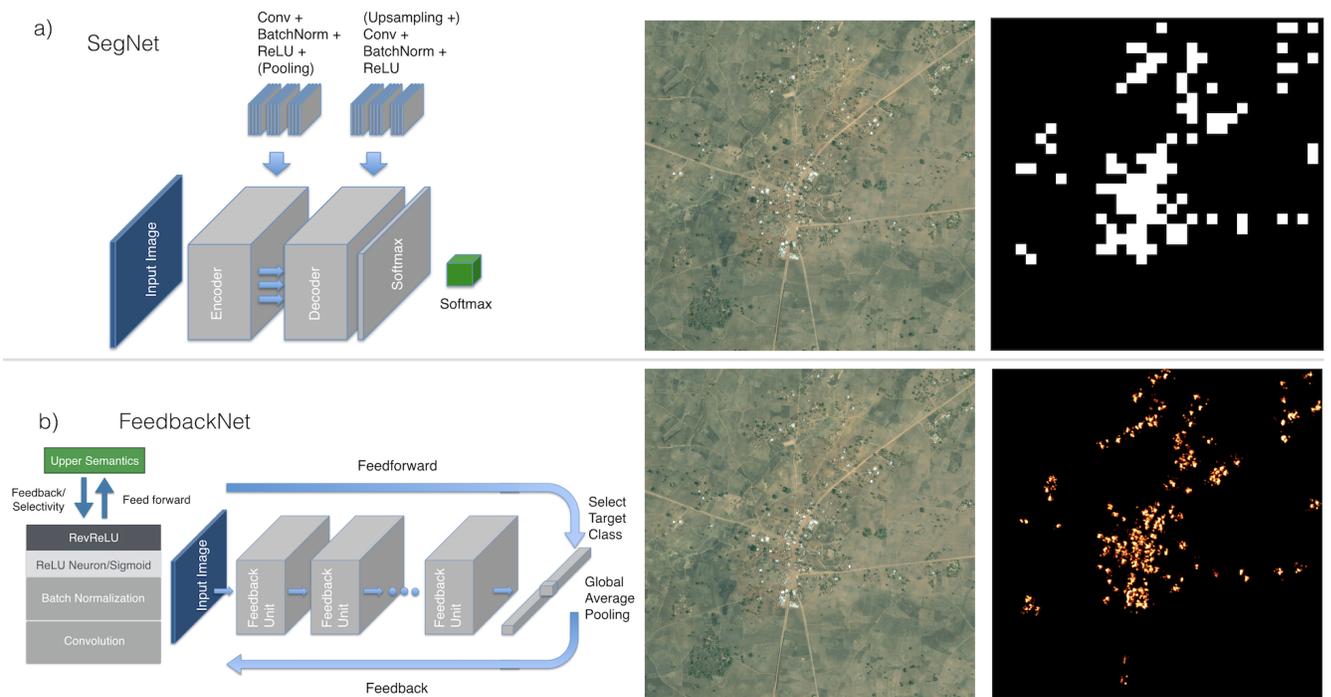


Figure 2: **CNNs for building detection** (a) for classification of *patches* an architecture based on SegNet[18] is used. (b) for image segmentation of the built up area within a *patch* we use a weakly-supervised FeedbackNet[19] architecture on 256x256 pixel images. This network outputs a feature map indicating the building footprint. Since this network is pixel-based it has more false positives than SegNet, which is accounted for by cascading the outputs of the two models.

