

Sensing Population Distribution from Satellite Imagery Via Deep Learning: Model Selection, Neighboring Effects, and Systematic Biases

Xiao Huang^{1b}, Di Zhu, Fan Zhang, Tao Liu, Xiao Li^{1b}, and Lei Zou

Abstract—The rapid development of remote sensing techniques provides rich, large-coverage, and high-temporal information of the ground, which can be coupled with the emerging deep learning approaches that enable latent features and hidden geographical patterns to be extracted. This article marks the first attempt to cross-compare performances of popular state-of-the-art deep learning models in estimating population distribution from remote sensing images, investigate the contribution of neighboring effect, and explore the potential systematic population estimation biases. We conduct an end-to-end training of four popular deep learning architectures, i.e., VGG, ResNet, Xception, and DenseNet, by establishing a mapping between Sentinel-2 image patches and their corresponding population count from the LandScan population grid. The results reveal that DenseNet outperforms the other three models, while VGG has the worst performances in all evaluating metrics under all selected neighboring scenarios. As for the neighboring effect, contradicting existing studies, our results suggest that the increase of neighboring sizes leads to reduced population estimation performance, which is found universal for all four selected models in all evaluating metrics. In addition, there exists a notable, universal bias that all selected deep learning models tend to overestimate sparsely populated image patches and underestimate densely populated image patches, regardless of neighboring sizes. The methodological, experimental, and contextual knowledge this article provides is expected to benefit a wide range of future studies that estimate population distribution via remote sensing imagery.

Index Terms—Deep learning, end-to-end architecture, population estimation, satellite imagery, systematic biases.

I. INTRODUCTION

FINE knowledge of the spatial contribution of human activity is essential for a wide range of fields, such as public health [1]–[3], urban planning [4]–[6], disaster management

[7], [8], resource allocation [9], economic evaluation [10], and migration [11], [12]. As stated by the global sustainable development goals, understanding where and how people are distributed is of great importance to make “cities and human settlements inclusive, safe, resilient, and sustainable.”

Census data, when linked with accurate administrative boundary data, can provide spatially explicit population distribution. In the U.S., for instance, official population data on various geographical levels [e.g., the Decennial Census records and the American Community Survey (ACS)] are repetitively released by the U.S. Census Bureau. Similar agencies that regularly release population distribution data include the Office for National Statistics of the U.K., the National Bureau of Statistics of China, and the Statistics Bureau of Japan, to list a few. Despite the authority of census-based population distribution released by the officials, it owns several intrinsic limitations, making it ill-suitable for many spatial problems. First, population distribution is with great heterogeneity [13]; therefore, it cannot be assumed uniformly distributed within predefined geographical units. Second, the census-based population suffers from the modifiable areal unit problem (MAUP) [14] due to its arbitrarily imposed boundaries that are rarely consistent with other boundaries in practical applications [15]. Third, census-based population distribution is often with poor temporal resolutions that preclude temporal-dynamic population estimations, and recent and reliable population data at fine scales can often be lacking, especially in resource-poor settings [16], [17]. Given the above limitations, scholars start to explore various means to improve the aggregated census-based population, one notable effort of which is to derive fine-grained, spatially-continuous population grids.

A population grid refers to a geographically referenced lattice of square cells, with the value of each cell representing a population count at its location. Population grids are generally constructed based on census unit-based data via dasymetric modeling [18], [19], or other statistical approaches that intelligently assign population to grids by establishing relationships between population and supporting auxiliary variables [20], [21]. The derived population grids not only capture the heterogeneity of population distribution, but also ensure the aggregated numbers at census units match the official records [22]. Despite that population distribution delivered in fine-grained grided format achieves spatial heterogeneity and largely mitigates the MAUP, the production of accurate population grids is temporally

Manuscript received March 2, 2021; revised April 4, 2021 and April 25, 2021; accepted April 27, 2021. Date of publication April 29, 2021; date of current version May 27, 2021. (Corresponding author: Xiao Huang.)

Xiao Huang is with the Department of Geosciences, University of Arkansas, Fayetteville, AR 72701 USA (e-mail: xh010@uark.edu).

Di Zhu is with the Department of Geography, Environment, and Society, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: dizhu@umn.edu).

Fan Zhang is with the Department of Urban Studies and Planning, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: zhangfan@mit.edu).

Tao Liu is with the College of Forest Resources and Environmental Science, Michigan Technological University, Houghton, MI 49931 USA (e-mail: taoliu@mtu.edu).

Xiao Li is with the Texas A&M Transportation Institute, Texas A&M University, College Station, TX 77843 USA (e-mail: xiao.li@tamu.edu).

Lei Zou is with the Department of Geography, Texas A&M University, College Station, TX 77843 USA (e-mail: lzou@exchange.tamu.edu).

Digital Object Identifier 10.1109/JSTARS.2021.3076630

restricted to the release of census data [16] and is heavily dependent on the availability of a diverse set of inputs (often large in number), leading to its questionable sustainability. With the development of statistical modeling techniques, questions arise: can we obtain a mapping between the gridded population and the gridded latent variables? If so, can we achieve a relatively robust mapping using as few and as easily-accessible variables as possible?

The rapid development of remote sensing techniques facilitates the consistency and sustainability of auxiliary variables, which are often preferred and needed in population distribution modeling, as remote sensing imagery, with improving imaging capability over time, provides rich, large-coverage, and high-temporal information of the ground. The advantages of remote sensing imagery are further expanded by the emerging supervised deep learning approaches with the capacity of extracting both low- and high-level latent features and capturing hidden hierarchies of geographical patterns from images [23]–[25], thus forming a stable supervised end-to-end mapping [26]. The coupling of remote sensing imagery and deep learning algorithms undoubtedly establishes a new venue that potentially advances traditional population modeling. Numerous efforts have been made to harvest the strong mapping capability of deep learning in estimating both population count and population density directly from satellite images [27]–[30]. These efforts differ in architecture design, model selection, and evaluation metrics. Despite the aforementioned attempts, gaps still exist in evaluating the performance of state-of-the-art deep learning models to provide guidelines of model selection for future studies, thoroughly investigating the potential contribution of neighboring image patches to the population estimation in the center image patch (neighboring effects) to provide guidance for the selection of suitable neighboring sizes, and exploring the potential systematic population estimation biases resulting from the intrinsic limitations of remote sensing images.

To fill the above gaps, we perform end-to-end training on popular deep learning architectures by establishing a mapping between satellite image patches and their corresponding population count from an existing gridded population product. Taking the Metropolitan Atlanta (Metro Atlanta) and Metropolitan Dallas (Metro Dallas) as the study region, we obtain remote sensing images from Sentinel-2 with a spatial resolution of 10 m and derive the 24-h ambient gridded population distribution from LandScan (<https://landscan.ornl.gov/>). Models are trained, validated, and tested in Metro Atlanta and further applied to Metro Dallas to evaluate the overall generalizability. The main contributions of this article are summarized as follows.

1) We adopt transfer learning techniques by modifying and fine-tuning several popular deep learning architectures toward the population estimation task. We further cross-compare estimated population distribution from these architectures quantitatively and qualitatively, providing guidelines for the model selection in future studies.

2) We investigate the neighboring effects by progressively extending the selection of neighboring patches to estimate the population residing in the center image patch, offering guidance for the selection of suitable neighboring sizes.

3) We explore the potential systematic biases in population estimation directly from remote sensing images via deep learning approaches, discuss reasonable causes that lead to these biases, and provide applicable solutions to mitigate these biases.

II. RELATED WORKS

A. Existing Population Grid Products and Their Issues

Gridded population contain regulated cells with population count, serving as an ideal input for training a mapping between image patches and the corresponding population statistics. Numerous global- and regional-focused population grid products have been developed and released with various spatiotemporal scales. Popular population products that contain the U.S. include (but is not limited to): Gridded Population of the World (GPW) [31], Global Rural Urban Mapping Project (GRUMP) [21], Global Human Settlement–Population (GHS-POP) [32], World Population Estimate (WPE) [33], LandScan-USA [34], and Building-based Population Grid USA (BPG-USA) [13]. Among them, GPW, GRUMP, GHS-POP, and WPE are at a global scale, while LandScan-USA and BPG-USA specifically target the U.S. and Conterminous U.S., respectively. Dasymeric techniques are often applied to allocate population count to habitable cells within each census unit, but with different input of auxiliary variables and weighting scenarios. Commonly used auxiliary variables involve land use/land cover (usually derived from remote sensing imagery), nighttime light intensity, distribution of infrastructures (e.g., roads, buildings, and point of interest), environmental variables, and restrictions (e.g., water body and protected areas), to list a few. Population grid products generally ensure that aggregated numbers at census units match the official records, at the same time, capture the heterogeneity of population distribution using regulated cells. However, massive inputs, especially for accurately modeled gridded population, need to be prepared and updated in accordance with the release of the new census population [35], posing a great challenge to the continuous product release, as auxiliary variables with differing spatiotemporal scales might be difficult to collect [16]. In addition, census-disaggregated population grids largely rely on the availability of census data. Such an issue is exaggerated in data-poor settings, where census data and required auxiliary variables are obsolete or unavailable at all.

The development of remote sensing platforms with improving imaging capability leads to the easy acquisition of rich, large-coverage, and high-temporal information of the ground, serving as an ideal set of variables that benefits continuous population distribution mapping. Deep learning techniques further enhance the ability to extract high-level latent features with hidden hierarchies of geographical patterns. Once a stable mapping is established between the gridded population and the corresponding image patch, population distribution mapping can be achieved only with remote sensing images, which are easily accessible and regularly updatable.

B. Popular Deep Learning Architectures and Transfer Learning

Most modern deep learning models are based on artificial neural networks, specifically convolutional neural networks

(CNNs) that were inspired by the visual system’s structure and hierarchically composed of input, output, and multiple hidden layers [36]. AlexNet by Krizhevsky *et al.* [37] was a pioneering CNN that transcended traditional approaches in various computer vision tasks. Further, the VGG by Simonyan and Zisserman [38] improved the performance by stacking simple convolutional operations, forming a deep architecture. Diverging from the mainstream of stacking layers via a sequential structure, GoogLeNet [39] achieved great performance by introducing a building block, i.e., the inception module, that largely reduces the number of parameters and operations. Numerous updated versions of GoogLeNet were proposed: Inception-V2 and -V3 [40]; Inception-V4 [41]; and Xception [42]. ResNet [43] is well-known for its depths and the introduction of residual blocks that implement identity skip connections to promote gradient propagation. From another perspective in solving the problem of vanishing gradient, DenseNet [44] facilitates the acquisition of “collective knowledge” by allowing each layer to obtain additional inputs from all preceding layers.

Training the above architectures from scratch is not often feasible due to the limited sample size and the temporal inefficiency for models with randomly initialized parameters to converge. Studies have proved that using pretrained weights from similar tasks greatly benefits the model training process [45], [46]. In this article, we adopt transfer learning techniques by loading the weights pretrained from ImageNet¹ as initialization. We select four popular deep learning architectures and compare their performances in deducing population counts given the corresponding remote sensing image patches. These modern architectures, including VGG, ResNet, Xception, and DenseNet, are widely applied in various computer vision tasks and generally diverge from each other in the overall architecture design.

