Deep Landscape Features for Improving Vector-borne Disease Prediction

Nabeel Abdur Rehman
New York University
nabeel@nyu.edu

Umar Saif
UNESCO Chair for ICTD
umar@csail.mit.edu

Rumi Chunara
New York University
rumi.chunara@nyu.edu

Abstract

The global population at risk of mosquito-borne diseases such as dengue, yellow fever, chikungunya and Zika is expanding. Infectious disease models commonly incorporate environmental measures like temperature and precipitation. Given increasing availability of high-resolution satellite imagery, here we consider including landscape features from satellite imagery into infectious disease prediction models. To do so, we implement a Convolutional Neural Network (CNN) model trained on Imagenet data and labelled landscape features in satellite data from London. We then incorporate landscape features from satellite image data from Pakistan, labelled using the CNN, in a well-known Susceptible-Infectious-Recovered epidemic model, alongside dengue case data from 2012-2016 in Pakistan. We study improvement of the prediction model for each of the individual landscape features, and assess the feasibility of using image labels from a different place. We find that incorporating satellite-derived landscape features can improve prediction of outbreaks, which is important for proactive and strategic surveillance and control programmes.

1. Introduction

Vector-borne diseases cause more than 700,000 deaths every year globally. Mosquitoes are the best known disease vector, others include ticks, flies, and fleas [29]. Dengue virus is the most ubiquitous mosquito-transmitted disease, and is transmitted primarily by Aedes aegypti mosquitoes, a vector which also transmits Zika, chikungunya and yellow fever [22]. The virus disproportionately affects urban areas in developing countries, which often have limited resources for containment and intervention activities [8].

Spatial patterns of land use and land cover play an important role in infectious disease dynamics. For example, landscape alterations, such as residential developments or agricultural land use, may result in higher contact rates between humans and disease vectors [23]. Other types of land cover may relate to disease prevalence as they can create favorable conditions for vectors and/or hosts. On the other hand,

human-induced landscape alterations may induce habitat loss, species extinction, altered nutrient dynamics, invasive species colonization, and other ecosystem processes leading to a change in disease incidence rates [11].

Remote sensing data is frequently utilized for land use/land cover classification. Organizations such as the National Oceanic and Atmospheric Administration and the U.S. Geological Survey also publish landcover data, which generally use a mix of spectral change for classification, automated and expert classification of satellite data [28]. The resolution of such data varies, but typically can range from ~10's to 100's of meters [2]. Moreover, landscape features in relation to disease incidence, have generally been related to disease through regression models, and do not account for underlying non-linear relationships between environmental features and disease incidence, nor the partial nature of the observations.

More recently, studies using Google Maps images take advantage of higher resolution data from satellites such as DigitalGlobe (0.15-1.24m resolution). Such publicly available remote sensing data has recently been harnessed for many social applications. From monitoring wheat fungus in crops [17], to predicting poverty [12], crop yields [26] and correlating with obesity [14]. The higher resolution potentially provides the opportunity to distinguish more specific landscape features, such as boundaries of homes and roads, and presence of individual trees, which can have a potential impact on the transmission of diseases.

In this paper, we present a deep learning model for acquisition of deep landscape features from high-dimensional satellite imagery, and also demonstrate integration of these deep features in an infectious disease prediction model and assess their performance for improving infectious disease prediction. To overcome the lack of training data on Pakistan images, we use a transfer learning approach. We also assess the added value of each landscape feature. Finally, given the expected limitation(s) of using training data from another location, we study how the transfer performs in urban versus rural regions of the Pakistani cities. Our results show that landscape features derived from automatically learned features improve prediction of dengue case data

(even with transferring information from landscape features from another location). This work serves as a valuable proof of concept for integration of new data sources into infectious disease modeling, and directions for future work. Given that the satellite data is freely-available, this work is easily scalable, which is especially relevant for places affected by vector borne diseases.