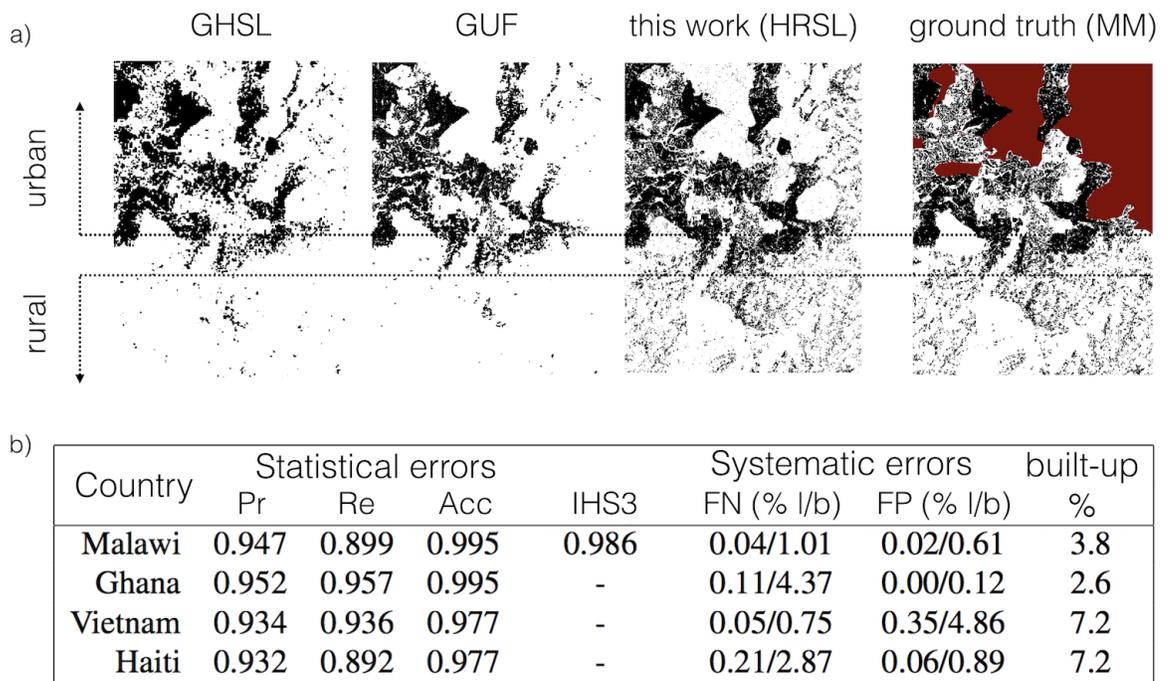


Figure 3: **Validation and error analysis** (a) comparison of GUF, GHSL, HRSL and human labeled ground truth (Missing Maps project, not mapped areas are indicated red). Our dataset accurately captures rural areas in the southern part of this region which are omitted in the GUF and GHSL methods (see text). (b) statistical and systematic errors, the false negative (FN) and false positive (FP) values are shown as percentage of the total landmass (l) and total built up area (b).

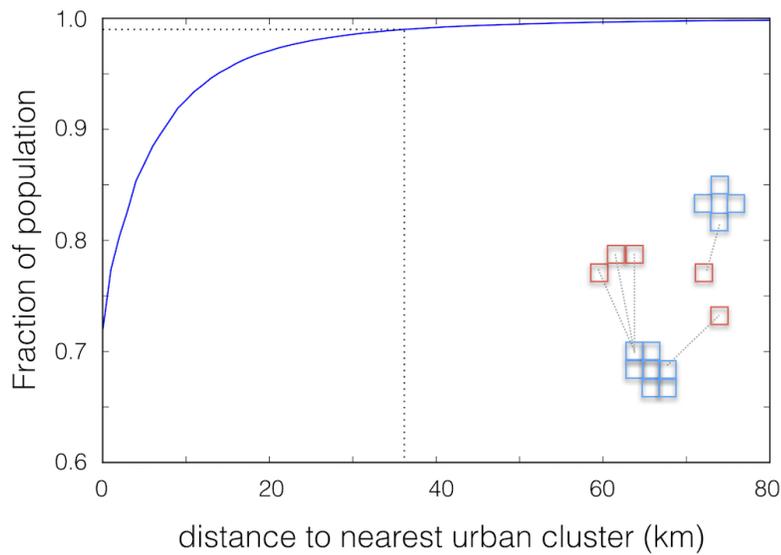


Figure 4: **Distribution of population near urban clusters.** By clustering the population estimates urban clusters are identified and for all population the distance to the nearest urban cluster is determined. Of the studied population (18 countries, covering 29% of the world population) 99% of lives in or within 36 km of an urban cluster. (inset) shows a schematic of the distance calculation; 1x1km gridcells are classified as an urban cluster (blue) or rural (red). The distance is calculated for all population to the nearest urbanized gridcell.

imagery ©2016 DigitalGlobe. Funding for World Bank affiliates has been provided by The World Bank Living Standards Measurement Study ? Integrated Surveys on Agriculture (LSMS-ISA) project, and The World Bank Strategic Research Program Projects: (i) Use of Satellite Data and CDRs for Census-Independent Spatial Distribution of Populations and (ii) Survey Imputation for Improved Poverty and Shared Prosperity Diagnostics. Funding for CIESIN affiliates has been provided by Facebook. The boundary and census data collection were developed with funding from the National Aeronautics and Space Administration under Contract NNG13HQ04C for the Socioeconomic Data and Applications Distributed Active Archive Center (DAAC). Funding for Facebook affiliates was provided by Facebook, Inc.

## References and Notes

- [1] United Nations Department of Economic and Social Affairs (2015)
- [2] Balk, D.L., U. Deichmann, G. Yetman, F. Pozzi, S.I. Hay, and A. Nelson , *Advances in Parasitology* (2006) **62** p. 119-156.
- [3] M. Pesaresi, *et al.* (2015) *European Commission, Joint Research Centre (JRC) [Dataset]* PID: [http://data.europa.eu/89h/jrc-ghsl-ghs\\_built\\_ldsmt\\_globe\\_r2015b](http://data.europa.eu/89h/jrc-ghsl-ghs_built_ldsmt_globe_r2015b)
- [4] M. Pesaresi, *et al.*, *Publications Office of the European Union* UR 27741 EN (2016).
- [5] Esch, T., *et al.*, *ISPRS Journal of Photogrammetry and Remote Sensing* (2017) **134** pp. 30-42. <https://doi.org/10.1016/j.isprsjprs.2017.10.012>
- [6] Esch, T., *et al.*, *IEEE Transactions on Geoscience and Remote Sensing*, (2011) **49**, pp. 1911-1925. <https://doi.org/10.1109/TGRS.2010.2091644>.
- [7] A. J. Tatem, *Sci. Data* **4**, 170004 (2017).