C. Deriving Population Distribution Directly From Satellite Imagery Via Deep Learning

Numerous attempts have been made to derive population distribution from satellite imagery supported by the strong capability of feature extraction in deep learning. Doupe *et al.* [30] converted satellite images into population density estimates using a VGG-like architecture that includes collections of “convolutional layer (Conv), pooling layer (Pool) and rectified linear unit” followed by two fully-connected layers, each with 4096 neurons. They trained and implemented their model in Tanzanian and Kenyan and found that their model achieved great performance with decent generalizability. Robinson *et al.* [29] adopted a similar VGG-like sequential architectural design to Doupe *et al.* [30] but regarded the task as a classification problem instead of a regression problem. They aimed to classify image patches to labels that derive from the power level of the population count (14 classes in total). The results suggested that the predictions from their model were consistent with the ground-truthing labels with an R^2 over 0.9 in selected areas in the U.S. Hu *et al.* [27] built a customized CNN architecture to predict population density in India from multisource imagery fused by the implementation of 1×1 Conv layers. Their results, again,

demonstrated the feasibility of producing accurate population estimates directly from satellite imagery, especially for rural areas. A more recent effort was by Xing *et al.* [28], who implemented a ResNet-based architecture with the consideration of neighboring effects, termed as neighbor-ResNet. The proposed neighbor-ResNet architecture used ResNet as a backbone and aimed to estimate the population residing in the center image patch, with the support of the eight surrounding neighboring patches (within a 3×3 patch space). The results proved the superiority of neighbor-ResNet compared to regular ResNet that ignores the neighboring effects.

Despite the above efforts, few studies conduct comparisons among state-of-the-art deep learning architectures in terms of evaluating their performances in the mapping between satellite imagery and population. The recent effort by Xing *et al.* [28] identified the positive contribution of neighboring patches but with a fixed 3×3 patch space. Efforts are needed to explore the dynamics of performance when a different size of a surrounding region is considered.

III. DATASETS AND STUDY AREAS

A. Datasets

1) *Population Grid (LandScan USA Population Database)*: In this article, we selected the population grid from LandScan² as ground-truthing population distribution, serving as a reference that allows an end-to-end mapping to be established by deep learning architectures between satellite image patches and the corresponding population count. LandScan belongs to Oak Ridge National Laboratory, the largest science and energy national laboratory in the Department of Energy. LandScan’s population grid has been widely recognized as one of the community standard products [47] and applied in a variety of domains [48]–[50]. LandScan team adopts multivariate dasymetric modeling frameworks using the best available demographic (census) data, remote sensing imagery, and other supporting variables to disaggregate census population counts into grids [34]. The population product we selected is the LandScan USA Population Database 2019, currently hosted by homeland infrastructure foundation-level data. It provides estimated population counts at 3 arc-second resolution for nighttime and daytime for Conterminous U.S., Hawaii, Alaska, as well as other U.S. territories. We averaged the daytime and nighttime population grids to form a 24-h ambient population grid, as we believe information from remote sensing images does not target a specific category of population distribution (daytime or nighttime), but an ambient one. Note that LandScan USA Population Database 2019 is a static baseline population estimate, which does not include transitory populations, such as business travelers and tourists [51]. Given the fact that the population grid with a 3 arc-second resolution presents great heterogeneity that a training patch with medium-resolution satellite imagery might fail to capture (a 3 arc-second grid only contains roughly 10×10 pixels of Sentinel-2 imagery), we aggregated the original population grid to 30 arc-second (approximately 1 km at the equator), a spatial resolution adopted

¹Online. [Available]: <http://www.image-net.org/>

²Online. [Available]: <https://landscan.ornl.gov/>

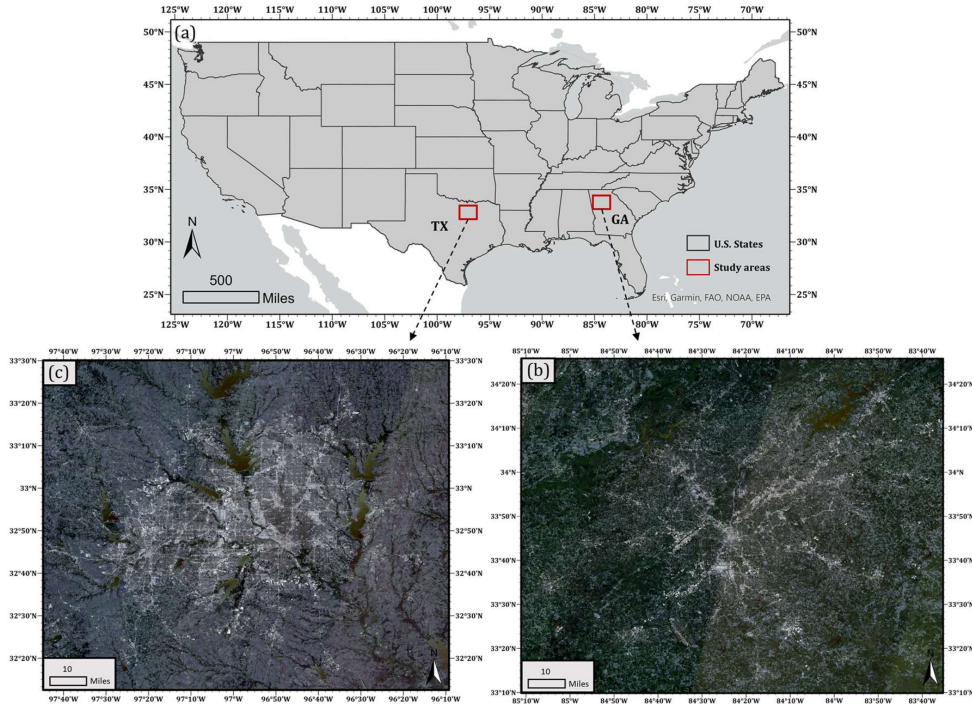


Fig. 1. Study areas with two sites. (a) Conterminous U.S. (b) Metro Atlanta (model training, validating, and testing). (c) Metro Dallas (evaluating generalizability).

by many popular population products (i.e., GPW, GRUMP, GHS-POP, etc.) and studies similar to this article [27]–[29].

2) *Satellite Imagery (Sentinel-2)*: The satellite imagery used in this study was derived from Sentinel-2, a wide-swath, fine-resolution, multispectral imaging mission of the European Space Agency developed in the framework of the European Union Copernicus program [52]. We selected the Sentinel-2 Level-2A product that had been atmospherically corrected (bottom of atmosphere reflectance) and orthorectified. We retrieved four bands, i.e., red (R), green (G), blue (B), and near-infrared band (NIR), all with a spatial resolution of 10 m. Studies have demonstrated that the physical and chemical characteristics of various types of land use and land cover can be well reflected by these four bands [53], [54]. Thus, we expected a stable mapping that associates population count with these bands to be formed by deep learning architectures through end-to-end training. In addition, NIR and visible spectrum that includes R, G, and B are available in most multispectral sensors. Serving as model inputs, the easy accessibility of these bands greatly promotes generalizability and sustainability. We queried Sentinel-2 imagery via Google Earth Engine (GEE). The imagery covers Metro Atlanta and Metro Dallas, with a temporal coverage from January 1 to December 31, 2019, consistent with the time span of the 2019 LandScan population grid. We implemented a standard Sentinel-2 cloud mask in GEE and selected the median value for each pixel with overlapping values among different scenes.

B. Study Areas

We selected two metropolitan regions within the Conterminous U.S. [see Fig. 1(a)] as our study areas, i.e., Metro Atlanta

[see Fig. 1(b)] and Metro Dallas [see Fig. 1(c)]. The article site that covers Metro Atlanta is bounded with latitude from 33.1750° N to 34.4250° N, and longitude from 85.1750° W to 85.1750° W. The article site that covers Metro Dallas is bounded with latitude from 32.2325° N to 33.4825° N, and longitude from 96.1500° W to 97.7334° W. These two metropolitan areas are both densely populated metropolises but with observable discrepancies in their distribution patterns of land use and land cover. Metro Atlanta is characterized by its sprawl-out urban fabrics [see Fig. 1(b)]. In comparison, the urban fabrics in Metro Dallas are comparably centralized, evidenced by its distinctly dense urban core [see Fig. 1(c)]. We further divided the two study sites into 30 arc-second grids, leading to a total of 28 500 (190 × 150) grids in each site for model training purposes. We trained deep learning architectures using patches from Metro Atlanta and evaluated their generalization capability using patches from Metro Dallas. We believe that the dissimilarity between these two sites benefits us in observing potential overfitting issues and systematic biases.

IV. METHODOLOGY

A. Preprocessing and Problem Formation

We first resized the daytime and nighttime population grids from LandScan to 30 arc-second and averaged them to form 24-h ambient population grids. Let p_i^{day} and p_i^{night} respectively denote the daytime and nighttime population count in grid i , the 24-h ambient population count in grid i (p_i) is calculated as

$$p_i = \begin{cases} \frac{p_i^{\text{day}} + p_i^{\text{night}}}{2}, & \text{if } \frac{p_i^{\text{day}} + p_i^{\text{night}}}{2} \geq 1 \\ 0, & \text{else} \end{cases} \quad (1)$$

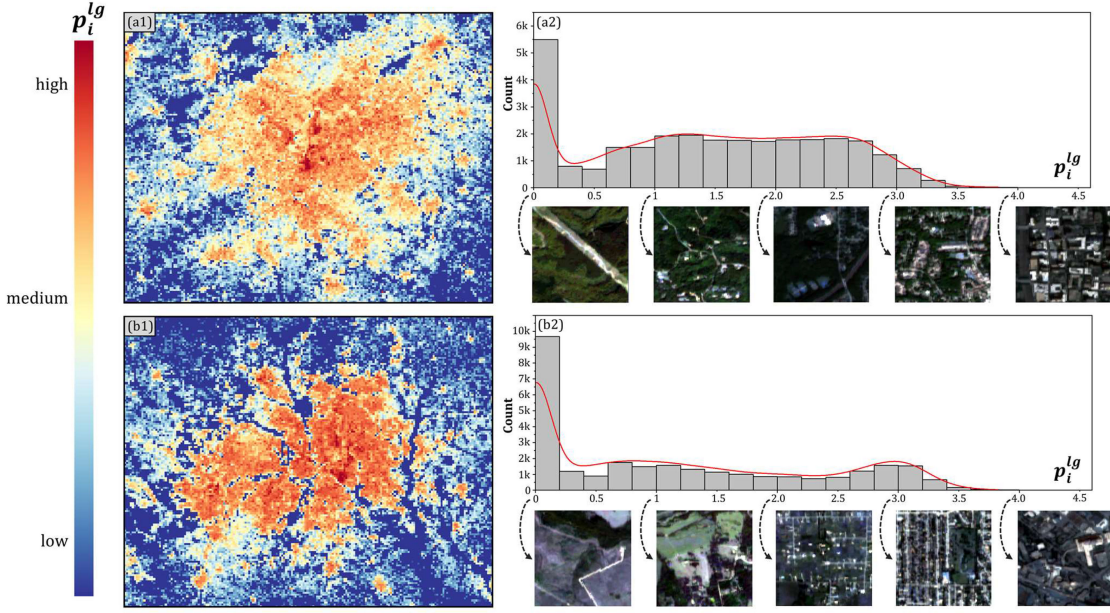


Fig. 2. The distribution of the logarithmized population within each grid (p_i^{lg}) in (a1) Metro Atlanta and (b1) Metro Dallas. Histograms and examples of image patches corresponding to the logarithmized population in (a2) Metro Atlanta and (b2) Metro Dallas.