2. Related Work

Landscape features, collected at various spatial and temporal resolutions, have been used in the past to identify disease risk. Studies which either use pre-labelled landcover maps, primitive image segmentation or classification approaches to label satellite imagery include Vanwambeke et. al, who analyzed how changes in landscape features such as forests, orchards and dam construction changes amount of malaria and dengue spreading mosquitoes in Thailand. The study uses coarse-grained Landstat data collected at a resolution of 30m [24]. Another work by Vanwambeke et. al also using low resolution satellite imagery data, studied the relationship between abundance of Aedes and Anopheles larvae, and 5 land cover features: forest, irrigated fields, orchards, peri-urban housing and villages [25]. Nakhapakorn et. al classified low resolution satellite imagery data from Landsat, into 4 classes: build-up, water, agriculture and forest areas using Maximum Likelihood Classification, and identified the percentage of dengue incidence occurring in each class in a province in Thailand [15]. These studies often rely on pre-labelled maps, which hinders scalability of the approach. Moreover, given the localized nature of disease activity, there is a need to resolve landscape features such as individual houses, trees and roads, which require a higher resolution than those described above, but can potentially impact the spread of disease. In addition, these studies have modelled the coarse-grained landscape features as part of linear models which often do not account for the underlying dynamics of disease, as opposed to disease transmission models which aim to capture such dynamics.

Some studies have used vegetation metrics, such as enhanced vegetation index (EVI) and normalized difference vegetation index (NDVI), extracted from satellite images, to model dengue incidence [13, 6]. While these aggregated metrics have shown to correlate with dengue risk, in the absence of data on individual spectral bands, it is difficult to measure non-vegetation landscape features such as building and water bodies from them. Similarly, nightlight imagery has been used as a proxy of changes in populations to model measles in Niger [5]. This study does not use any framework to first extract features from the nightlight images, and instead directly uses averaged pixel intensity values.

More recently, studies have started using deep learning frameworks to classify satellite images for disease and related predictions. The high dimensional output of a VGG neural network, as individual features, was used to predict the prevalence of obesity in a neighborhood [14]. Satellite imagery data has also been used to predict poverty levels in five countries in Africa. The study uses a VGG architecture to infer representation of nightlight images, from corresponding daylight images, and use them to predict poverty [12]. While these studies perform well, the lack of knowledge about which landscape features are represented through the high dimensional output of a neural network, makes it difficult to assess the relationship between individual landscape features and the problem under consideration.

3. Datasets

Here we use publicly available high resolution satellite imagery data from Google Maps API. Using boundaries for the cities of Lahore and Rawalpindi in Pakistan, we download satellite imagery data at a zoom level of 17 with each image scaled to a factor of 2 (n =8,632 images 6,476 for Lahore, 2,156 for Rawalpindi). Each downloaded image is of native resolution 1280 by 1280 pixels.

Given the absence of semantic segmentation maps for the city of Pakistan, we use labeled data from the Kaggle DSTL (Defence Science and Technology Laboratory) competition (https://www.kaggle.com/c/dstl-satellite-imagery-feature-detection) for training purposes. The datasets consists of 25 images of resolution 3600 by 3600. The dataset also contains segmentation maps (pixel-wise labels) for 10 labeled classes corresponding to each image. We selected landscape classes for which there is some prior knowledge regarding mechanistic relevance to vector-borne disease spread. Roads, buildings and crops were included given that the movement of individuals and urbanicity in a location impacts dengue transmission [27]. Additionally, trees and water bodies like standing water and waterways provide good breeding sites for dengue transmitting mosquitoes [1] and hence were also included.

Confirmed dengue incidence data from both cities was received from the Punjab Information Technology Board. Details of confirmed dengue incidence from all public hospitals in Pakistan's Punjab province are aggregated in a centralized server. The dataset consisted of 10,888 cases reported in the cities of Rawalpindi (n=7,890 between January 1, 2014 and December 31, 2017) and Lahore (n=2,998 between January 1, 2012 and December 31, 2017). City-wide daily mean temperature and precipitation were obtained from the Pakistan Meteorological Department [3]. High resolution population data was retrieved from World-Pop [4], and previously published work [21].
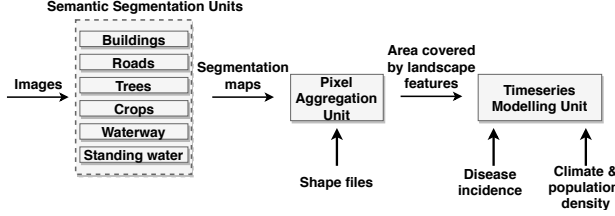
Figure 1. Pipeline of the multi-step methodology used in the study.