- [8] B. Bhaduri, E. Bright, P. Coleman, M. L. Urban, *GeoJournal* **69**, pp. 103 (2007).
- [9] Red Cross, Missing Maps project, <http://download.geofabrik.de/africa/malawi-latest.osm.pbf>. [Online; accessed 1-Aug-2017].
- [10] A. Krizhevsky, I. Sutskever, G. E. Hinton, *Advances in neural information processing systems* (2012), pp. 1097–1105.
- [11] C. Szegedy, *et al.*, *Proceedings of the IEEE conference on computer vision and pattern recognition* (2015), pp. 1–9.
- [12] K. He, X. Zhang, S. Ren, J. Sun, *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 770–778.
- [13] F. Pacifici, M. Chini, W. J. Emery, *Remote Sensing of Environment* (2009) **113**, pp. 1276
- [14] J. Yuan, (2016) arXiv:1602.06564
- [15] D. Potere, A. Schneider, S. Angel and D. L. Civco, *International Journal of Remote Sensing* (2009) **30**, 24 pp. 6531-6558
- [16] P. Gamba, M. Herold, *CRC Press* (2009)
- [17] Y. Ban, P. Gong, C. Giri, *ISPRS J. Photogramm. Rem. Sens.* (2015) **103**, pp. 1-6.
- [18] V. Badrinarayanan, A. Kendall, R. Cipolla, (2015) arXiv:1511.00561 (2015).
- [19] X. Liu, A. Zhang, T. Tiedeke, A. Gros, T. S. Huang, (2016) arXiv:1612.02766 (2016).
- [20] J. Long, E. Shelhamer, T. Darrell, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2015), pp. 3431–3440.
- [21] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, A. L. Yuille, (2014) arXiv:1412.7062

- [22] J. R. Uijlings, K. E. Van De Sande, T. Gevers, A. W. Smeulders, *International journal of computer vision* (2013) **104**, pp. 154
- [23] S. Ren, K. He, R. Girshick, J. Sun, *Advances in neural information processing systems* (2015), pp. 91–99.
- [24] R. Desimone, J. Duncan, *Annual review of neuroscience* **18**, 193 (1995).
- [25] C. Cao, *et al.*, *Proceedings of the IEEE International Conference on Computer Vision* (2015), pp. 2956–2964.
- [26] A. De Sherbinin, S. Adamo, *United Nations Expert Group Meeting on Strengthening the Demographic Evidence Base for the Post-2015 Development Agenda. New York, UN-DESA.* (2015).
- [27] E. Doxsey-Whitfield, *et al.*, *Papers in Applied Geography* **1**, 226 (2015).
- [28] By combining satellite imagery with census data, there is no association with any information that would lead to re-identification of individuals, and the granularity of the data is not enough for individual identification. Additionally, no Facebook user data has been used in this project.
- [29] Dijkstra, Lewis, H. Poelmann (2014), vol. European Commission Urban and Regional Policy. Working paper 1 (2014).

## **Supplementary Information for** *Mapping the world population one building at a time*

### **1. Source satellite imagery**

#### **1.1. Availability and requirements**

Our models are based on satellite imagery from Digital Globe. The data is largely from the Vivid+ product, augmented with grayscale data in areas with high cloud cover and in very rare cases the Global Basemap RGB data. All data has 50cm resolution, which is required to achieve sensitivity to individual houses. We tested the sensitivity of our approach on the resolution of the imagery by repeating our approach on downsampled imagery, which is subsequently re-sampled by applying the super resolution technique [1] to reconstruct the high resolution imagery. We find that individual houses can still be recognized when downsampling by 3x (1.5m resolution), albeit with reduced accuracy. At 2m resolution and less individual houses are not recognized anymore. At the time of writing, the only commercially available product with spatial resolution of 1.5m or better and (near) global coverage is provided by Digital Globe.

In order to achieve cloud free coverage, the satellite imagery naturally is composed of images taken over a spread of time, typically multiple years per country. Of the data used for the countries listed below, 90% was taken between 2011 and 2015. Further limitations in the performance arise from seasonal differences, however, when sufficient data of different seasons is included in the training dataset this is mitigated.

#### **1.2. Denoising neural network**

Noise in source images is a big concern for a convolutional neural network to generate accurate predictions[2]. Unfortunately, in a mosaic of satellite images a certain fraction of the images are noisy due to the normalization or calibration in pre-processing steps, this leads to a spatial dependency of the noise which can lead to spatially determined systematic errors. To handle this problem considering both efficiency and effectiveness, we implemented an Image Denoising Neural Network (DeepDenoiser), as shown in Figure S1a.

The network is trained end-to-end in a regression setting, with a generator producing random noise applied to original input dynamically. Compared with the state-of-the-art (BM3D[3]), our approach does not require a large search space and can achieve similar performance with at least 10 times efficiency gains in speed. This allows us to run the denoising at large scale. As shown in Figure S1b, the DeepDenoiser improves the data quality significantly compared with the original raw input data; while Figure S1c shows the detection results before and after denoising on Lagos, Nigeria, where the denoising helps find the missed city caused by low data quality.

