

POPULATION MAPPING IN INFORMAL SETTLEMENTS WITH HIGH-RESOLUTION SATELLITE IMAGERY AND EQUITABLE GROUND-TRUTH

Konstantin Klemmer
University of Warwick & NYU *

Godwin Yeboah
University of Warwick

João Porto de Albuquerque
University of Warwick &
The Alan Turing Institute

Stephen A. Jarvis
University of Warwick &
The Alan Turing Institute

ABSTRACT

We propose a generalizable framework for the population estimation of dense, informal settlements in low-income urban areas—so called ‘slums’—using high-resolution satellite imagery. Precise population estimates are a crucial factor for efficient resource allocations by government authorities and NGO’s, for instance in medical emergencies. We utilize *equitable* ground-truth data, which is gathered in collaboration with local communities: Through training and community mapping, the local population contributes their unique domain knowledge, while also maintaining agency over their data. This practice allows us to avoid carrying forward potential biases into the modeling pipeline, which might arise from a less rigorous ground-truthing approach. We contextualize our approach in respect to the ongoing discussion within the machine learning community, aiming to make real-world machine learning applications more inclusive, fair and accountable. Because of the resource intensive ground-truth generation process, our training data is limited. We propose a gridded population estimation model, enabling flexible and customizable spatial resolutions. We test our pipeline on three experimental sites in Nigeria, utilizing pre-trained and fine-tune vision networks to overcome data sparsity. Our findings highlight the difficulties of transferring common benchmark models to real-world tasks. We discuss this and propose steps forward.

1 INTRODUCTION

As urbanization accelerates around the globe, cities increasingly struggle to provide affordable and adequate housing for their growing population. This has forced many people seeking opportunities and employment in the cities into informal settlements, also referred to as ‘slums’. These informal settlements are often highly deprived and their populations face various problems, e.g. the lack of healthcare access (van de Vijver et al., 2015) or basic infrastructure (Castells-Quintana, 2017). The poor living conditions in informal settlements in combination with their continuous growth have led the United Nations (UN) to specifically mention “slums” as a challenge to their sustainable development goals (SDG) (United Nations, 2018). It is thus crucial to understand informal settlements and their development. In order to optimize humanitarian aid, precise “slum” detection and population estimates are of particular interest. Deep learning approaches built atop satellite image data have recently become more and more popular for this task (Jean et al., 2016), as they scale well on high-dimensional data and outperform traditional methods of remote sensing (Persello & Stein, 2017). However, these approaches mostly originate from the remote sensing, geographic information science (GIS) and human geography communities (Kuffer et al., 2016). We believe that the machine learning community should assume a more active role tackling these problems, especially since it is at the forefront of computer vision, pattern recognition and scalable deep learning methods. But

*Corresponding author: k.klemmer@warwick.ac.uk

Publication	Data (input)	Data (labels)	Output res.	Valid. res.	Method	Eq. GT
Wardrop et al. 2018	census data soc.-ec. covariates built environment	-	<i>flex. grid</i>	-	Spatial Disaggregation	-
Zong et al. 2019	coarse labels POI data	cellular data	<i>flex. grid</i>	cell-tower	Deep Learning	-
Weber et al. 2018	high-res satellite image microcensus	survey data	<i>flex. grid</i>	-	Image Segmentation	-
Hu et al. 2019	high-res satellite image	census data	<i>flex. grid</i>	village	Deep Learning	-
Xie et al. 2016	nighttime light satellite image	poverty index	<i>flex. grid</i>	household group	Deep Learning	-
Ours	high-res satellite image	building mask	<i>flex. grid</i>	building	Deep Learning	✓

Table 1: Comparing recent population estimation methods: the *Data* columns describe input and validation labels (if present), *Output res.* and *Validation res.* provide the spatial units of the estimation output and the validation labels respectively, *Method* gives the quantitative approach for population estimation and *Eq. GT* indicates, if the validation labels constitute equitable ground-truth.

while image classification is one of the most successful applications of machine learning, it often struggles to generalise to real-life cases such as multi-label images and poor-resolution images. In remote sensing and satellite image classification, we see this effect demonstrated in competitions such as Deep Globe Demir et al. (2018) in which building detection for Las Vegas was 85% whereas detection for Khartoum was only 54%. This also highlights another problem: current methods work better on cities in the Western world, while cities in the Global South, which may benefit especially from these applications, remain neglected.