For each grid i , there exists a corresponding image patch i , denoted as Pch_i . To investigate the potential contribution of neighboring patches to the population residing in the center image patch, we extended Pch_i to include its neighboring patches. Given image patch Pch_i , the extension, i.e., $\text{Pch}_i^{\text{nbr}(n)}$, that includes its $n \times n$ neighborhood can be formulated as

$$\text{Pch}_i^{\text{nbr}(n)} = \begin{bmatrix} \text{Pch}_j & \cdots & \text{Pch}_{j+n-1} \\ \vdots & \text{Pch}_i & \vdots \\ \text{Pch}_k & \cdots & \text{Pch}_{k+n-1} \end{bmatrix} \quad (2)$$

where n denotes the size of the neighborhood and $n \in \{1, 3, 57, 911\}$.

As we observed that the distribution of p_i was heavily tailed (even excluding the value of 0), we took the logarithms of the population count at grid i , denoted as $\log p_i$, following the works by Xing *et al.* [28] and Hu *et al.* [27]: $p_i^{lg} = \log_{10}(p_i)$. Specifically, if $p_i = 0$, we set p_i^{lg} as 0. Fig. 2 presents the distribution of $\log p_i$ in the two study areas: Metro Atlanta [see Fig. 2(a1)] and Metro Dallas [see Fig. 2(b1)]. Their histograms and examples of image patches corresponding to the logarithmized population count are presented in Fig. 2(a2) and 2(b2), respectively. We observed a spike of 0 values in both study areas, suggesting the existence of massive uninhabited grids. The logarithm operation flattened the curve for values above 0, leading to a considerably balanced target set that benefits model training.

We view population distribution estimation as a mapping problem. Given the extended image patch centered at Pch_i , i.e., $\text{Pch}_i^{\text{nbr}(n)}$, and the logarithmized population count at grid i , i.e., $\log p_i$, we aim to build a mapping function between them

$$\text{Pch}_i^{\text{nbr}(n)} \rightarrow p_i^{lg} \quad (3)$$

where symbol \rightarrow denotes the mapping to be learned by deep learning architectures. Note that $\text{Pch}_i^{\text{nbr}(n)}$ contains four channels, i.e., R, G, B, and NIR. We detailed this end-to-end architecture in the next session.

B. End-To-End Framework

We implemented an end-to-end training framework to facilitate the establishment of the mapping from $\text{Pch}_i^{\text{nbr}(n)}$ to p_i^{lg} (see Fig. 3). Considering the dynamic yet elusive relationship between remotely observed imagery and population distribution, deep neural networks are often regarded as preferred choices in modeling such a nonlinear, complex relationship [28]. In this article, we selected four widely-adopted deep learning architectures with diverging concepts in design: VGG (VGG-16); ResNet (ResNet-50); Xception; and DenseNet (DenseNet-121). We aim to compare their performances in deducing population count given the corresponding image patches with different neighboring considerations. As shown in Fig. 3, we first reconstructed $\text{Pch}_i^{\text{nbr}(n)}$ by considering the n by n neighborhood of Pch_i . We then resized $\text{Pch}_i^{\text{nbr}(n)}$ to a fixed size of 256×256 . Each reconstructed $\text{Pch}_i^{\text{nbr}(n)}$ has four channels, i.e., R, G, B, and NIR. Aiming to fine-tune weights of backbone models pre-trained from ImageNet, we implemented 1×1 Conv layers, commonly-used layers for image fusion and dimensionality reduction [43], to reduce the number of channels from four to three, consistent with the situation where the initial weights from ImageNet were derived. Further, we fed the rechanneled $\text{Pch}_i^{\text{nbr}(n)}$ to the backbones (with top fully connected layers cut off) of the four selected deep learning architectures. We appended a global average pooling layer to each architecture backbone, flattened the resulted neurons with 0.5 as dropout

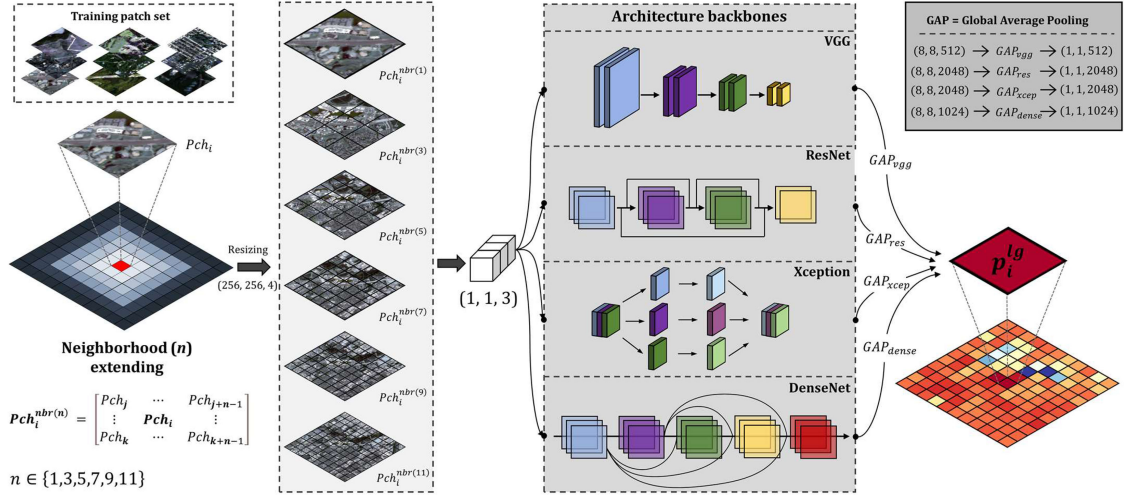


Fig. 3. End-to-end framework to derive the mapping between remote sensing imagery and population distribution. This framework tests the performances of popular deep learning architectures with different neighboring considerations. A dropout ratio of 0.5 is implemented after the results from the global average pooling layers are flattened. Note that $Pch_i^{nbr(n)}$ contains four channels, i.e., R, G, B, and NIR.

ratio before connecting them to the output neuron, whose linear activation suggests the estimated p_i^{lg} , denoted as \widehat{p}_i^{lg} .

C. Training Strategies

We chose log hyperbolic cosine (log-cosh) loss for back-propagation learning and weight updating

$$\mathcal{L} \left(p_i^{lg}, \widehat{p}_i^{lg} \right) = \sum_{i=1}^s \log_{10} \left(\cosh \left(\widehat{p}_i^{lg} - p_i^{lg} \right) \right) \quad (4)$$

where $\mathcal{L}(p_i^{lg}, \widehat{p}_i^{lg})$ denotes the log-cosh loss function, given the ground-truthing p_i^{lg} and estimated \widehat{p}_i^{lg} , s denotes the batch size (set as 32), and $\cosh(x)$ denotes the hyperbolic cosine function, i.e., $\frac{e^x + e^{-x}}{2}$. The log-cosh is similar to the mean squared error loss (i.e., the $L2$ loss), but is more tolerant of abnormal predictions [55].

We divided the two study areas into 30 arc-second grids, leading to a total of 28 500 (190×150) grids in each site. Image patches in Metro Atlanta were randomly divided into a training set (17 100 patches, 60%), a validating set (5700 patches, 20%), and a testing set (5700 patches, 20%). The spatial distribution of Metro Atlanta's image patches in different sets is presented in Fig. 10 (see Appendix). Image patches in Metro Dallas were used to evaluate the framework's generalization capability, considering the heterogeneous patterns of land use/land cover revealed from these two study areas.

Hyperparameters were tuned empirically based on the performance of the validating set. Adam optimizer was adopted with a learning rate initialized to 0.0001 ($\beta_1 = 0.9, \beta_2 = 0.999$), the batch size was set to 32, and epochs were capped at 10 000, before which all selected architectures reached a stable performance. The framework run on a desktop with two NVIDIA GTX 1080Ti GPUs, a 3.20 GHz Intel Core i7-8700 CPU, and 32 GB RAM. We implemented the framework using Tensoreflow 2.0 library with Python 3.6.5 under Windows 10, CUDA 10.1, and CUDNN 7.0 system.

D. Evaluation Metrics

We used three common quantitative indices to reveal the discrepancy between p_i^{lg} (ground-truthing) and \widehat{p}_i^{lg} (estimated) in the testing set: coefficient of determination (R^2); coefficient of efficiency (CoE); and modified index of agreement (MIoA). CoE, ranging from minus infinity to 1, suggests the proportion of initial variance accounted for by a model [56]. The higher the CoE value, the better agreement a model reaches. R^2 and MIoA are widely adopted metrics to suggest the general goodness of fit of a model. Specifically, MIoA, less sensitive to the proportional difference compared with R^2 , is a modified version of the original index of agreement proposed by Willmott *et al.* [57]. With a range from 0 to 1, MIoA with a higher value indicates better agreement. The calculations for the three indices follow:

$$R^2 = \left(\frac{\sum_{i=1}^m (p_i^{lg} - \overline{p_i^{lg}}) (\widehat{p}_i^{lg} - \overline{\widehat{p}_i^{lg}})}{\sqrt{\sum_{i=1}^m (p_i^{lg} - \overline{p_i^{lg}})^2} \sqrt{\sum_{i=1}^m (\widehat{p}_i^{lg} - \overline{\widehat{p}_i^{lg}})^2}} \right)^2 \quad (5)$$

$$COE = 1 - \frac{\sum_{i=1}^m (p_i^{lg} - \widehat{p}_i^{lg})^2}{\sum_{i=1}^m (p_i^{lg} - \overline{p_i^{lg}})^2} \quad (6)$$

$$MIoA = 1 - \frac{\sum_{i=1}^m |p_i^{lg} - \widehat{p}_i^{lg}|}{\sum_{i=1}^m \left(|\widehat{p}_i^{lg} - \overline{\widehat{p}_i^{lg}}| + |p_i^{lg} - \overline{p_i^{lg}}| \right)} \quad (7)$$

where m denotes the number of testing patches, $\overline{p_i^{lg}}$ denotes the average value of p_i^{lg} , i.e., $\overline{p_i^{lg}} = \frac{\sum_{i=1}^m p_i^{lg}}{m}$, and $\overline{\widehat{p}_i^{lg}}$ denotes the average value of \widehat{p}_i^{lg} , i.e., $\overline{\widehat{p}_i^{lg}} = \frac{\sum_{i=1}^m \widehat{p}_i^{lg}}{m}$.