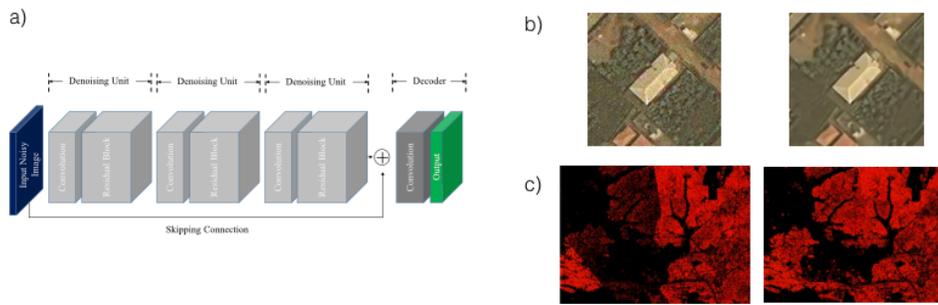


Figure S1. Image denoiser neural network. a) the network architecture. b) and c) show the performance of the denoiser in Lagos, Nigeria. On the left noisy images result in omission of a large area in the west of Lagos, while after applying the denoising network (right) the missed buildings are recovered. Satellite images © 2016 Digital Globe.

## 2. Evaluation of neural networks

### 2.1. Training and testing data labeling

To train and evaluate machine learning models for building detection, we sample 150K 64x64 patches detected from straight-line detector to label. For the training data, we adopt a bootstrapping method to make positives and negatives balanced. While for testing data, we use randomly sampling over spatial location to ensure the testing data is of the identical distribution with the real-world data.

All samples are labeled by at least two different labelers, and the ones with dispute are removed to ensure the data correctness. Even though, we still observe large portion of incorrectness, especially in rural areas, where the imagery data quality is usually bad and density of buildings is low. As a consequence, labelers tend to label “negative” since there are far more negative examples than positives. Accuracy of single human labeler is calculated by using golden dataset, which is around 85% over all labelers. Considering the accuracy of machine learning model at a global scale being over 90%, we can make a conclusion that the machine learning model is superior than human in performing this particular task.

### 2.2. Precision and recall results by country

In Table S1 we list the precision and recall numbers by country, as well as the fraction of the landmass which is covered in our analysis. Coverage less than 100% is due the lack of the availability of cloud-free satellite imagery.

Country	Precision	Recall	Coverage (%)	Urban population (%)	p90 (km)	p95 (km)	p99 (km)
Algeria	0.948	0.913	100	66.5	10.3	17.9	67.7
Burkina Faso	0.940	0.919	100	33.2	26.6	33.2	48.2
Cambodia	0.918	0.941	97.2	51.1	16.4	28.2	48.8
Ghana	0.952	0.957	99.0	56.8	40.6	127.5	171.9
Haiti	0.932	0.892	100	11.6	93.7	111.5	134.0
India	0.993	0.961	99.1	74.3	4.6	8.7	20.5
Indonesia	0.944	0.883	96.8	78.3	8.3	21.3	79.5
Ivory Coast	0.950	0.960	96.1	42.4	53.1	104.0	172.2
Madagascar	0.926	0.862	100	30.0	52.6	67.5	102.3
Malawi	0.947	0.899	100	46.3	15.8	22.4	42.4
Nigeria	0.922	0.871	97.5	63.7	16.7	32.4	80.9
Philippines	0.954	0.901	100	73.7	7.5	23.2	84.9
Rwanda	0.912	0.912	100	85.7	<1.0	<1.0	4.8
South Africa	0.984	0.876	100	71.3	22.5	38.4	88.9
Sri Lanka	0.977	0.933	100	73.7	8.1	41.1	72.9
Thailand	0.931	0.892	99.2	48.4	19.9	27.4	48.2
Uganda	0.989	0.926	95.0	58.3	10.4	16.7	29.6
Vietnam	0.934	0.936	83.0	80.5	2.5	7.2	22.3

Table S1. Precision, recall and coverage by country.

### 3. Allocating population to buildings

Census unit boundaries and associated population estimates from the Gridded Population of the World collection [4], version 4 (GPWv4) are used to distribute population to the binary classification data at a resolution of 1 arc-second, approximately 30x30m at the equator. The census unit boundaries collected for GPWv4 vary in the level of spatial detail and date of estimate; Table S2 shows the administrative level, number of units, average unit size, and census year and for each country. The GPWv4 census data include population estimates for five time periods: 2005, 2010, 2015, and 2010; these estimates are interpolated or extrapolated from at least two census dates for each country[5]. The 2015 estimate is used in the population model to create comparable results across countries.