4. Methods

4.1. Overview

We use a multi-step approach to understand if the addition of landscape features improves the prediction of dengue transmission in the cities of Lahore and Rawalpindi in Pakistan. To achieve this, we first use a deep learning approach to create segmentation maps for the six landscape classes, from satellite images from both cities. Given the localized nature of dengue transmission, it is helpful to model the transmission of dengue at a sub-city spatial resolution in each city. Hence the segmentation maps of each landscape class, extracted from the deep learning approach, are aggregated at a sub-city spatial unit level in each city to identify the percentage of a spatial unit covered by a particular landscape class. The landscape features are then included in a common time series disease modelling framework that accounts for population mixing, disease specific properties and exogenous parameters and the processes underlying relations between those factors and disease incidence [7, 13]. In addition to the extracted landscape features, weather parameters and population density, which are commonly used, are included to model the transmission of dengue and predict disease incidence over time. Figure 1 shows the pipeline of the multi-step approach used here.

4.2. Semantic Segmentation

4.2.1 Architecture

To generate segmentation maps (pixel-wise labels) of individuals landscape features from satellite imagery, we use a U-Net architecture, a class of CNN, given they tend to perform well for semantic segmentation tasks with limited training data [19]. The architecture consists of a series of successive convolutional and maximum pooling layers, followed by a series of convolutional and up-convolutional layers (Figure 2). The encoder extracts features, while the decoder reconstructs dense representation of pixel-wise classification using the extracted features from the encoder. Skip connections, present in a U-net architecture from the encoder to decoder, allow incorporation of both global and local features present inside an image in the final output to improve the pixel-wise classification [19].

Given that both the encoder of U-Net and VGG neural networks consist of a series of convolutional and max pooling layers, and given the lack of a large training dataset available to train the neural network from scratch, we use a VGG16 neural network, trained on ImageNet data, as the encoder of the architecture. A similar methodology, but with VGG11 architecture as an encoder has been described elsewhere [10]. ImageNet data consists of millions of labelled images from over a thousand categories [20], and networks trained on this dataset are known to be good extractors of visual features such as edges and corners [12]. The fully connected layers at the end of the VGG16 architecture are discarded, given they represent the decoder part of the VGG architecture. The remaining convolutional layers are then used as a pre-trained encoder. The decoder, to complement this pre-trained encoder, is then constructed with randomly initialized weights.

4.2.2 Training, optimization and prediction

We train separate networks for each landscape feature, given that landscape classes in the training dataset are not mutually exclusive [9]. We use an input dimension of 256 by 256 by 3 pixels to the network, where the third dimension represent the red green blue (RGB) band of an image. The input size is chosen to ensure that the network does not cause memory issues when training and is comparable to that used in other work harnessing similar features from satellite imagery [30]. In addition the size is chosen as a multiple of 1024 to ensure that the images, to be predicted, can be subset into non-overlapping sub images without padding white spaces on boundary images [19]. The output of the network is 256 by 256 by 1 representing the probability of each pixel belonging to a particular landscape class for which the model is trained. We use an Adam optimizer with the default learning rate of 0.001 and optimize the loss function C, which is defined as: $C = B - log(J)$, where B is the binary entropy between predicted and actual segmentation maps and J is the Jaccard similarity index between predicted and actual segmentation maps.

To learn the weights of the decoder we use data from the Kaggle DSTL competition. Pixels in each image in the dataset are normalized using the combined mean and the standard deviation of the pixels in the training dataset. The images and corresponding pixel-wise labels in the dataset are then partitioned into non-overlapping blocks of sub-images of size equal to the input size of the network. The network is then trained using the sub-images.