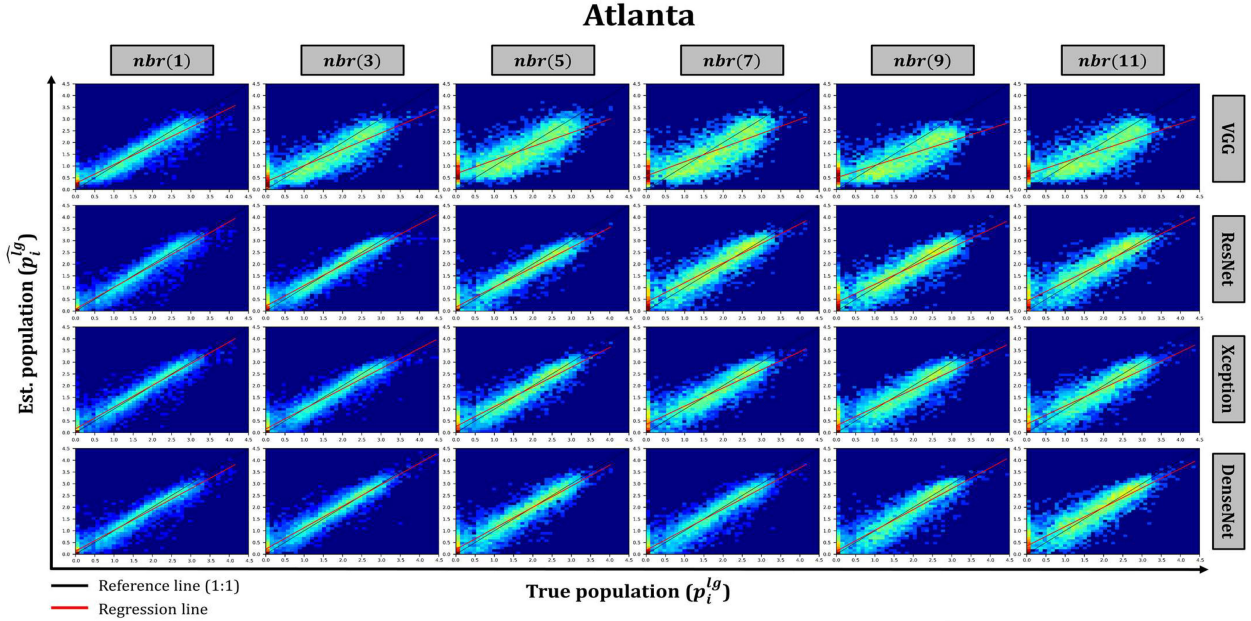


Fig. 4. Scatterplot of Metro Atlanta testing patches between the true population (p_i^{lg}) and the estimated population (\widehat{p}_i^{lg}) from VGG, ResNet, Xception, and DenseNet, under different neighboring scenarios.

V. RESULTS

A. Model Performances of Testing Patches in Metro Atlanta

After the model training process, we applied well-trained models to the testing patches (5700) in Metro Atlanta, aiming to reveal the discrepancies of model performances, investigate the influence of neighboring sizes, and identify the potential systematic biases. The model performances of testing patches in Metro Atlanta under different neighborhood scenarios can be found in Table I. In general, DenseNet stands out, as it outperforms the other three models in all evaluation metrics under all selected neighboring sizes. In comparison, VGG has the worst performance, evidenced by its lowest values in evaluating metrics among the four selected models. ResNet and Xception have a similar performance lying between DenseNet and VGG. The scatterplots of the true population (p_i^{lg}) and the estimated population (\widehat{p}_i^{lg}) (see Fig. 4) also confirm this conclusion, as the scatterplot for DenseNet presents a more clustered pattern in each neighboring scenario compared to the other three models. In comparison, p_i^{lg} and \widehat{p}_i^{lg} are scatteredly distributed from the 1:1 reference line in VGG, suggesting its poor predicting performance.

As for neighboring scenarios, the increase of neighboring sizes leads to reduced population estimation performance, which is found universal for all four selected models in all evaluating metrics (see Table I). For instance, the R^2 for DenseNet with nbr(1) is 0.915, which is gradually reduced when increasing neighboring patches: 0.904 with nbr(3), 0.874 with nbr(5), 0.853 with nbr(7), 0.836 with nbr(9), and 0.811 with nbr(11). Such performance reduction in DenseNet with an increased neighboring size is also supported by CoE and MIoA metrics: CoE gradually reduces from 0.901 to 0.773 and MIoA from 0.877 to 0.791, when neighboring size increases from 1 to 11.

The scatterplots validate this claim, as the p_i^{lg} and \widehat{p}_i^{lg} in all selected models present a more scattered distribution pattern when more neighboring patches are considered (see Fig. 4), suggesting that increasing neighboring sizes, to a certain degree, confuses models, leading to a negative impact on model prediction. The above results contradict the findings by Xing *et al.* [28], who found an improved predicting performance can be achieved with the consideration of more neighboring patches. We assume that such a phenomenon can be attributed to the heterogeneous nature of population distribution and the diminishing proportion of information in the central patch resulting from the requirement of a fixed input size for deep learning models.

To explore whether systematic biases exist in population estimation from remote sensing images via deep learning models, we investigated the relationship between the true population, i.e., p_i^{lg} , and the difference in estimated population and true population, i.e., $\widehat{p}_i^{lg} - p_i^{lg}$. Fig. 5 presents the scatterplots of p_i^{lg} and $\widehat{p}_i^{lg} - p_i^{lg}$ from VGG, ResNet, Xception, and DenseNet, under different neighboring scenarios. The results reveal a notable, universal bias, evidenced by the negative slope (β) and negative Pearson's r values significant at 0.001 significance level from all models under all neighboring scenarios. The results reveal a notable, universal bias, evidenced by the negative slope (β) and negative Pearson's r values significant at 0.001 significance level from all models under all neighboring scenarios (Table III). The above results demonstrate that all selected deep learning models tend to overestimate sparsely populated image patches and underestimate densely populated image patches, regardless of neighboring sizes. The choice of neighboring sizes plays a trivial role in the strength of such bias, as Pearson's r and β fluctuate with the increase of neighboring sizes. Among all the neighboring scenarios from nbr(1) to nbr(11), VGG,

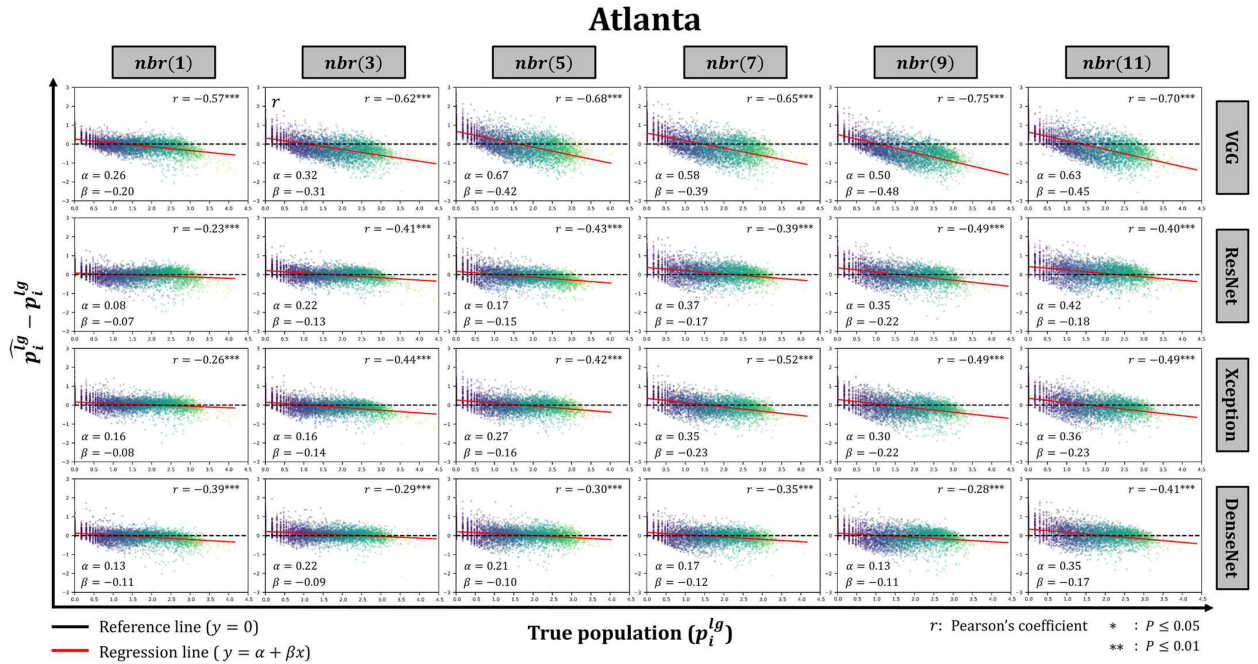


Fig. 5. Scatterplot of Metro Atlanta testing patches between the true population (p_i^{lg}) and the difference in estimated population and true population ($\widehat{p}_i^{lg} - p_i^{lg}$) from VGG, ResNet, Xception, and DenseNet, under different neighboring scenarios.

ResNet, and Xception have the least bias with $\text{nbr}(1)$, i.e., the central patch. For DenseNet, however, the least bias occurs with $\text{nbr}(3)$, i.e., the central patch with its eight nearby neighboring patches. The identified systematic bias can be partly explained by the intrinsic limitation of multispectral remote sensing images: lacking the vertical observation, which presumably leads to the underestimation of holding capacity (volume) of high-rise buildings in densely populated areas.

B. Model Performances in Metro Dallas

Given the observable discrepancies in distribution patterns of land use and land cover in Metro Dallas compared with Metro Atlanta, we evaluated the generalization capability by applying models trained using patches from Metro Atlanta to image patches in Metro Dallas. We aimed to investigate whether similar patterns still hold in Metro Dallas regarding model performances, the impact of neighboring sizes, and the identified systematic bias.