Country	Administrative Level	Number of Units	Census Year
Ghana	2-Parliamentary Constituency	170	2010
Haiti	4-Section Communale	560	2003
Malawi	3-Enumeration Area	12,645	2008
Vietnam	3-District	688	2009

Table S2 Characteristics of source census data. Data sources are as follows: Ghana: [10,11], Haiti: [12,13]; Malawi: [14]; Vietnam: [15].

Population data are modeled using a minimally modeled approach of proportional allocation [5,6], distributing population equally to settled areas in the binary classification that fall inside census units.<sup>1</sup> We use the minimally modeled approach because it is easily understood, suitable for use as an independent variable in studies that use other physical data, and amenable to additional modeling with additional data. The model is used to create a surface of population counts for each country that sums to match the country population estimate in 2015. The population surfaces represent settled areas with homogeneous population counts within census units. Variability in population distribution are a direct result of settled areas identified in the binary classification and the size of the administrative units. South Africa, with much smaller census units on average, provides a more detailed distribution of population than Ghana.

#### 4. Validation of population redistribution

In gauging the accuracy of the population maps generated by the proportional allocation over the building data as described in Section 3, we first assess the maps at the level of the census enumeration area (EA)<sup>2</sup>. The analysis is conducted for Ghana, Malawi, and Vietnam. First, using the EA shapefiles obtained from the respective national statistical agencies, we sum pixel-specific population counts from the high-resolution population maps at the EA-level. Subsequently, we document the census EA-level correspondence between the population estimates according to the population maps and the latest census-based population counts that are updated with country- and location-specific annual population growth rates that are used by CIESIN for generating the GPWv4.<sup>3,4</sup>

Our analysis is based on Ordinary Least Squares (OLS) regressions of the following form:

$$(1) y_i = \alpha + \beta P_i + \varepsilon_i$$

where  $i$  denotes the census EA;  $y$  denotes the 2015 census-based population projection;  $P$  denotes the population estimate summed across the high-resolution population map pixels at the EA-level, and  $\alpha$  and  $\varepsilon$  denote constant and error terms, respectively. Equation 1 is estimated in the national,

---

<sup>1</sup> A coastal buffer of 100m is used to include settled land areas that fall slightly outside of the coastline represented in the census units.

<sup>2</sup> Each country is divided into enumeration areas by the respective national statistical office, specifically for the purpose of the population and housing census. An EA could encompass one part of a ward/village, a complete ward/village or a collection of wards/villages, and constitutes the highest-resolution at which aggregate individual and household population numbers are produced from a population and housing census (PHC).

<sup>3</sup> The latest PHC was conducted in 2008, 2009 and 2010 in Malawi, Vietnam, and Ghana, respectively.

<sup>4</sup> The administrative unit level at which CIESIN annual population growth rates are available differs by country. The rates are defined at the traditional authority-level (administrative unit level 2 out of 3) in Malawi; the parliamentary constituency-level (administrative unit level 2 out of 4) in Ghana; and the district-level (administrative unit level 3 out of 4) in Vietnam.

rural, and urban domains, separately, in Malawi and Vietnam<sup>5</sup>, and only in the urban domain in Ghana because of the urban-only availability EA-level population counts from the latest census.<sup>6</sup>

Given the cross-country differences in the resolution of the census-based population input data underlying the high-resolution population maps, the selected countries provide an opportunity to explore the implications of such variation on the EA-level correspondence between the census-based population projection and the Facebook population estimate.

In the specific case of Malawi, the census-based population input data for the high-resolution map is already provided by the Malawi National Statistical Office (NSO) to CIESIN at the EA-level, i.e., the highest resolution at which input data could be defined. For this reason, we generate two simulated versions of the population estimates in Malawi – under the artificially-created scenarios of using population input data defined at the traditional authority (TA)-level (administrative unit level 2) and, separately, at the district-level (administrative unit level 1), as opposed to at the EA-level (administrative unit level 3).<sup>7</sup>

The results of the regression analyses are summarized in Table S3. Each estimation of Equation 1 uses the population estimate, and the resolution of the available census population data underlying the high-resolution population map. We report the  $R^2$  value associated with each estimation. We have performed this analysis for allocating the population proportional to the binary classification within 1x1 arcsecond gridcell as well as proportional to the builtup fraction within that gridcell (see main text). We find no significant differences between both methods at the lengthscales of administrative unit as performed in this analysis and therefore restrict ourselves to the binary classification results.

Under the simulated population input data scenarios in Malawi, we see that population input data defined at the traditional authority-level or district-level would have a significant bearing on the accuracy of the population estimate. For example, across the nation as a whole and under the TA-level population input data scenario, the EA-level population estimate would explain 37.5 percent of the variation in the census-based population projection. The comparable statistic under the even coarser district-level population input data scenario would decline to 21.8 percent.

These results provide insight in the errors induced by the of the coarseness of the source population data used to allocate to our buildings data. At the national-level, the reduction of the  $R^2$  values for Malawi correspond to an estimated error in the population data (standard deviation) of a factor  $\sigma_{TA}$

---

<sup>5</sup> In Vietnam, the analysis is restricted to the EAs whose center point latitudes were greater than 10 and less than 17.3. This restriction leads us to work in central Vietnam; in an area in which the built-up area estimates were derived with confidence. The satellite imagery for the rest of the country was otherwise not useable due to perpetual cloud cover. A negligible number of observations were dropped due to the lack of correspondence between the shapefiles and the census location information.