To generate pixel-wise labels for images in Lahore and Rawalpindi, we first re-size all images from Google API in our set to 1024 by 1024 pixels. Pixels in the images are then normalized using the combined mean and the standard deviation of pixels in the downloaded images from Google
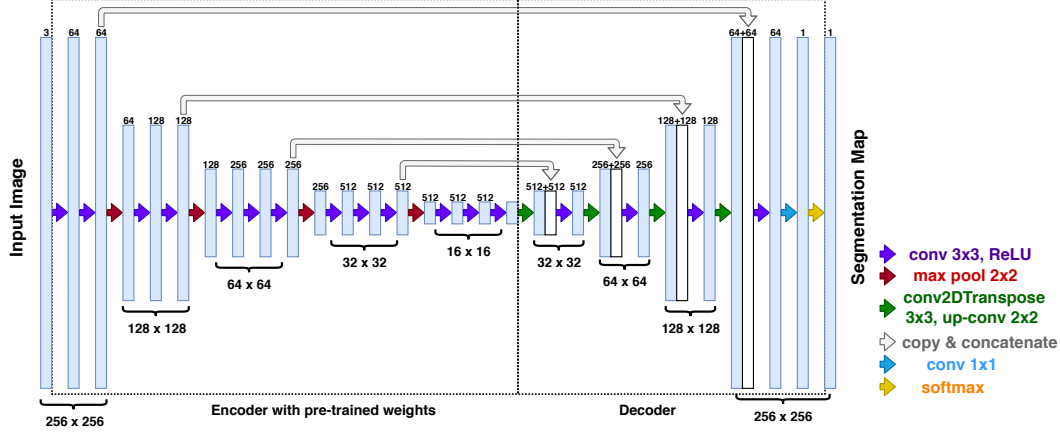
Figure 2. Architecture of U-net used for semantic segmentation. Dotted lines enclose the encoder and decoder. The encoder consists of 13 3x3 convolutions each followed by a rectified linear unit (ReLU), and 4 2x2 max pooling operations with strides of 2 to reduce the map size by 2. Weights of all layers in the encoder are pre-trained on ImageNet dataset. The decoder consists of 4 3x3 convolutions each followed by a rectified linear activation unit (ReLU) and 1 1x1 convolution. The decoder also consists of 4 3x3 transposed convolutions with strides of 2 to upsample maps by 2. The upsampled maps are concatenated with skip connections from the encoder. The architecture has a final softmax activation which predicts the probability of each pixel in the image belonging to the landscape class.

API, as opposed to images from the DSTL challenge. This ensures that any differences, due to variation in exposure or lightning, when capturing the DSTL and Google Maps images, are minimized. Each image is then partitioned into non-overlapping sub-images of the size of input of the network, and predictions made using the pretrained models for each landscape feature. Keras library in Python, with a Tensorflow backend is used for semantic segmentation.

4.3. Timeseries Modelling

We model the dengue transmission using a timeseries susceptible infected recovered (TSIR) model of viral incidence which has been widely used in epidemiology to reconstruct dynamics of diseases [7, 13]. Each city is divided into sub-city spatial units (towns) to model localized transmission (n=10 spatial units in Lahore and n=14 spatial units in Rawalpindi). Spatial units for the city of Lahore are the sub-city resolution administrative boundaries, while for the city of Rawalpindi we use boundaries as defined by Rehman *et al.* [18] based on clusters of dengue incidence in the city. These spatial units span a median area of 57 km^2 (min=24, max= 437, n=10) for the city of Lahore and a median area of 37 km^2 (min=9, max= 84, n=14) for the city of Rawalpindi. This spatial unit size balanced locality and density of cases and is reasoned in other work [18].

Given that only the most recent satellite imagery data was available, we model landscape features as time-invariant in the dengue transmission model. We calculate the proportion of area, in each spatial unit, covered by each landscape feature. To achieve this, we first select the images belonging to each spatial unit using the boundaries of

the spatial units. The predicted pixel-wise labels of the selected images are then used to calculate the ratio of total number of pixels representing the presence of a landscape feature, and the total number pixels in the selected images.