In general, DenseNet still outperforms the other three models in all evaluation metrics under all neighboring scenarios, while VGG has the worst performance (see Table II). Comparing Tables I and II, there exist overall performance reductions, mostly slight ones, for all models when they are applied to a different metropolis. For instance, the MIOA for DenseNet with $\text{nbr}(1)$ in Metro Atlanta is 0.877, but it is reduced to 0.838 in Metro Dallas. Similar patterns can be found for MIOA in DenseNet under other neighboring scenarios: from 0.867 to 0.808 with $\text{nbr}(3)$, from 0.841 to 0.793 with $\text{nbr}(5)$, from 0.831 to 0.768 with $\text{nbr}(7)$, from 0.814 to 0.741 with $\text{nbr}(9)$, and from 0.791 to 0.710 with $\text{nbr}(11)$. The above findings can be confirmed by comparing Figs. 4 and 6, where scatterplots of p_i^{lg} and \widehat{p}_i^{lg} are

presented. Evidently, more scattered distribution of p_i^{lg} and \widehat{p}_i^{lg} along the 1:1 reference line can be found in Metro Dallas (see Fig. 6) comparing to the corresponding ones in Metro Atlanta (see Fig. 4). The above results suggest that models trained in Metro Atlanta can be generalized to Metro Dallas. However, the differences between these two metropolitan areas pose certain challenges for all models in population estimation.

As for neighboring scenarios, the pattern identified from the testing patches in Metro Atlanta still holds: increased neighboring sizes lead to universal reduced population estimation performances for all four selected models, evidenced by gradually reduced values in all evaluating metrics (see Table II) and scattered distribution of p_i^{lg} and \widehat{p}_i^{lg} (see Fig. 6) when neighboring scenarios move from $\text{nbr}(1)$ to $\text{nbr}(11)$.

The relationship between the p_i^{lg} and $\widehat{p}_i^{lg} - p_i^{lg}$ in Dallas reveals that the same bias exists in Metro Dallas (see Fig. 7 and Table III) and the bias is, to some extent, intensified. For instance, the Pearson's r values between $\widehat{p}_i^{lg} - p_i^{lg}$ and p_i^{lg} for DenseNet with $\text{nbr}(1)$ in Metro Atlanta and Metro Dallas are -0.39 and -0.57, respectively, suggesting a strengthened negative correlation. Similar intensified biases can be found for DenseNet under other neighboring scenarios: from -0.29 to -0.68 with $\text{nbr}(3)$, from -0.30 to -0.45 with $\text{nbr}(5)$, from -0.35 to -0.64 with $\text{nbr}(7)$, from -0.28 to -0.50 with $\text{nbr}(9)$, and from -0.41 to -0.60 with $\text{nbr}(11)$. Other models that include VGG, ResNet, and Xception, also experience intensified systematic biases under varying neighboring scenarios when estimating population using image patches in Metro Dallas. The intensified bias (underestimating densely populated image patches) presumably results from the more centralized urban pattern and the denser urban core of Metro Dallas compared with Metro Atlanta (see Fig. 1).

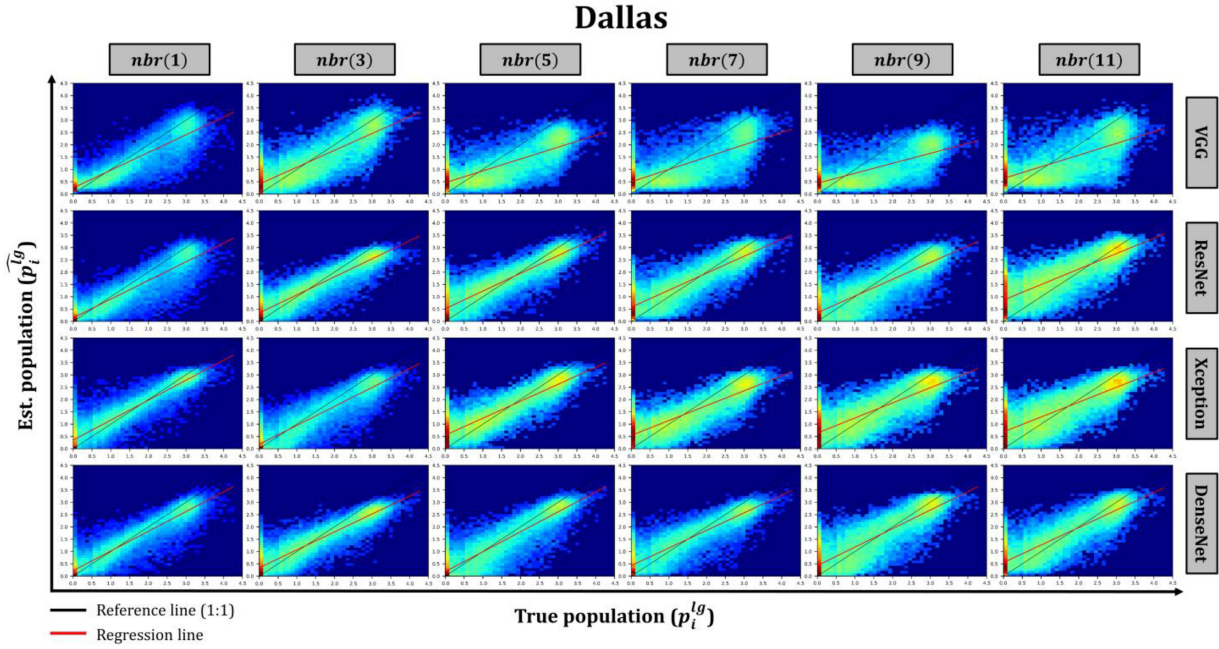


Fig. 6. Scatterplot of all Metro Dallas patches between the true population (p_i^{lg}) and the estimated population ($p_i^{\hat{lg}}$) from VGG, ResNet, Xception, and DenseNet, under different neighboring scenarios.

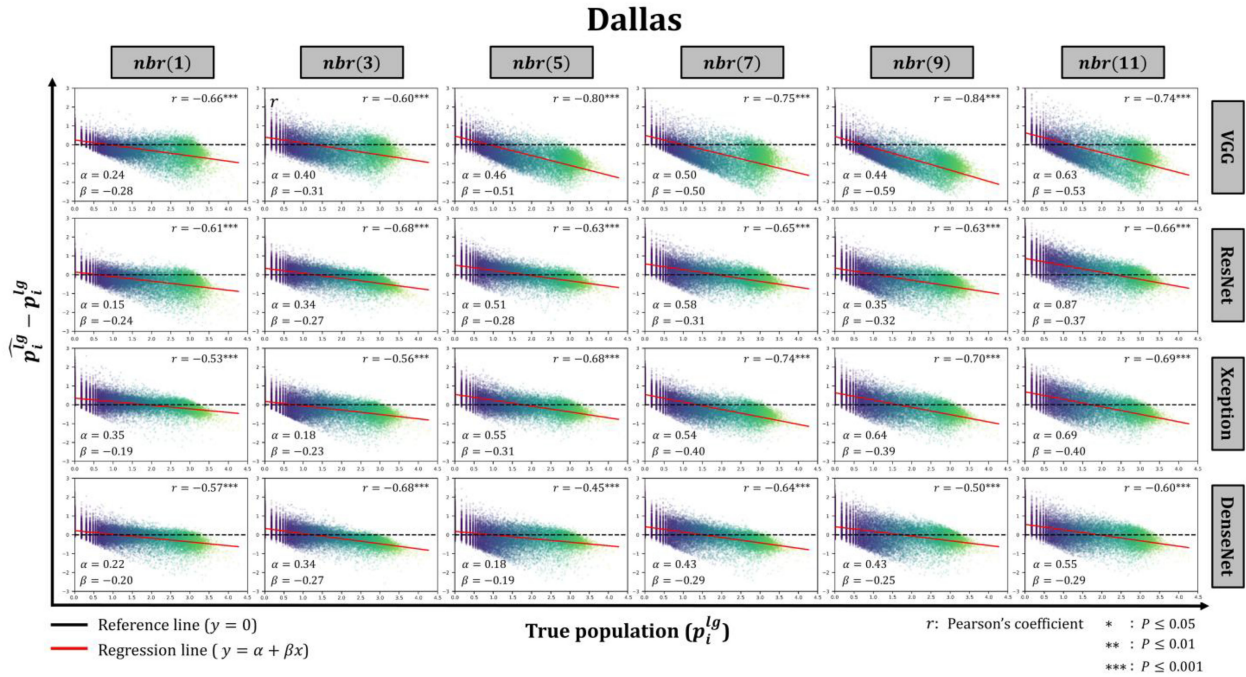


Fig. 7. Scatterplot of all Metro Dallas patches between the true population (p_i^{lg}) and the difference in estimated population and true population ($p_i^{\hat{lg}} - p_i^{lg}$) from VGG, ResNet, Xception, and DenseNet, under different neighboring scenarios.

Given that all models achieve the best performance under the neighboring scenario of $nbr(1)$, we present the estimated population distribution in Metro Dallas with $nbr(1)$ via DenseNet [see Fig. 8(b)], ResNet [see Fig. 8(c)], Xception [see Fig. 8(d)], VGG [see Fig. 8(e)], as well as ground-truthing population distribution [see Fig. 8(a)]. In general, estimated population distributions from all models present a considerably similar

pattern to that of the true population distribution, suggesting their overall capability of capturing the heterogeneity of population distribution from remote sensing images. However, The issue of population underestimation in urban areas is evident for all models, especially in the urban core (highlighted by the purple rectangle), where densely populated grids (dark red grids) are markedly underestimated. Our investigation reveals that the

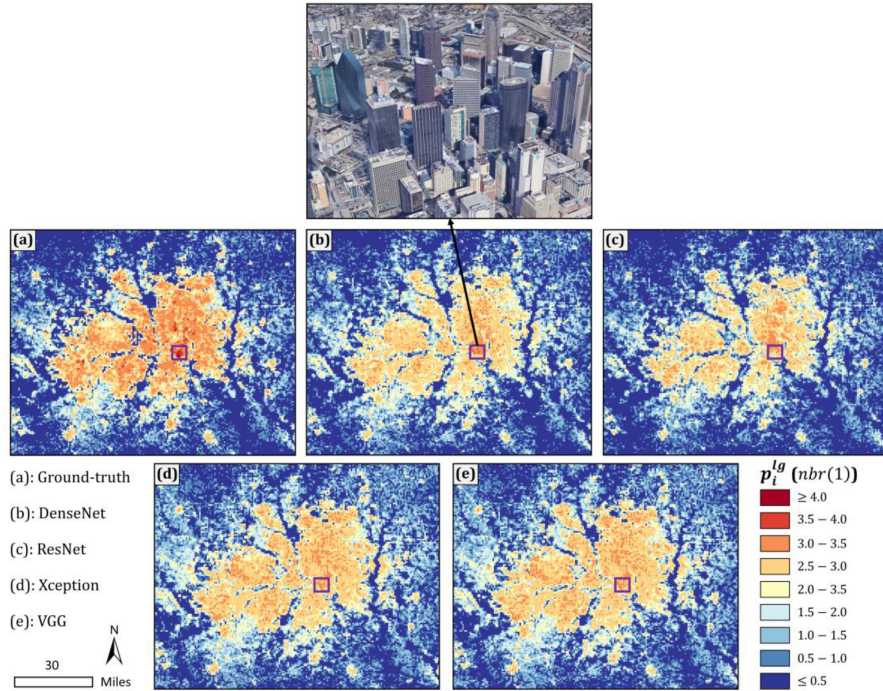


Fig. 8. Ground-truthing population. (a) Population estimated via DenseNet. (b) ResNet. (c) Xception. (d) VGG (e) under the neighboring scenario of $nbr(1)$. 3-D imagery covered by the purple rectangle was derived from Google map.