In this study, we propose a flexible framework for estimating the population of informal settlements. Our proposed method builds on high-resolution satellite imagery and utilizes equitable ground-truth: Building-level shapes and labels are mapped and curated in collaboration with local universities and communities. Through this cooperative approach, the local population takes agency over mapping their own neighborhood; furthermore leveraging their unique domain knowledge to generate high-quality data. While satellite images have previously been used to map informal settlements, our equitable ground-truth data presents a novel approach to population mapping. We further comment on the differences with existing population mapping approaches. Our proposed pipeline is based on a gridded population estimation model with adaptable spatial resolution. Methodologically, we have to overcome a lack of training data, stemming from the expensive ground-truth generation process, and resulting proneness to overfitting. We experiment with fine-tuned versions of pre-trained computer vision algorithms like *ResNet50* and *MobileNetV2* to tackle this issue. However, we find a surprising lack of capability of these powerful networks. We use these preliminary finding as a starting point to discuss the challenges of successfully implementing our approach in terms of (1) the ethical and qualitative data requirements and (2) the methodological requirements in the light of inadequate, existing benchmark models. Lastly, we outline a research agenda to address these issues.

2 RELATED WORK

Informal settlement mapping is an increasingly prominent research objective (Hofmann et al., 2015), but poses distinct challenges, as the morphology and structure of informal settlements varies substantially—across different cultural spaces and even within cities (Taubenböck et al., 2018). As of recent, machine learning approaches provide promising solutions to account for these complexities (Maiya & Babu, 2018; Wurm et al., 2019). Yet, few studies are built around active community engagement, such as participatory mapping of informal settlements (Yuan et al., 2018). We believe integrating local communities into work such as ours is not only helpful for leveraging the locals’ expertise, but also provides them with agency over their data, enables further conversation on the equitability and performance of resulting applications and prevents model bias stemming from poor-quality ground-truth. Informal settlements can be seen as self-contained systems, including their own places of worship, entertainment and trade. Many structures are not inhabited, but rather




Figure 1: Proposed population estimation framework with equitable ground-truth through participatory online digital mapping with on-site validation; Figure partially adapted from Porto de Albuquerque et al. (2019).

being used for other purposes (e.g. shops or religious buildings). Precise population estimates are thus often hard to obtain. Recently, Wardrop et al (2018) highlighted the need for better and more granular population estimates in informal settings. Traditionally, spatial disaggregation approaches utilizing auxiliary information such as road networks or vegetation have been used (Wardrop et al., 2018), while more recent machine learning focused work has proposed disaggregation using points-of-interest (Zong et al., 2019) or combining satellite images with existing census data (Weber et al., 2018; Hu et al., 2019). A further, methodological challenge is that few deep learning methods tailored to satellite imagery exist. Together with the fact that applicable training data and reliable ground-truth can be sparse and expensive to produce, building models which can be trained on limited data but perform well on a global scale are the "holy grail" of informal settlement analysis. So far, this task has been proven to be exceptionally challenging: Wurm et al. (2019) point out that data quality and resolution differences are a substantial problem in their transfer learning task on informal settlement satellite imagery. However, Xie et al. (2016) find that rather simple data such as nighttime lights can be used for successful transfer learning and poverty mapping. They also note that pre-training on standard datasets such as *ImageNet* can be useful. Table 1 provides an overview on data requirements, spatial resolution, methods and equitability of existing approaches, as compared to our proposed approach.

3 DATA & EQUITABLE GROUND-TRUTH

Our experiments are conducted using data from three sites: two settlements in Ibadan, Nigeria (Sasa, Idikan) provide training data and one in Lagos, Nigeria (Bariga) test data. Our data can broadly be split into two categories: Satellite image data (*input*) and ground-truth maps (*labels*):

Input: We acquire very high-resolution (VHR) satellite image data with RGB channels and a pixel length of 0.02 meters for each of the three test sites from two major mapping services via the *SAS Planet* software ¹, obtaining two images per informal settlement. Nevertheless, this sub-meter resolution is not indicative of quality; rather we find the images to be quite noisy, especially compared to VHR images of high-income cities such as New York. Satellite data was collected in 2019, the ground-truth labels were mapped in 2018 (note that the images used to create the ground-truth-maps are different from the images we use for model training). We find that the temporal difference in accessing the respective data does not affect the match between both datasets. This implies that all three informal settlements did not change substantially during the period between community-mapping and accessing satellite data. Nevertheless, informal settlements are inherently dynamic and if different imagery is used for mapping and modeling, the overlap between them should always be assessed to guarantee consistent results.