The entire modeling approach is described below succinctly, however we emphasize that such approaches, parameters chosen, etc. are the standard in infectious disease modeling. Further details on incidence data and this modelling approach including correction for under-reporting of cases and discussion on accuracy of data, can be found in Rehman *et al.*'s work [18] from where the specific dengue TSIR modeling methodology has been adapted. The general TSIR model, for each spatial unit, i, is defined via:

$$I_i(t+1) = \beta_i(t)\frac{S_i(t)}{N_i(t)}I_i^{\alpha_i}(t)\epsilon \qquad (1)$$

where an interval of 2-weeks is used for t. $I_i(t)$, $S_i(t)$, and $N_i(t)$ are the infected, susceptible and total population during time step t in spatial unit i, α_i is the mixing coefficient in i, and $\beta_i(t)$ is the transmission rate during time step t. The error term ϵ is assumed to be an independent and identically log-normally distributed random variable. Weather parameters (bi-weekly mean temperature and number of rainfall days) and population density are commonly incorporated in such models given they are known to affect dengue transmission [31, 16]. Time invariant landscape features, and time-variant weather parameters and population density are modelled as part of the transmission rate. Equation 1 is rearranged as:

$$\begin{aligned} log(\beta_i(t)) + \alpha_i * log(I_i(t)) = \\ log(I_i(t+1)) + log(N_i(t)) - log(S_i(t)) - \epsilon \end{aligned} \qquad (2)$$
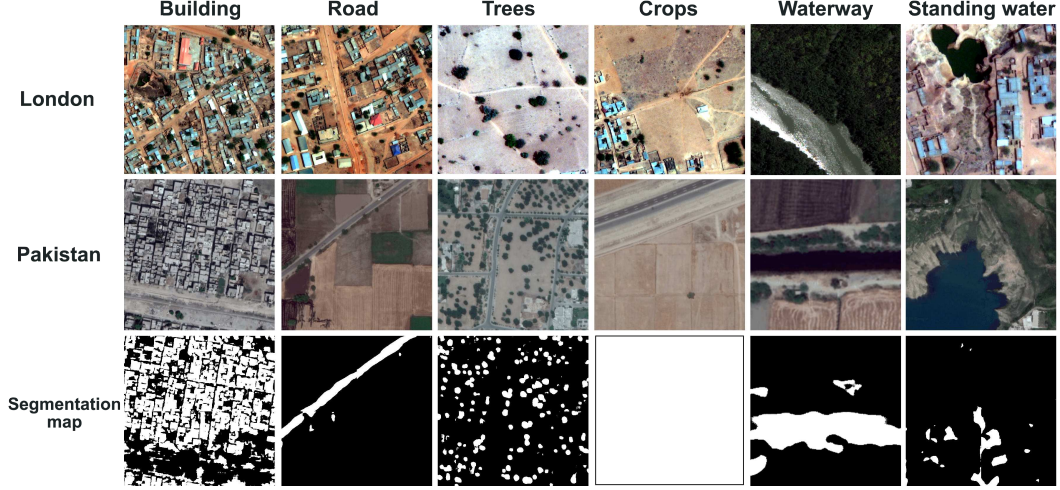
Figure 3. Sample images of landscape features present in DSTL data (row 1) and Google Maps API data (row 2), and predicted segmentation maps corresponding to the Google Maps API data (row 3). Landscape features from left to right: building, roads, trees, crops, waterway and standing water. Higher intensity corresponds to a higher probability of a pixel belonging to a landscape class. Images shown are constructed from varying number of sub-images of size 256 x 256 pixels to ensure individual landscape features are clearly visible.

and the transmission rate is substituted with:

$$log(\beta_i(t)) = \sum_a \theta_a L_{i,a} + \sum_j \theta_j E_j(t - l_j) + \theta_p D_i(t)$$

(3)

where l_j are the delays in time steps which are added to weather parameters j to account for vector life cycle. $L_{i,a}$ is the percentage of area of spatial unit i covered by landscape feature a. $E_j(t - l_j)$ is the value of weather parameter j during time step $(t - l_j)$. $D_i(t)$ is the population density in spatial unit i, which is known to also drive transmission, β_i, and is computed by dividing the population of the spatial unit with the area of the spatial unit. θ_a, θ_j and θ_p are parameters which relate landscape features, weather and population density to β.

We use a linear model to fit the relationship in equation 2 and then study how the fit varies for model with and without the addition of landscape features. In addition, we also study the fit of the model by incorporating each type of resolved landscape feature individually in the model.