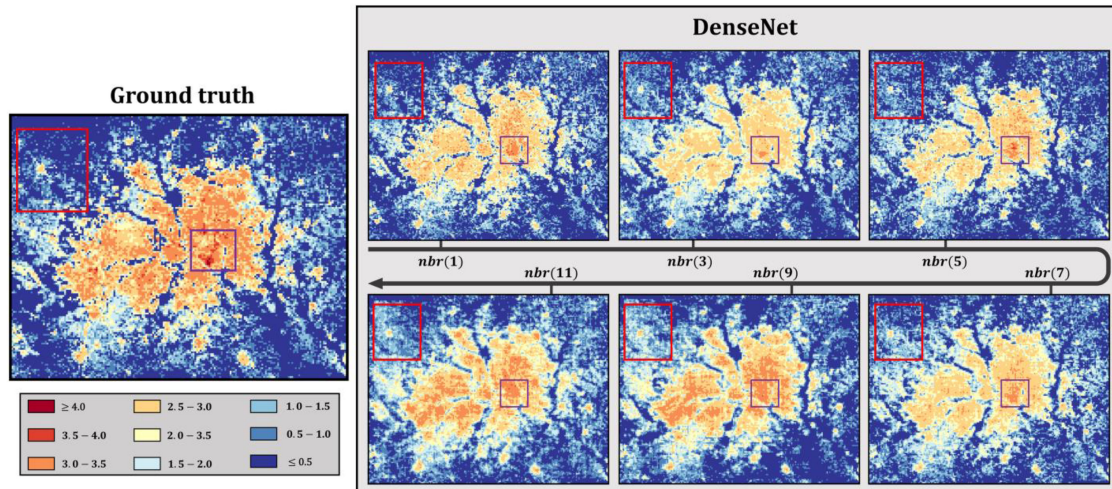


Fig. 9. Ground-truthing population and population estimated via DenseNet under different neighboring scenarios: $nbr(1)$; $nbr(3)$; $nbr(5)$; $nbr(7)$; $nbr(9)$; and $nbr(11)$.

underestimated grids in the urban core contain many high-rise buildings (see Fig. 8). As remote sensing imagery (Sentinel-2 images in this study) fails to obtain information in the vertical dimension, the holding capacity of these high-rise buildings is significantly underestimated, responsible for the systematic underestimation of population distribution in this region from all deep learning models. The estimated population from the best model, i.e., DenseNet, under different neighboring scenarios are presented in Fig. 9. It can be observed that increased neighboring sizes lead to blurring estimations, which is expected as the consideration of nearby neighbors essentially applies a “filter” by downplaying the importance of the central image patch, thus

diminishing the grid-level heterogeneity. We can also conclude that smaller neighboring sizes are able to retain fine details in sparsely populated areas (see areas highlighted by red rectangles in Fig. 9).

VI. DISCUSSION

Obtaining fine-grained population distribution is of great importance to a variety of fields that demand such spatial knowledge. The rapid development of remote sensing techniques provides rich, large-coverage, and high-temporal information of the ground, which can be coupled with the emerging deep

TABLE I
MODEL PERFORMANCES OF TESTING PATCHES IN METRO ATLANTA UNDER DIFFERENT NEIGHBORING SCENARIOS

Evaluating metrics	VGG	ResNet	Xception	DenseNet
R^2 (0~1)				
<i>nbr</i> (1)	0.875	0.906	0.911	0.915
<i>nbr</i> (3)	0.743	0.905	0.894	0.904
<i>nbr</i> (5)	0.617	0.866	0.856	0.874
<i>nbr</i> (7)	0.606	0.807	0.808	0.853
<i>nbr</i> (9)	0.546	0.803	0.788	0.836
<i>nbr</i> (11)	0.534	0.773	0.760	0.811
CoE ($-\infty$ ~1)				
<i>nbr</i> (1)	0.827	0.902	0.905	0.901
<i>nbr</i> (3)	0.590	0.887	0.870	0.895
<i>nbr</i> (5)	0.371	0.837	0.838	0.863
<i>nbr</i> (7)	0.350	0.771	0.747	0.857
<i>nbr</i> (9)	0.130	0.741	0.723	0.825
<i>nbr</i> (11)	0.105	0.730	0.707	0.773
MloA (0~1)				
<i>nbr</i> (1)	0.830	0.871	0.865	0.877
<i>nbr</i> (3)	0.738	0.862	0.852	0.867
<i>nbr</i> (5)	0.672	0.837	0.832	0.841
<i>nbr</i> (7)	0.666	0.801	0.792	0.831
<i>nbr</i> (9)	0.617	0.791	0.785	0.814
<i>nbr</i> (11)	0.589	0.781	0.769	0.791

TABLE II
MODEL PERFORMANCES IN METRO DALLAS UNDER DIFFERENT NEIGHBORING SCENARIOS

Evaluating metrics	VGG	ResNet	Xception	DenseNet
R^2 (0~1)				
<i>nbr</i> (1)	0.817	0.830	0.835	0.878
<i>nbr</i> (3)	0.732	0.841	0.820	0.841
<i>nbr</i> (5)	0.563	0.778	0.756	0.802
<i>nbr</i> (7)	0.546	0.742	0.711	0.791
<i>nbr</i> (9)	0.476	0.742	0.654	0.744
<i>nbr</i> (11)	0.452	0.674	0.624	0.722
CoE ($-\infty$ ~1)				
<i>nbr</i> (1)	0.702	0.748	0.808	0.831
<i>nbr</i> (3)	0.574	0.748	0.747	0.765
<i>nbr</i> (5)	0.210	0.653	0.626	0.762
<i>nbr</i> (7)	0.150	0.572	0.510	0.667
<i>nbr</i> (9)	0.072	0.532	0.427	0.624
<i>nbr</i> (11)	-0.312	0.436	0.396	0.601
MloA (0~1)				
<i>nbr</i> (1)	0.785	0.812	0.825	0.838
<i>nbr</i> (3)	0.732	0.805	0.801	0.808
<i>nbr</i> (5)	0.621	0.776	0.767	0.793
<i>nbr</i> (7)	0.612	0.749	0.714	0.768
<i>nbr</i> (9)	0.523	0.731	0.685	0.741
<i>nbr</i> (11)	0.501	0.679	0.648	0.710

learning approaches that enable latent features and hidden geographical patterns to be extracted. In this article, we establish an end-to-end framework to evaluate the performances of four popular deep learning architectures, VGG, ResNet, Xception, and DenseNet, in estimating population distribution directly from remote sensing image patches via transfer learning and fine-tuning techniques. In addition, we conduct a thorough investigation on the neighboring effects and the potential systematic biases in population estimation using remote sensing images via deep learning approaches.

Our results reveal that in all evaluation metrics under all selected neighboring sizes, DenseNet outperforms the other three models, while VGG has the worst performances. ResNet and Xception have similar performance, lying between DenseNet and VGG. We believe the superior performance of DenseNet can be attributed to its architecture design. Layers in DenseNet receive all preceding layers as input, thus creating diversified

features with richer patterns and facilitating the acquisition of “collective knowledge.” These diversified features and patterns assist in capturing the hidden mapping between image patches and gridded population, leading to more accurate population estimation. In addition, thanks to the channel-wise concatenation throughout its Dense Block, DenseNet maintains both low- and high-level complexity features (unlike the standard stacks of ConvNets that use mostly high-level features). The capability of DenseNet in handling multi-level features contributes to the better summarization of the heterogeneity in population distribution. In comparison, the poor performance of VGG is expected, as its deep architecture formed by plainly stacking simple convolutional operations fails to promote multi-level complexity and diversity.

In terms of the neighboring effect, our results suggest that the increase of neighboring sizes leads to reduced population estimation performance, which is found universal for all four selected models, in all evaluating metrics, and in both Metro Atlanta and Metro Dallas, contradicting a recent study by Xing *et al.* [28], who found 3×3 neighboring scenario outperformed 1×1 via ResNet architecture. Two possible reasons are responsible for the above findings. First, numerous studies have proved that population, as a fundamental agent in urban and suburban ecosystems, is distributed with great heterogeneity [58], [59]. For instance, a densely populated image patch can be surrounded by uninhabited patches, such as green space and water bodies. Such neighboring information does not necessarily benefit population estimation in the central patch but introduces certain uncertainty (noises) to the model prediction. Second, for the selected four deep learning models, as well as for most existing deep learning models, a fixed input shape of images is required. Considering neighboring patches unavoidably leads to a diminishing proportion of information in the central patch of an input image to the model. For instance, the proportion of information contained in the central patch with a neighboring scenario of *nbr*(1), i.e., 1×1 , is 100%. This proportion reduces to 11.11% with a neighboring scenario of *nbr*(3), i.e., 3×3 , and to 0.83% with *nbr*(11), i.e., 11×11 . Such diminishing proportion of information in the central patch with the increase of neighboring consideration poses challenges for models to predict the population residing in the central patch.

In terms of biases, our results reveal a notable, universal bias for all models under all neighboring scenarios, indicating that all selected deep learning models tend to overestimate sparsely populated image patches and underestimate densely populated image patches, regardless of neighboring sizes. This systematic bias can be attributed to the limitation of multi-spectral remote sensing sensors that lack the ability to obtain information in the vertical dimension. In densely populated urban centers, there exist a large number of high-rise buildings whose holding capacity cannot be properly measured without the vertical information (i.e., building height), as they tend to hold an unexpectedly larger population than their building footprint size has suggested. Efforts have been made to extract the vertical dimension of buildings from light detection and

TABLE III
CORRELATION BETWEEN THE p_i^{lg} AND $\widehat{p_i^{lg}} - p_i^{lg}$ IN METRO ATLANTA AND METRO DALLAS UNDER DIFFERENT NEIGHBORING SCENARIOS

	VGG		ResNet		Xception		DenseNet	
	Atlanta	Dallas	Atlanta	Dallas	Atlanta	Dallas	Atlanta	Dallas
Pearson's r								
$nbr(1)$	-0.57	-0.66	-0.23	-0.61	-0.26	-0.53	-0.39	-0.57
$nbr(3)$	-0.62	-0.60	-0.41	-0.68	-0.44	-0.56	-0.29	-0.68
$nbr(5)$	-0.68	-0.80	-0.43	-0.63	-0.42	-0.68	-0.30	-0.45
$nbr(7)$	-0.65	-0.75	-0.39	-0.65	-0.52	-0.74	-0.35	-0.64
$nbr(9)$	-0.75	-0.84	-0.49	-0.63	-0.49	-0.70	-0.28	-0.50
$nbr(11)$	-0.70	-0.74	-0.40	-0.66	-0.49	-0.69	-0.41	-0.60
α								
$nbr(1)$	0.26	0.24	0.08	0.15	0.16	0.35	0.13	0.22
$nbr(3)$	0.32	0.40	0.22	0.34	0.16	0.18	0.22	0.34
$nbr(5)$	0.67	0.46	0.17	0.51	0.27	0.55	0.21	0.18
$nbr(7)$	0.58	0.50	0.37	0.58	0.35	0.54	0.17	0.43
$nbr(9)$	0.50	0.44	0.35	0.35	0.30	0.64	0.13	0.43
$nbr(11)$	0.63	0.63	0.42	0.87	0.36	0.69	0.35	0.55
β								
$nbr(1)$	-0.20	-0.28	-0.07	-0.24	-0.08	-0.19	-0.11	-0.20
$nbr(3)$	-0.31	-0.31	-0.13	-0.27	-0.14	-0.23	-0.09	-0.27
$nbr(5)$	-0.42	-0.51	-0.15	-0.28	-0.16	-0.31	-0.10	-0.19
$nbr(7)$	-0.39	-0.50	-0.17	-0.31	-0.23	-0.40	-0.12	-0.29
$nbr(9)$	-0.48	-0.59	-0.22	-0.32	-0.22	-0.39	-0.11	-0.25
$nbr(11)$	-0.45	-0.53	-0.18	-0.37	-0.23	-0.40	-0.17	-0.29

All Pearson's r values are significant at a 0.001 significance level.

ranging aerial photogrammetry, and synthetic aperture radar [60], [61]. Future studies need to incorporate the building height information to better estimate building volume, potentially leading to better population estimation in densely populated urban areas.