Labels: Active community engagement is essential for avoiding biased or misinformed labels which might transfer into biased models. We use data from the University of Warwick’s *Improving Health in Slums Collaborative* (IHSC, 2019), which is derived from a participatory community mapping framework that aims to achieve equitable, empowering community engagement while at the same time conforming to rigorous, reliable and generalizable data standards (de Albuquerque et al., 2019).

¹Accessible here: <http://www.sasgis.org/>




Figure 2: Random sample of training data. Highlighted on the lower right is an exemplary image for high residency

This framework is based on the premise that the generation of data by citizens can also be an opportunity for mutual learning and to develop a new critical consciousness on the local realities that are being mapped in the data (de Albuquerque & de Almeida, 2020). While the *IHSC* paper should be consulted for details, we want to outline the key steps of this approach briefly (also, see Figure 1): In preparation for the mapping, local stakeholders are consulted to identify key local partners. These partners are essential for understanding the complexities of the physical and social environments in informal settlements. The next step involves the creation of a base map from high-resolution satellite imagery. This mapping process can be conducted by locals and remote participants using open-source collaborative software such as the *HOT Tasking Manager* (see: <https://tasks.hotosm.org/>). Lastly, "ground-truthing" is conducted at the site using participatory mapping by specially trained local citizens. This validates and enriches the data gathered in the previous step. The final map can then be contributed as open geographic data to volunteered geographic information (VGI) services like *OpenStreetMap*. *OpenStreetMap* is a "wikipedia of maps"; a collaborative project aim at creating a free editable map of the world. Consequently, our definition of "equitable ground-truth" (and how we compare it to existing work, as in Table 1 entails all the principles outlined above: *pre-assessment & local partners, community participation & education, double validation (remote and on-site)* and *transparency & accessibility*.

4 DISCUSSION

In order to explore how we may learn differently with this community data, we decided to look at the process of classifying buildings into inhabited vs non-inhabited from satellite images. This is a difficult but important question: Building an understanding of where people live is crucial for estimating needs in emergency scenarios and for general public service provision and planning. Below we give some of the steps and results of our preliminary analysis.

First we split the satellite images into a grid of 36x36 meter tiles. We then rescale each image from 1600x1600 pixels to 224x224 pixels - given how noisy the images are there was not much loss of information with this compression. For one of the settlements—Lagos-Bariga—we also removed images with cloud cover, calculating the ratio of white pixels as a proxy - we set threshold at 5%, which gives us as many images as possible. After creating the grid, each image is assigned a respective residential occupancy level (see Figure 1), calculated as the percentage of pixels occupied by residential buildings. We transform this into a classification problem by declaring images with 30% or more residential occupancy level as residential (= 1) and less than that as non-residential (= 0). We chose this threshold through a preceding analysis, comparing each image to its respective building masks and assessing visual cues in the images, while avoiding further class imbalances. Even though this might seem a fairly low residential occupancy level, this threshold corresponds to the true ratio of occupied versus non-occupied buildings, which we observe to be very similar in all three sites. About 1/3 of the data are assigned a "residential" label and 2/3 a "non-residential" label.

For our baseline analysis, we trained a *ResNet50* model (150 epochs) (He et al., 2016) and a *MobileNetV2* (30 epochs) (Sandler et al., 2018) using data from Ibadan-Sasa and Ibadan-Idikan for training- which gave us a total of 1056 training examples - and from Lagos-Bariga for validation - 646 validation examples. We found that *MobileNetV2*, trained with only our dataset performed better than both *MobileNetV2* with pre-trained weights from the *Imagenet* dataset. Validation performance converged after running 30/50 epochs. In this set up we got a training accuracy of 75% and validation accuracy of 63%. In several studies in machine learning, transfer learning has proven

to be a great tool to learn under constrained resources Zoph et al. (2016); Kolesnikov et al. (2019). In our case we find that in questions as complex as that of building separability, standard methods of pre-training with benchmark datasets does not help us much and in fact harms accuracy. This is not a novel finding Ngiam et al. (2018), but it once again underlines the importance of curating relevant datasets, especially for applied work. While our preliminary results may suggest that we are not close to a practicable solution, we gained crucial insights to guide future research directions. Further they allow us to comment on lessons-learned and best-practices when working with equitable ground-truth data:

- During the data preparation, we realised that there were spatial cues in images such as how square a roof is or how square neighbouring roofs are, that help the classification task. However, we have learned automating this separation is a lot more complex. That said, this might just be a function of the limited dataset.
- Ethical considerations regarding the detachment between the "analysts" (e.g. researchers) and "subjects" (informal settlement population) need to take a more central role in machine learning based population mapping tasks. Although different organisations are increasingly discussing issues around bias and adopting frameworks based on principles for "responsible AI", there is still a lack of clear processes for engaging with the local communities that are represented in the data used in machine learning systems. We hope that by having the local community as part of the annotators, we have provided the community with hitherto unavailable data to understand their contexts, while simultaneously generating a more representative training dataset. However, we also conclude that deeper analyses of the informal and often undocumented implications of machine learning systems are needed. Failing to do so would not only exclude or misrepresent large population parts ², but also may lead to severe problems with downstream tasks. It is hence crucial that the machine learning community actively works with practitioners and local stakeholders. Further, best-practice guidelines for generating and working with equitable ground-truth data should be developed, both in general and domain-specific, such as the recent proposal of an integrated deprived area mapping system (Kuffer et al., 2020).
- We observe a certain lack of transferability of existing, pre-trained vision models to our task. In particular, we see that larger—and presumably more powerful—models perform worse than smaller models. This raises the question of how much real-world relevance existing benchmarks hold. This also serves as a call-to-action for developing tailored methodologies and domain-specific benchmarks for common applications of machine learning in the development context.

5 CONCLUSION

Government authorities and NGOs require precise population estimates for informal settlements to optimize their services and health interventions. Often the subject of these estimates, the population, is excluded from the data gathering process. In this paper we discuss the merits and difficulties of deploying machine learning for population estimation in informal settlements, building on equitable ground-truth data, gathered in collaboration with local population and stakeholders. This approach aims at empowering local communities to take agency over their own data and at eliminating potential root causes of biases further down the modelling pipeline. We argue that this method is not only more ethical, but also leads to a higher data quality, as local knowledge is explicitly leveraged throughout the annotation process.

We run experiments using pre-trained and fine-tuned vision models. Interestingly, we observe that pre-trained models don't fare as well as may be expected, and smaller models such as *MobileNetV2* perform best. This raises concerns on the value ascribed to common computer vision benchmark models and their transferability to real world tasks. We hope our findings encourage further investigations in this area within the ML for development community.

²In Nigeria, for example, the informal sector accounts for 50-65% of GDP. IMF 2017 working paper: <https://www.imf.org/en/Publications/WP/Issues/2017/07/10/The-Informal-Economy-in-Sub-Saharan-Africa-Size-and-Determinants-45017>

ACKNOWLEDGMENTS

First, we wish to that thank Nyalleng Moorosi for her guidance and contribution to this project. The authors also gratefully acknowledge funding from the UK Engineering and Physical Sciences Research Council, the EPSRC Centre for Doctoral Training in Urban Science (EPSRC grant no. EP/L016400/1); The Alan Turing Institute (EPSRC grant no. EP/N510129/1). The authors would like to express gratitude to all project members of the NIHR Improving Health in Slums Collaborative for their invaluable input, to informal settlement residents and to the mappers of the OpenStreetMap community who were active locally and remotely to create the data used in this paper.