5. Results

First, we use a transfer learning approach to generate segmentation maps of six landscape features from the satellite imagery data of the cities. The trained architectures for landscape classes show good performance when predicting the segmentation maps on the 20% held-out labelled data from the DSTL challenge. Specifically, the architecture provided a higher Jaccard similarity for buildings (0.628), roads (0.660), trees (0.557), crops (0.813), and waterway (0.618). Performance on standing water was lower (0.357). The estimated Jaccard index values on the held out dataset

are slightly lower than those reported in a previous work which uses DSTL data [9], but in our work we only use the RGB bands to train the architectures, and not the additional 8 multi-spectral and 1 short-wave infrared bands as we will be predicting on satellite image which only use RGB bands.

Though the encoder of the architectures were pre-trained on ImageNet data, and decoder on DSTL challenge data from London, our goal was to eventually make predictions on data from Pakistan. Visually analyzing predicted segmentation maps for satellite imagery data in cities of Pakistan (Figure 3) shows that despite using no labelled data from Pakistan, the architectures were able to identify the segmentation maps of buildings, roads and trees with a reasonable quality. Results for waterways provided medium quality. Identification of crops and standing water landscape features on Pakistan data were not meaningful.

Epidemic model results which included proportion landscape features, derived from the satellite imagery, by area result in time series of predicted dengue cases. For evaluation we use the adjusted R^2 metric, which allows the comparison of the fit of various models while accounting for the number of parameters being used in the model. The base model trained only on weather parameters and population density provide a good fit with the incidence of dengue, with an adjusted R^2 value of 0.728 in Lahore and 0.774 in Rawalpindi. Addition of the landscape features in the model improved the fit (0.747 in Lahore and 0.808 in Rawalpindi).

To examine how the London training data generalized to different types of Pakistan landscapes, we assessed the goodness of fit in more versus less urban spatial units in both cities, as this is a natural potential difference between

Model	Lahore			Rawalpindi		
	all towns	more urban	less urban	all towns	more urban	less urban
Environment only	0.728	0.736	0.715	0.774	0.801	0.660
All landscape & environment	**0.747**	**0.757**	**0.734**	**0.808**	**0.832**	**0.717**
Building & environment	0.732	0.746	0.725	0.784	0.824	0.693
Road & environment	0.731	0.746	0.719	0.781	0.804	0.678
Trees & environment	0.730	0.742	0.725	0.785	0.813	0.678
Crops & environment	0.728	0.736	0.715	0.774	0.801	0.660
Waterway & environment	0.737	0.745	0.725	0.796	0.817	0.702
Standing water & environment	0.730	0.736	0.723	0.778	0.808	0.663

Table 1. Adjusted R-squared values of model fit using only i) weather and population density (environment parameters), ii) one landscape feature and environment parameters, and iii) all landscape features and environment parameters. The model fit values are shown across i) all towns, ii) more urban towns, and iii) less urban towns. Best model results in bold.

the two places. For each city, we identified spatial units which have more areas covered by buildings compared to the average area covered by buildings in the spatial units. We found that the fit of model is better in more urban areas as compared to less urban areas, across both cities. (0.757 as compared to 0.734 in Lahore and 0.832 as compared to 0.717 in Rawalpindi). Incorporation of landscape features individually showed varied improvement in the fit of the model. Overall we find incorporation of buildings and waterways to provide the most improvement in results, followed by roads and trees. Incorporation of crops and standing water provided little or no improvement in the fit. Table 1 summarizes the results across all models in both cities.