<sup>6</sup> The urban/rural attribute is defined at the EA-level, based on the population and housing census.

<sup>7</sup> To derive the first version, we compute the total percent built-up area in each traditional authority; divide the cell-level percent built-up area by the total traditional authority-level built up area to compute the cell-level allocation ratio; compute the total population estimate in each traditional authority; multiply each cell's allocation ratio by the total traditional authority-level population; and compute the new EA-level population estimate by summing across the new population estimates across the cells in each EA. The second version is computed in an identical fashion, with traditional authority-related parts of the computation above replaced with the computations at the district-level.

$= 1/\sqrt{0.415} = 1.6$  and  $\sigma_D = 1/\sqrt{0.232} = 2.1$  for the TA and district-level population input data respectively. We also observe a slightly worse performance in urban areas compared to rural areas. We attribute this to the expected larger diversity of building heights and usage (residential versus non-residential) in urban areas. We illustrate this difference further by looking at the errors in population allocation as a function of EA size. Figure S2 focuses on Accra Metropolitan Area in Ghana. Panel A depicts the extent of EA-level difference between our population data and independent census-based population projection. We see a clear pattern of over-allocation of population to larger (presumably less dense) EAs, and under-allocation to smaller EAs, likely because our model does not account for building height and therefore it does not properly account for changes in population per structure.

Overall, this analysis shows that the accuracy population estimates is limited by the available population input data and its redistribution rather than our building identification. Future research will focus on how to more accurately distribute population counts to building structures by, for example, classification of the building type as well as third-party geospatial data, including those on agroecology, climate, and land use.

Domain	Country	Level of Administrative Unit for Input Data †	Level of Administrative Unit for Validation Data	Observations	R <sup>2</sup>
<b>National</b>	Vietnam	3	4	3,170	0.508
	Malawi	2	3	12,530	0.415
	Malawi	1	3	12,530	0.232
<b>Rural</b>	Vietnam	3	4	2,486	0.485
	Malawi	2	3	11,059	0.351
	Malawi	1	3	11,059	0.282
<b>Urban</b>	Ghana	2	4	7,279	0.102
	Vietnam	3	4	684	0.432
	Malawi	2	3	1,471	0.454
	Malawi	1	3	1,471	0.144

*Table S3. Summary of regression results. There are 3 levels of administrative units in Malawi, and 4 in each of Vietnam and Ghana, where enumeration area corresponds to the highest level of administrative unit. † In Malawi, the comparisons using CIESIN input data defined at the administrative levels 1 and 2 are relying on simulated population estimates, as described above.*

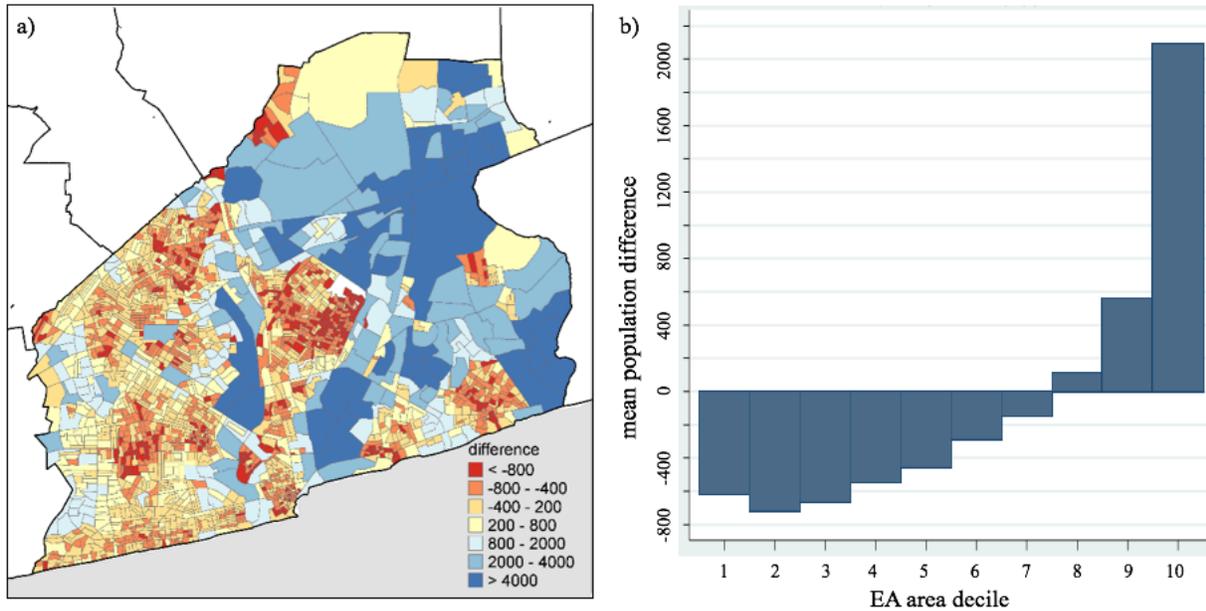


Figure S2. Differences in population attribution (HRSL population estimate vs. EA population) versus EA size in the metropolitan area of Accra, Ghana. a) shows the difference by EA, note that the input census data for Ghana (the Parliamentary Constituency) is here indicated by the black line, i.e. only one unit of the Parliamentary Constituency is redistributed over 7279 Enumeration Areas. b) shows the same data, but grouped by EA area size, indicating a systematic under estimation in small EA areas in urban areas, likely due to the fact that our model does not take building height into account.