REFERENCES

- David Castells-Quintana. Malthus living in a slum: Urban concentration, infrastructure and economic growth. *Journal of Urban Economics*, 98:158–173, mar 2017. ISSN 00941190. doi: 10.1016/j.jue.2016.02.003. URL <https://www.sciencedirect.com/science/article/pii/S0094119016000115>.
- João Porto de Albuquerque and André de Almeida. Modes of engagement: reframing ‘sensing’ and data generation in citizen science for empowering relationships. In Alice Mah and Thom Davies (eds.), *Toxic Expertise: Environmental Justice and Citizen Science in a Post Truth Age*. Manchester University Press, Manchester, 2020.
- João Porto de Albuquerque, Godwin Yeboah, Vangelis Pitidis, and Philipp Ulbrich. Towards a participatory methodology for community data generation to analyse urban health inequalities: a multi-country case study. In *Proceedings of the 52nd Hawaii International Conference on System Science*, jan 2019. ISBN 978-0-9981331-2-6. doi: 10.13140/RG.2.2.13710.20804. URL <https://scholarspace.manoa.hawaii.edu/handle/10125/59829>.
- Ilke Demir, Krzysztof Koperski, David Lindenbaum, Guan Pang, Jing Huang, Saikat Basu, Forest Hughes, Devis Tuia, and Ramesh Raska. DeepGlobe 2018: A challenge to parse the earth through satellite images. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, volume 2018-June, pp. 172–181, 2018. ISBN 9781538661000. doi: 10.1109/CVPRW.2018.00031.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2016-Decem, pp. 770–778, 2016. ISBN 9781467388504. doi: 10.1109/CVPR.2016.90. URL <http://image-net.org/challenges/LSVRC/2015/>.
- Peter Hofmann, Hannes Taubenbock, and Christian Werthmann. Monitoring and modelling of informal settlements - A review on recent developments and challenges. In *2015 Joint Urban Remote Sensing Event (JURSE)*, pp. 1–4. IEEE, mar 2015. ISBN 978-1-4799-6652-3. doi: 10.1109/JURSE.2015.7120513. URL <http://ieeexplore.ieee.org/document/7120513/>.
- Wenjie Hu, Jay Harshadbhai Patel, Zoe-Alanah Robert, Paul Novosad, Samuel Asher, Zhongyi Tang, Marshall Burke, David Lobell, and Stefano Ermon. Mapping Missing Population in Rural India: A Deep Learning Approach with Satellite Imagery. In *AAAI/ACM Conference on AI, Ethics, and Society (AIES 2019)*, may 2019. doi: 10.1145/3306618.3314263. URL <http://arxiv.org/abs/1905.02196><http://dx.doi.org/10.1145/3306618.3314263><http://arxiv.org/abs/1905.02196>{%}0A<http://dx.doi.org/10.1145/3306618.3314263>.
- IHSC. A protocol for a multi-site, spatially-referenced household survey in slum settings: methods for access, sampling frame construction, sampling, and field data collection. *BMC medical research methodology*, 19(1):109, may 2019. ISSN 14712288. doi: 10.1186/s12874-019-0732-x.
- Neal Jean, Marshall Burke, Michael Xie, W Matthew Davis, David B Lobell, and Stefano Ermon. Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790–794, aug 2016. ISSN 10959203. doi: 10.1126/science.aaf7894. URL <http://www.ncbi.nlm.nih.gov/pubmed/27540167>.

- Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Large scale learning of general visual representations for transfer, 2019.
- Monika Kuffer, Karin Pfeffer, and Richard Sliuzas. Slums from space-15 years of slum mapping using remote sensing. *Remote Sensing*, 8(6):455, may 2016. ISSN 20724292. doi: 10.3390/rs8060455. URL <http://www.mdpi.com/2072-4292/8/6/455>.
- Monika Kuffer, Dana R. Thomson, Gianluca Boo, Ron Mahabir, Täis Grippa, Sabine Vanhuyse, Ryan Engstrom, Robert Ndugwa, Jack Makau, Edith Darin, João Porto de Albuquerque, and Caroline Kabaria. The role of earth observation in an integrated deprived area mapping “system” for low-to-middle income countries. *Remote Sensing*, 12(6), 2020. ISSN 2072-4292. doi: 10.3390/rs12060982. URL <https://www.mdpi.com/2072-4292/12/6/982>.
- Shishira R Maiya and Sudharshan Chandra Babu. Slum Segmentation and Change Detection : A Deep Learning Approach. In *NeurIPS Workshop on Machine Learning for Development (ML4D)*, 2018. URL <https://arxiv.org/abs/1811.07896><http://arxiv.org/abs/1811.07896>.
- Jiquan Ngiam, Daiyi Peng, Vijay Vasudevan, Simon Kornblith, Quoc V. Le, and Ruoming Pang. Domain adaptive transfer learning with specialist models, 2018.
- Claudio Persello and Alfred Stein. Deep Fully Convolutional Networks for the Detection of Informal Settlements in VHR Images. *IEEE Geoscience and Remote Sensing Letters*, 14(12):2325–2329, dec 2017. ISSN 1545598X. doi: 10.1109/LGRS.2017.2763738. URL <http://ieeexplore.ieee.org/document/8094993/>.
- Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang Chieh Chen. MobileNetV2: Inverted Residuals and Linear Bottlenecks. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 4510–4520, 2018. ISBN 9781538664209. doi: 10.1109/CVPR.2018.00474.
- H. Taubenböck, N.J. Kraff, and M. Wurm. The morphology of the Arrival City - A global categorization based on literature surveys and remotely sensed data. *Applied Geography*, 92:150–167, mar 2018. ISSN 0143-6228. doi: 10.1016/J.APGEOG.2018.02.002. URL <https://www.sciencedirect.com/science/article/pii/S0143622817309955?via%3Dihub>.
- United Nations. The Sustainable Development Goals Report 2018. Technical report, 2018. URL <https://www.un.org/development/desa/publications/the-sustainable-development-goals-report-2018.html>.
- Steven van de Vijver, Samuel Oti, Clement Oduor, Alex Ezech, Joep Lange, Charles Agyemang, and Catherine Kyobutungi. Challenges of health programmes in slums. *The Lancet*, 386(10008): 2114–2116, nov 2015. ISSN 01406736. doi: 10.1016/s0140-6736(15)00385-2. URL <http://www.ncbi.nlm.nih.gov/pubmed/26452707>.
- N A Wardrop, W C Jochem, T J Bird, H R Chamberlain, D Clarke, D Kerr, L Bengtsson, S Jurran, V Seaman, and A J Tatem. Spatially disaggregated population estimates in the absence of national population and housing census data. *Proceedings of the National Academy of Sciences of the United States of America*, 115(14):3529–3537, apr 2018. ISSN 1091-6490. doi: 10.1073/pnas.1715305115. URL <http://www.ncbi.nlm.nih.gov/pubmed/29555739><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC5889633>.
- Eric M. Weber, Vincent Y. Seaman, Robert N. Stewart, Tomas J. Bird, Andrew J. Tatem, Jacob J. McKee, Budhendra L. Bhaduri, Jessica J. Moehl, and Andrew E. Reith. Census-independent population mapping in northern Nigeria. *Remote Sensing of Environment*, 204:786–798, jan 2018. ISSN 0034-4257. doi: 10.1016/J.RSE.2017.09.024. URL <https://www.sciencedirect.com/science/article/pii/S0034425717304364?via%3Dihub>.
- Michael Wurm, Thomas Stark, Xiao Xiang Zhu, Matthias Weigand, and Hannes Taubenböck. Semantic segmentation of slums in satellite images using transfer learning on fully convolutional neural networks. *ISPRS Journal of Photogrammetry and Remote Sensing*, 150:59–69, apr 2019. ISSN 09242716. doi: 10.1016/j.isprsjprs.2019.02.006. URL <https://www.sciencedirect.com/science/article/pii/S0924271619300383>.

Michael Xie, Neal Jean, Marshall Burke, David Lobell, and Stefano Ermon. Transfer Learning from Deep Features for Remote Sensing and Poverty Mapping. In *Thirtieth AAAI Conference on Artificial Intelligence*, mar 2016. URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/viewPaper/12196><http://arxiv.org/abs/1510.00098>.

Jiangye Yuan, Pranab K. Roy Chowdhury, Jacob McKee, Hsiuhan Lexie Yang, Jeanette Weaver, and Budhendra Bhaduri. Exploiting deep learning and volunteered geographic information for mapping buildings in Kano, Nigeria. *Scientific Data*, 5(1):180217, dec 2018. ISSN 2052-4463. doi: 10.1038/sdata.2018.217. URL <http://www.nature.com/articles/sdata2018217>.

Zefang Zong, Jie Feng, Kechun Liu, Hongzhi Shi, and Yong Li. DeepDPM: Dynamic Population Mapping via Deep Neural Network. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2019. URL <https://www.aaai.org/ojs/index.php/AAAI/article/view/3925><http://arxiv.org/abs/1811.02644>.

Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. Transfer learning for low-resource neural machine translation. *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2016. doi: 10.18653/v1/d16-1163. URL <http://dx.doi.org/10.18653/v1/D16-1163>.