6. Conclusion and Future Work

Despite differences in the types of roofs and color of roads in London and cities in Pakistan, architectures trained for buildings and roads were able to predict the segmentation maps for satellite data in Pakistan with a reasonable accuracy (Figure 3). Predictions for trees were the most accurate amongst all landscape features, given the consistency of color and shape of trees across the datasets. Given the variation in waterways in both datasets, despite relatively high Jaccard index for this feature, the trained architecture was only able to identify approximate location of waterways and not the exact outline on the Pakistan data. Results in Table 1 reflect not only the predictive power of the landscape features, but also how correct the generated segmentation maps were for the landscape features. Little to no improvement in the adjusted R^2 value with the addition of crops and standing water in the model does not conclude that these features are not predictive of dengue transmission. Lack of predictive performance in the model when including standing water feature can be attributed to the already low performance of the architecture on the held out training data. In contrast, the deep learning architecture trained for crops provided the highest predictive performance on held out data, yet given the lack of variation in shapes and color of crops in the training data, the architecture was not able to learn generalizable features for transfer to the Pakistan data; when applied to satellite imagery data from Pakistan, the model tended to predict every image as entirely consisting of crops (example subplot in Figure 3 consisting of all white pixels). Amongst all landscape features, buildings provided both a large improvement in the fit of the timeseries model and segmentation maps which were reasonably accurate.

Results consistently showed that addition of landscape features improved the fit of the model across both cities, and in both rural and urban areas. Base models for the cities provided a better fit in urban areas as compared to rural areas as dengue is more prevalent in urban areas. In addition models for the city of Rawalpindi provided a better fit compared to those for Lahore, given that the disease activity was higher in Rawalpindi during the study time period.

While results demonstrate the utility of adding landscape features in timeseries modelling of dengue, steps can be taken to make the results more robust. First, using labelled satellite imagery data from Pakistan, during training, can help improve the prediction of segmentation maps. Doing so can help answer more precise questions such as "how can landscapes be changed to reduce disease incidence". Second, given the high spatial resolution of satellite imagery data, modelling the transmission of dengue at an even higher resolution could potentially help identify more robust relationships between landscape features and disease transmission. Also, we only used satellite images from one time period in this study, based on availability. This may be fair as several considered landscape features may not have changed over the time period of the study. However for features where there is expected differences, it may be useful to resolve them with more temporal resolution. Finally, an unsupervised approach, using the high dimensional output of the neural network, as done in previous work [14, 12], may identify unknown features which can potentially impact disease transmission.

References

[1] Dengue and the aedes aegypti mosquito - cdc. https://www.cdc.gov/dengue/resources/30Jan2012/aegyptifactsheet.pdf. Accessed: May 14, 2019.

[2] Landsat 8 landsat science - nasa. https://landsat.gsfc.nasa.gov/landsat-data-continuity-mission/. Accessed: May 14, 2019.

[3] Pakistan meteorological department pmd. http://www.pmd.gov.pk/. Accessed: May 14, 2019.

[4] Worldpop. http://www.worldpop.org.uk/. Accessed: May 14, 2019.

[5] Nita Bharti, Andrew J Tatem, Matthew J Ferrari, Rebecca F Grais, Ali Djibo, and Bryan T Grenfell. Explaining seasonal fluctuations of measles in niger using nighttime lights imagery. *Science*, 334(6061):1424–1427, 2011.

[6] S. Bhatt, P. W. Gething, O. J. Brady, J. P. Messina, A. W. Farlow, C. L. Moyes, J. M. Drake, J. S. Brownstein, A. G. Hoen, O. Sankoh, M. F. Myers, D. B. George, T. Jaenisch, G. R. Wint, C. P. Simmons, T. W. Scott, J. J. Farrar, and S. I. Hay. The global distribution and burden of dengue. *Nature*, 496(7446):504–7, 2013.

[7] Brbel F Finkenstdt and Bryan T Grenfell. Time series modelling of childhood diseases: a dynamical systems approach. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 49(2):187–205, 2000.

[8] Scott B Halstead. Selective primary health care: strategies for control of disease in the developing world. xi. dengue. *Reviews of infectious diseases*, 6(2):251–264, 1984.

[9] Vladimir Iglovikov, Sergey Mushinskiy, and Vladimir Osin. Satellite imagery feature detection using deep convolutional neural network: A kaggle competition. *arXiv preprint arXiv:1706.06169*, 2017.

[10] Vladimir Iglovikov and Alexey Shvets. Ternausnet: U-net with vgg11 encoder pre-trained on imagenet for image segmentation. *arXiv preprint arXiv:1801.05746*, 2018.

[11] Laura E Jackson, Elizabeth D Hilborn, and James C Thomas. Towards landscape design guidelines for reducing lyme disease risk. *International Journal of Epidemiology*, 35(2):315–322, 2006.