## 5. Validation of building localization with household survey data

In this section, we assess the performance of the building identification at a granular level. To do so, we first overlay on the country's high-resolution population map the confidential GPS coordinates for the dwellings of the households interviewed by the Malawi Third Integrated Household Survey (IHS3)<sup>8</sup>. Subsequently, we compute the incidence of a Malawian household appearing in a pixel with no population according to the population map.

On the whole, there are minimal discrepancies between the high-resolution population and the household survey data. The incidence of finding Malawian households at a distance of more than 100m from the nearest populated pixel was only 1.7 percent at the national-level. In computing this statistic, we use a 100-meter tolerance in view of the (i) possible limitations in accuracy of

<sup>8</sup> The Third Integrated Household Survey (IHS3) was conducted by the Malawi National Statistical Office (NSO) from March 2010 to March 2011, with financial and technical support from the World Bank Living Standards Measurement Study – Integrated Surveys on Agriculture (LSMS-ISA) initiative. The IHS3 data are representative for key socioeconomic outcomes at the national-, urban/rural-, regional-, and district-levels. The sample includes 12,271 households, distributed across 768 EAs from 31 districts. The IHS3 anonymized unit-record data and documentation are publicly available on [www.worldbank.org/lsmis](http://www.worldbank.org/lsmis). The unit-record IHS3 data that are publicly available include EA-level GPS coordinates that are computed as an average of household GPS coordinates in each EA and that are randomly off-set to preserve community and respondent confidentiality by 0-2 kilometers in urban areas and 2-5 kilometers in rural areas. The confidential georeferenced dwelling locations that we use for our analyses are not publicly available. The research clearance was provided by the Malawi NSO, who, along with the World Bank Living Standards Measurement Study (LSMS) are the sole gatekeepers of the confidential data, which are not shared with anyone else inside or outside the World Bank, including the co-authors affiliated with Facebook and CIESIN.

GPS position, (ii) GPS measurements having been taken outside of the structures, and (iii) additional imprecision introduced by the gridding process itself (up to 1 gridcell of 30m). The incidence of a Malawian household appearing in a pixel with no population according to the high-resolution population map is 0.4 percent if one defines positive household identification as being within 500 meters of a populated pixel. To better understand the profiles of the “missed” IHS3 households, we reviewed each missed location manually on the Digital Globe imagery, and observed that these households were often 1) frequently occurring in clusters of three or more, and 2) more remote and with smaller structures compared to the rest of the sampled EA. The former may be linked to the mismatch between the dates of the imagery and the dates of the IHS3 fieldwork.

## **6. Comparison with GUF and GHSL**

Two independent classifications of settlements are used for validation of the binary classification: the Global Urban Footprint (GUF) [7] and the Global Human Settlement Layer (GHSL) built-up grid [8]. Both comparison classifications use different data sources and methods to identify settlements. The GUF data represent settlements detected by synthetic aperture with an unsupervised classification scheme applied to backscattering amplitude and texture information from RADAR data—TerraSAR-X/TanDEM-X image products collected in 2011 and 2012<sup>9</sup> at a resolution of 0.4 arc-seconds, approximately 12 metres at the equator [7,9], The GHSL data represent settlements detected by a supervised classification method based on symbolic machine learning from optical imagery—Landsat satellite data collected in 2014 at a resolution of 38 metres [8]. These independent classifications were chosen because of the similar spatial resolution, the proximity in time to the building dataset and the differing data sources.

The three datasets are compared at the resolution of the building dataset —1 arc-second—to create all possible combinations of absence or presence of a settled area. Single-pixel areas of disagreement and areas of cloud in the source DigitalGlobe imagery were removed. The remaining areas where our classification disagrees with both other classifications represented as contiguous vectors are compared with the DigitalGlobe optical imagery to identify false positive and negative units. Country totals of agreement and disagreement as a proportion of country area are shown in Table S4.