Finally, we need to acknowledge several limitations of this article. First, we only compared the performances of four widely-adopted deep learning architectures with diverging concepts in design: VGG; ResNet; Xception; and DenseNet. Despite their popularity, future efforts are needed to explore the performances of other deep learning architectures in population estimation from remote sensing images. Second, we chose to estimate population distribution at 30 arc-second grids (approximately 1 km at the equator), a spatial resolution adopted by many studies similar to this article [27]–[29]. However, we acknowledge that the heterogeneity of population distribution tends to vary in different scales. Thus, the conclusion in this study may not hold for population estimation at a different geographical scale (e.g., in sub-km population estimation). Future studies are needed to investigate model performances, neighboring effects, and potential biases by adopting a multi-scale estimation framework. Third, we trained (17 100 patches), validated (5700 patches), and tested (5600 patches) deep learning models in Metro Atlanta and further applied them to Metro Dallas (28 500 patches) to evaluate their overall generalizability. As the performance of deep learning models largely depends on the size of training samples, we encourage future studies to involve more training samples from multiple study sites for better generalization capability. Finally, we established a mapping between remote sensing imagery and population distribution using a band composition of R, G, B, and NIR. However, it is reasonable to assume that involving other bands will benefit the recognition of land use and land cover types, potentially leading to improved population distribution estimation. Thus, we encourage more investigations on the spectral domain of remote sensing images in population estimation.

VII. CONCLUSION

The coupling of remote sensing imagery and deep learning algorithms undoubtedly establishes a new venue that potentially advances traditional population modeling. This article marks the first attempt to cross-compare performances of popular state-of-the-art deep learning models in estimating population distribution from remote sensing images, investigate the contribution of neighboring effects, and explore the potential systematic population estimation biases.

The results suggest that DenseNet outperforms the other three models, while VGG has the worst performances in all evaluating metrics under all selected neighboring scenarios. ResNet and Xception have similar performance, lying between DenseNet, and VGG. The superior performance of DenseNet is presumably due to its architectural design that facilitates diversified features and patterns for capturing hidden mapping between image patches and gridded population. As for the neighboring effect, our results indicate that the increase of neighboring sizes leads to reduced population estimation performance, which is found universal for all four selected models, in all evaluating metrics, and in both Metro Atlanta and Metro Dallas, which contradicts a recent study by Xing *et al.* [28]. Such a phenomenon can be attributed to the heterogeneous nature of population distribution and the diminishing proportion of information in the central patch resulting from the requirement of a fixed input size for deep learning models. In addition, there exists a notable, universal bias that all selected deep learning models tend to overestimate sparsely populated image patches and underestimate densely populated image patches, regardless of neighboring sizes. This systematic bias can be explained by the intrinsic limitation of multispectral remote sensing images: lacking the vertical observation, thus leading to underestimating the holding capacity of high-rise buildings in densely populated urban cores. The methodological, experimental, and contextual knowledge this study provides is expected to benefit a wide range of future studies that estimate population distribution via remote sensing imagery.

APPENDIX

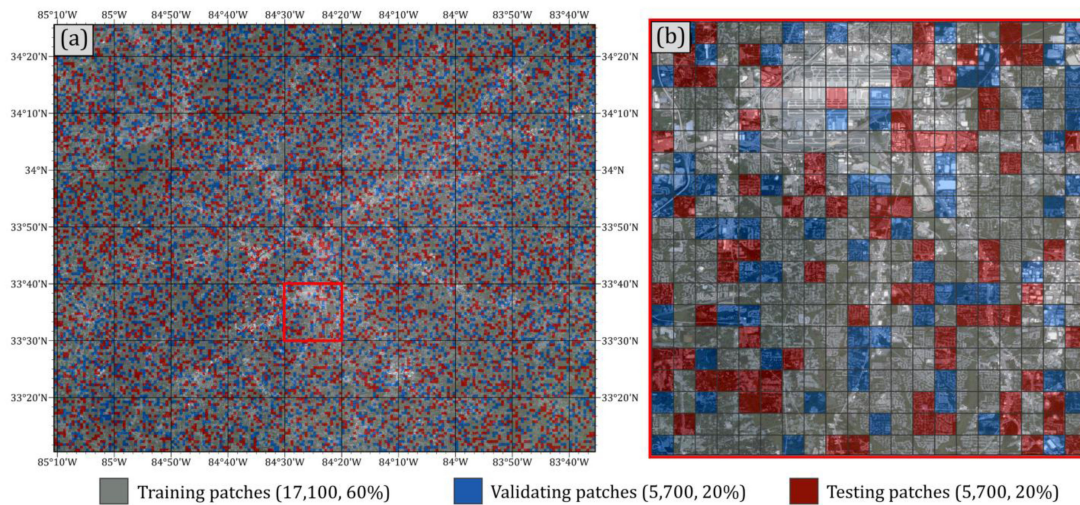


Fig. 10. Spatial distribution of Metro Atlanta's image patches in the training set (60%), validating set (20%), and testing set (20%).

REFERENCES

- [1] S. I. Hay, A. M. Noor, A. Nelson, and A. J. Tatem, "The accuracy of human population maps for public health application," *Tropical Med. Int. Health*, vol. 10, no. 10, pp. 1073–1086, 2005.
- [2] K. L. Frohlich, and L. Potvin, "Transcending the known in public health practice: The inequality paradox: The population approach and vulnerable populations," *Amer. J. Public Health*, vol. 98, no. 2, pp. 216–221, 2008.
- [3] A. J. Tatem, A. M. Noor, and S. I. Hay, "Defining approaches to settlement mapping for public health management in Kenya using medium spatial resolution satellite imagery," *Remote Sens. Environ.*, vol. 93, no. 1/2, pp. 42–52, 2004.
- [4] M. Langford, G. Higgs, J. Radcliffe, and S. White, "Urban population distribution models and service accessibility estimation," *Comput. Environ. Urban Syst.*, vol. 32, no. 1, pp. 66–80, 2008.
- [5] J. A. Maantay, A. R. Maroko, and C. Herrmann, "Mapping population distribution in the urban environment: The cadastral-based expert dasymetric system (CEDS)," *Cartogr. Geography Inf. Sci.*, vol. 34, no. 2, pp. 77–102, 2007.
- [6] C. Kang, Y. Liu, X. Ma, and L. Wu, "Towards estimating urban population distributions from mobile call data," *J. Urban Technol.*, vol. 19, no. 4, pp. 3–21, 2012.
- [7] P. Tenerelli, J. F. Gallego, and D. Ehrlich, "Population density modelling in support of disaster risk assessment," *Int. J. Disaster Risk Reduction*, vol. 13, pp. 334–341, 2015.
- [8] X. Song, Q. Zhang, Y. Sekimoto, T. Horanont, S. Ueyama, and R. Shibasaki, "Modeling and probabilistic reasoning of population evacuation during large-scale disaster," in *Proc. 19th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2013, pp. 1231–1239.
- [9] C. Linard, M. Gilbert, R. W. Snow, A. M. Noor, and A. J. Tatem, "Population distribution, settlement patterns and accessibility across Africa in 2010," *PLoS One*, vol. 7, no. 2, 2012, Art. no. e31743.
- [10] D. Murakami, and Y. Yamagata, "Estimation of gridded population and GDP scenarios with spatially explicit statistical downscaling," *Sustainability*, vol. 11, no. 7, p. 2106, 2019, doi: [10.3390/su11072106](https://doi.org/10.3390/su11072106).
- [11] J. Shen, "Internal migration and regional population dynamics in China," *Prog. Planning*, vol. 45, no. 3, pp. 123–188, 1996.
- [12] C. Tacoli, "Crisis or adaptation? Migration and climate change in a context of high mobility," *Environ. Urbanization*, vol. 21, no. 2, pp. 513–525, 2009.
- [13] X. Huang, C. Wang, Z. Li, and H. Ning, "A 100m population grid in the CONUS by disaggregating census data with open-source microsoft building footprints," *Big Earth Data*, vol. 5, no. 1, pp. 112–133, 2020.
- [14] A. S. Fotheringham, and D. W. S. Wong, "The modifiable areal unit problem in multivariate statistical analysis," *Environ. Planning A*, vol. 23, no. 7, pp. 1025–1044, 1991.
- [15] Z. Lu, J. Im, L. Quackenbush, and K. Halligan, "Population estimation based on multi-sensor data fusion," *Int. J. Remote Sens.*, vol. 31, no. 21, pp. 5587–5604, 2010.
- [16] N. A. Wardrop *et al.*, "Spatially disaggregated population estimates in the absence of national population and housing census data," *Proc. Nat. Acad. Sci. USA*, vol. 115, no. 14, pp. 3529–3537, 2018.
- [17] X. Huang, C. Wang, and J. Lu, "Understanding spatiotemporal development of human settlement in hurricane-prone areas on U.S. Atlantic and gulf coasts using nighttime remote sensing," *Nat. Hazards Earth Syst. Sci. Discuss.*, vol. 17, no. 10, pp. 2141–2155, 2019.
- [18] C. L. Eicher, and C. A. Brewer, "Dasymetric mapping and areal interpolation: Implementation and evaluation," *Cartogr. Geogr. Inf. Sci.*, vol. 28, no. 2, pp. 125–138, 2001.
- [19] F. Batista e Silva, J. Gallego, and C. Lavalle, "A high-resolution population grid map for Europe," *J. Maps*, vol. 9, no. 1, pp. 16–28, 2013.
- [20] F. R. Stevens, A. E. Gaughan, C. Linard, and A. J. Tatem, "Disaggregating census data for population mapping using random forests with remotely-sensed and ancillary data," *PLoS One*, vol. 10, no. 2, 2015, Art. no. e0107042.
- [21] D. L. Balk, U. Deichmann, G. Yetman, F. Pozzi, S. I. Hay, and A. Nelson, "Determining global population distribution: Methods, applications and data," *Adv. Parasitol.*, vol. 62, pp. 119–156, 2006.
- [22] J. Mennis, "Dasymetric mapping for estimating population in small areas: Dasymetric mapping for estimating population in small areas," *Geogr. Compass*, vol. 3, no. 2, pp. 727–745, 2009.
- [23] X. X. Zhu *et al.*, "Deep learning in remote sensing: A comprehensive review and list of resources," *IEEE Geosci. Remote Sens. Mag.*, vol. 5, no. 4, pp. 8–36, Dec. 2017.
- [24] Q. Zou, L. Ni, T. Zhang, and Q. Wang, "Deep learning based feature selection for remote sensing scene classification," *IEEE Geosci. Remote Sens. Lett.*, vol. 12, no. 11, pp. 2321–2325, Nov. 2015.
- [25] J. E. Ball, D. T. Anderson, and C. S. Chan, "Comprehensive survey of deep learning in remote sensing: Theories, tools, and challenges for the community," *J. Appl. Remote Sens.*, vol. 11, no. 4, p. 042609, 2017.
- [26] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
- [27] W. Hu *et al.*, "Mapping missing population in rural india: A deep learning approach with satellite imagery," in *Proc. 2019 AAAI/ACM Conf. AI, Ethics, Soc.*, 2019, pp. 353–359.
- [28] X. King *et al.*, "Mapping human activity volumes through remote sensing imagery," *IEEE J. Sel. Top. Appl. Earth Observ. Remote Sens.*, vol. 13, pp. 5652–5668, Sep. 14, 2020, doi: [10.1109/JSTARS.2020.3023730](https://doi.org/10.1109/JSTARS.2020.3023730).