[12] 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, 2016.

[13] M. U. Kraemer, T. A. Perkins, D. A. Cummings, R. Zakar, S. I. Hay, D. L. Smith, and Jr. Reiner, R. C. Big city, small world: density, contact rates, and transmission of dengue across pakistan. *J R Soc Interface*, 12(111):20150468, 2015.

[14] Adyasha Maharana and Elaine Okanyene Nsoesie. Use of deep learning to examine the association of the built environment with prevalence of neighborhood adult obesity. *JAMA network open*, 1(4):e181535–e181535, 2018.

[15] Kanchana Nakhapakorn and Nitin Kumar Tripathi. An information value based analysis of physical and climatic factors affecting dengue fever and dengue haemorrhagic fever incidence. *International Journal of Health Geographics*, 4(1):13, 2005.

[16] World Health Organization et al. Dengue guidelines for diagnosis, treatment, prevention and control: new edition. 2009.

[17] Reid Pryzant, Stefano Ermon, and David Lobell. Monitoring ethiopian wheat fungus with satellite imagery and deep feature learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 39–47, 2017.

[18] Nabeel Abdur Rehman, Henrik Salje, Moritz UG Kraemer, Lakshminarayanan Subramanian, Simon Cauchemez, Umar Saif, and Rumi Chunara. Quantifying the impact of dengue containment activities using high-resolution observational data. *bioRxiv*, page 401653, 2018.

[19] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

[20] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.

[21] Safdar Ali Shirazi. Spatial analysis of ndvi and density of population: a case study of lahore-pakistan. *Science International*, 24(3), 2012.

[22] Cameron P Simmons, Jeremy J Farrar, Nguyen van Vinh Chau, and Bridget Wills. Dengue. *New England Journal of Medicine*, 366(15):1423–1432, 2012.

[23] Katherine F Smith, Andrew P Dobson, F Ellis McKenzie, Leslie A Real, David L Smith, and Mark L Wilson. Ecological theory to enhance infectious disease control and public health policy. *Frontiers in Ecology and the Environment*, 3(1):29–37, 2005.

[24] Sophie O Vanwambeke, Eric F Lambin, Markus P Eichhorn, Stéphane P Flasse, Ralph E Harbach, Linda Oskam, Pradya Somboon, Stella van Beers, Birgit HB van Benthem, Cathy Walton, et al. Impact of land-use change on dengue and malaria in northern thailand. *EcoHealth*, 4(1):37–51, 2007.

[25] Sophie O Vanwambeke, Pradya Somboon, Ralph E Harbach, Mark Isenstadt, Eric F Lambin, Catherine Walton, and Roger K Butlin. Landscape and land cover factors influence the presence of aedes and anopheles larvae. *Journal of medical entomology*, 44(1):133–144, 2007.

[26] Anna X Wang, Caelin Tran, Nikhil Desai, David Lobell, and Stefano Ermon. Deep transfer learning for crop yield prediction with remote sensing data. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, page 50. ACM, 2018.

[27] A. Wesolowski, T. Qureshi, M. F. Boni, P. R. Sundsoy, M. A. Johansson, S. B. Rasheed, K. Engo-Monsen, and C. O. Buckee. Impact of human mobility on the emergence of dengue epidemics in pakistan. *Proc Natl Acad Sci U S A*, 112(38):11887–92, 2015.

[28] J Wickham, SV Stehman, and CG Homer. Spatial patterns of the united states national land cover dataset (nlcd) land-

cover change thematic accuracy (2001–2011). *International journal of remote sensing*, 39(6):1729–1743, 2018.

[29] World Health Organization. Vector-borne diseases. `https://www.who.int/news-room/fact-sheets/detail/vector-borne-diseases`. Accessed: May 14, 2019.

[30] 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*, 2016.

[31] Lei Xu, Leif C Stige, Kung-Sik Chan, Jie Zhou, Jun Yang, Shaowei Sang, Ming Wang, Zhicong Yang, Ziqiang Yan, Tong Jiang, et al. Climate variation drives dengue dynamics. *Proceedings of the National Academy of Sciences*, 114(1):113–118, 2017.