---

<sup>9</sup> Data gaps for approximately 7% of the global data are filled with more recent data from 2013 and 2014 (DLR, 2016)

Country	# Assessed	Area Assessed (km <sup>2</sup> )	# Omissions	# Valid	# Undetermined	Area of FN (km <sup>2</sup> )	Percent (by number)	Percent (by area)
Ghana	977	52.5	655	172	150	49.1	67	93.55
Haiti	444	3.6	113	298	33	1.7	25	48.25
Malawi	692	5.7	300	243	149	3.3	43	58.48
Vietnam	1238	69.3	676	297	265	40.0	55	57.67

	# Assessed	Area Assessed (km <sup>2</sup> )	# False Positives	# Valid	# Undetermined	Area of FP (km <sup>2</sup> )	Percent (by number)	Percent (by area)
Ghana	748	50.7	11	660	77	0.31	1.47	0.62
Haiti	3207	234.6	45	3153	9	3.6	1.40	1.53
Malawi	895	56.0	43	848	4	0.63	4.80	1.13
Vietnam	1054	168.3	163	824	67	37.2	15.46	22.12

*Table S4 Country totals of systematic disagreement. The areas correspond to the cumulative areas of disagreement at 1 arcsecond resolution.*

## References:

- [1] Jiwon Kim, Jung Kwon Lee, Kyoung Mu Lee, “Accurate Image Super-Resolution Using Very Deep Convolutional Networks”, arXiv:1511.04587 (2015)
- [2] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, Rob Fergus, “Intriguing properties of neural networks”, arXiv:1312.6199 (2013)
- [3] K. Dabov, A. Foi, V. Katkovnik and K. Egiazarian, “Image Denoising by Sparse 3-D Transform-Domain Collaborative Filtering,” in IEEE Transactions on Image Processing, vol. 16, no. 8, pp. 2080-2095, Aug. 2007
- [4] Center for International Earth Science Information Network – CIESIN – Columbia University. 2016. Gridded Population of the World, Version 4 (GPWv4): Population Count. Palisades, NY: NASA Socioeconomic Data and Applications Center (SEDAC). <http://dx.doi.org/10.7927/H4X63JVC>
- [5] Doxsey-Whitfield, E., K. MacManus, S. Adamo, L. Pistolesi, J. Squires, O. Borkovska and S. Baptista. 2015. Taking Advantage of the Improved Availability of Census Data: A First Look at the Gridded Population of the World, Version 4. Papers in Applied Geography, 1(3):226-234. <http://www.tandfonline.com/doi/full/10.1080/23754931.2015.1014272>
- [6] De Sherbinin, A. and S. Adamo. 2015. CIESIN’s experience in mapping population and poverty. United Nations Expert Group Meeting on Strengthening the Demographic Evidence Base for the Post-2015 Development Agenda. New York, UNDESA. <http://www.un.org/en/development/desa/population/events/pdf/expert/23/ShortNotes/Alex%20de%20Sherbinin%20and%20Susana%20Adamo.pdf>
- [7] DLR, 2016. Global Urban Footprint. Received 11/23/2016.
- [8] Pesaresi, Martino; Ehrlich, Daniele; Florczyk, Aneta J.; Freire, Sergio; Julea, Andreea; Kemper, Thomas; Soille, Pierre; Syrris, Vasileios (2015): GHS built-up grid, derived from Landsat, multitemporal (1975, 1990, 2000, 2014). European Commission, Joint Research Centre (JRC). PID: [http://data.europa.eu/89h/jrc-ghsl-ghs\\_built\\_ldsmt\\_globe\\_r2015b](http://data.europa.eu/89h/jrc-ghsl-ghs_built_ldsmt_globe_r2015b), Downloaded 7/20/2016.
- [9] Esch, T., Marconcini, M., Felbier, A., Roth, A., Heldens, W., Huber, M., Schwinger, M., Taubenböck, H., Müller, A., Dech, S. (2013): Urban Footprint Processor – Fully Automated

Processing Chain Generating Settlement Masks from Global Data of the TanDEM-X Mission. IEEE Geoscience and Remote Sensing Letters, Vol. 10, No. 6, pp. 1617-1621. <http://dx.doi.org/10.1109/LGRS.2013.2272953>.

- [10] Geocommons, <http://geocommons.com/overlays/201941>, Accessed 11/04/2013.
- [11] Ghana Statistical Service, 2010 Population and Housing Census (PHC), Table 2: Population by District and Sex, [http://www.ghanadistricts.com/pdfs/2010\\_pop\\_census\\_districts.pdf](http://www.ghanadistricts.com/pdfs/2010_pop_census_districts.pdf), Accessed 07/12/2013.
- [12] United States Census Bureau, Demobase project using Haiti Population and Housing Census 2003, <https://www.census.gov/population/international/data/mapping/demobase.html>, Accessed 4/28/2014.
- [13] United States Census Bureau, Demobase Haiti, Updated 01/21/2010, [http://www.census.gov/population/international/files/demobase/Haiti\\_sections\\_with\\_2003\\_census\\_data.zip](http://www.census.gov/population/international/files/demobase/Haiti_sections_with_2003_census_data.zip), Accessed 05/06/2014.
- [14] National Statistical Office of Malawi, 2008 Census, Received, 02/18/2014.
- [15] General Statistics Office of Vietnam, The 2009 Vietnam Population and Housing Census: Completed Results, Table 2: Population by Urban/Rural Residence, Sex and District Administration Level, Published 01/04/2009, <http://www.gso.gov.vn/default.aspx?tabid=512&idmid=5&ItemID=10798>, Accessed 04/09/2014.