- [29] C. Robinson, F. Hohman, and B. Dilkina, "A deep learning approach for population estimation from satellite imagery," in *Proc. 1st ACM SIGSPATIAL Workshop Geospatial Humanities*, 2017, pp. 47–54.
- [30] P. Doupe, E. Bruzelius, J. Faghmous, and S. G. Ruchman, "Equitable development through deep learning: The case of sub-national population density estimation," in *Proc. 7th Annu. Symp. Comput. Develop.*, 2016, pp. 1–10.
- [31] E. Doxsey-Whitfield *et al.*, "Taking advantage of the improved availability of census data: A first look at the gridded population of the world, version 4," *Papers Appl. Geogr.*, vol. 1, no. 3, pp. 226–234, 2015.
- [32] "GHS-POP R2015A - Datasets - EU data portal - Europa Eu," Accessed: Mar./Feb., 2021. [Online]. Available: https://data.europa.eu/euodp/en/data/dataset/jrc-ghsl-ghs_pop_gpww4_globe_r2015a.
- [33] C. Frye, E. Nordstrand, D. J. Wright, C. Terborgh, and J. Foust, "Using classified and unclassified land cover data to estimate the footprint of human settlement," *Data Sci. J.*, vol. 17, p. 20, 2018, doi: [10.5334/dsj-2018-020](https://doi.org/10.5334/dsj-2018-020).
- [34] B. Bhaduri, E. Bright, P. Coleman, and M. L. Urban, "LandScan USA: A high-resolution geospatial and temporal modeling approach for population distribution and dynamics," *GeoJ.*, vol. 69, no. 1/2, pp. 103–117, 2007.
- [35] S. Eichhorn, "Disaggregating population data and evaluating the accuracy of modeled high-resolution population distribution—The case study of Germany," *Sustainability*, vol. 12, no. 10, p. 3976, 2020, doi: [10.3390/su12103976](https://doi.org/10.3390/su12103976).
- [36] A. Vouloimos, N. Doulamis, A. Doulamis, and E. Protopapadakis, "Deep learning for computer vision: A brief review," *Comput. Intell. Neurosci.*, vol. 2018, 2018, Art. no. 7068349.
- [37] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, no. 6, pp. 84–90, 2017.
- [38] K. Simonyan, and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014, *arXiv:1409.1556*.
- [39] C. Szegedy *et al.*, "Going deeper with convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1–9.
- [40] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 2818–2826.
- [41] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. Alemi, "Inception-v4, Inception-resnet and the impact of residual connections on learning," *Proc. AAAI Conf. Artif. Intell.* vol. 31, no. 1, 2017.
- [42] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1251–1258.
- [43] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [44] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 4700–4708.
- [45] A. Ahmed, K. Yu, W. Xu, Y. Gong, and E. Xing, "Training hierarchical feed-forward visual recognition models using transfer learning from pseudo-tasks," in *Lecture Notes in Computer Science*, Berlin, Germany: Springer, 2008, pp. 69–82.
- [46] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?," in *Proc. 27th Int. Conf. Neural Inf. Process. Syst.*, vol. 2, pp. 3320–3328, 2014.
- [47] "Oak ridge national laboratory," Oak Ridge Nat. Lab., Oak Ridge, TN, USA. Accessed: Feb./Mar. 2021. [Online]. Available: <http://web.ornl.gov/>
- [48] J. E. Dobson, E. A. Bright, P. R. Coleman, R. C. Durfee, and B. A. Worley, "LandScan: A global population database for estimating populations at risk," *Photogramm. Eng. Remote Sens.*, vol. 69, no. 7, pp. 849–857, 2000.
- [49] P. P. Simarro *et al.*, "Estimating and mapping the population at risk of sleeping sickness," *PLoS Negl. Tropical Dis.*, vol. 6, no. 10, 2012, Art. no. e1859.
- [50] A. Smith, P. D. Bates, O. Wing, C. Sampson, N. Quinn, and J. Neal, "New estimates of flood exposure in developing countries using high-resolution population data," *Nat. Commun.*, vol. 10, no. 1, pp. 1–7, 2019, doi: [10.1038/s41467-019-09282-y](https://doi.org/10.1038/s41467-019-09282-y).
- [51] "LandScan USA," *Arcgis.com*, 2020. Accessed: Feb./Mar. 2021. [Online]. Available: <https://hifld-geoplatform.opendata.arcgis.com/datasets/e431a6410145450aa56606568345765b>
- [52] M. Drusch *et al.*, "Sentinel-2: ESA's optical high-resolution mission for GMES operational services," *Remote Sens. Environ.*, vol. 120, pp. 25–36, 2012.
- [53] M. Herold, D. A. Roberts, M. E. Gardner, and P. E. Dennison, "Spectrometry for urban area remote sensing—Development and analysis of a spectral library from 350 to 2400 nm," *Remote Sens. Environ.*, vol. 91, no. 3/4, pp. 304–319, 2004.
- [54] Z. Jiang *et al.*, "Analysis of NDVI and scaled difference vegetation index retrievals of vegetation fraction," *Remote Sens. Environ.*, vol. 101, no. 3, pp. 366–378, 2006.
- [55] "Log hyperbolic cosine loss improves variational auto-encoder," Accessed: Mar./Feb., 2021 [Online]. Available: <https://openreview.net/forum?id=mrklvsC9Ym>
- [56] A. N. Mandeville, P. E. O'Connell, J. V. Sutcliffe, and J. E. Nash, "River flow forecasting through conceptual models part III - The Ray catchment at grendon underwood," *J. Hydrol.*, vol. 11, no. 2, pp. 109–128, 1970.
- [57] C. J. Willmott *et al.*, "Statistics for the evaluation and comparison of models," *J. Geophys. Res.*, vol. 90, no. C5, pp. 8995, 1985.
- [58] D. Azar *et al.*, "Spatial refinement of census population distribution using remotely sensed estimates of impervious surfaces in haiti," *Int. J. Remote Sens.*, vol. 31, no. 21, pp. 5635–5655, 2010.
- [59] X. Huang, C. Wang, and Z. Li, "High-resolution population grid in the CONUS using microsoft building footprints: A feasibility study," in *Proc. 3rd ACM SIGSPATIAL Int. Workshop Geospatial Humanities*, 2019, pp. 1–9, doi: [10.1145/3356991.3365469](https://doi.org/10.1145/3356991.3365469).
- [60] C. Stal, F. Tack, P. De Maeyer, A. De Wulf, and R. Goossens, "Airborne photogrammetry and lidar for DSM extraction and 3D change detection over an urban area – A comparative study," *Int. J. Remote Sens.*, vol. 34, no. 4, pp. 1087–1110, 2013.
- [61] D. Brunner, G. Lemoine, L. Bruzzone, and H. Greidanus, "Building height retrieval from VHR SAR imagery based on an iterative simulation and matching technique," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 3, pp. 1487–1504, Mar. 2010.



Xiao Huang received the Ph.D. degree in geography from the University of South Carolina, Columbia, SC, USA, in 2020.

He is currently an Assistant Professor with the Department of Geosciences, University of Arkansas, Fayetteville, AR, USA. His research interests include GeoAI, Big Data, natural hazards, data science/visualization, and remote sensing.



Di Zhu received the Ph.D. degree in cartology and GIS from Peking University, Beijing, China.

He is currently an Assistant Professor with the Department of Geography, Environment and Society, University of Minnesota, Twin Cities, Minneapolis, MN, USA. His research interests include GeoAI, spatial analytics, social sensing, and urban complexities.



Fan Zhang received the Ph.D. degree from the Chinese University of Hong Kong, Hong Kong, in 2017

He is currently a Postdoctoral Research Associate at Senseable City Lab, Massachusetts Institute of Technology. His research interests include spatiotemporal data mining, urban visual intelligence, and data-driven approaches for urban studies.



Tao Liu received the Ph.D. degree in geomatics from the University of Florida, Gainesville, FL, USA.

He was a Researcher with the Oak Ridge National Laboratory, Oak Ridge, TN, USA. He is currently an Assistant Professor with Michigan Technological University, Houghton, MI, USA, with expertise in remote sensing, and in using artificial intelligent (AI) techniques and remote sensing techniques to solve regression or classification problems to address issues in sustainability development in urban environment and natural systems.



Lei Zou received the Ph.D. degree in environmental science from Louisiana State University, Baton Rouge, LA, USA.

He is currently an Assistant Professor with the Department of Geography, Texas A&M University, College Station, TX, USA. His research interests include GIScience, social sensing, spatial big data analysis and modeling, disaster resilience, and coastal sustainability.



Xiao Li received the Graduate degree in 2019 from Texas A&M University, College Station, TX, USA, in 2019 and the Ph.D. degree in geography.

He is currently an Associate Transportation Researcher with Texas A&M Transportation Institute, Houston, TX, USA. His research interests include geospatial machine learning in crash analysis, imagery analysis for transportation facilities detection, and individualized traffic safety estimation through mobile data mining.