

U-net for learning and inference of dense representation of multiple air pollutants from satellite imagery

Jacquelyn Shelton*
jacquelyn.ann.shelton@gmail.com
Hong Kong Polytechnic University
Hung Hom, Hong Kong SAR, China

Przemyslaw Polewski
TomTom Location Technology
Germany GmbH
Berlin, Germany
Hong Kong Polytechnic University
Hung Hom, Hong Kong SAR, China

Wei Yao
Hong Kong Polytechnic University
Hung Hom, Hong Kong SAR, China

ABSTRACT

Air pollution is an important topic on countless fronts and is an active area of research. The goal of this work is to provide a machine learning model for learning and inference of pollution concentrations and air quality measures, namely Particulate Matter 2.5, NO_3 , Nitrate Pollution, and NH_4 , Atmospheric Ammonium, with high granularity by using easily obtainable satellite imagery data. In order to achieve this, we propose the fully convolutional network U-net that, unlike previous work, can predict these pollutant values at a pixel-level high-resolution instead of being able only to predict a single value for an entire geographical region. We demonstrate that this approach can reconstruct the considered pollutant concentrations on ground-truth data and can predict the concentrations and their spatial structure reasonably well, even for data that the network has temporally not yet seen. Finally, we illustrate that the model's pollutant predictions can offer valuable insights into the current *COVID-19* pandemic.

CCS CONCEPTS

• Applied computing → Environmental sciences.

KEYWORDS

U-net, dense regression, pollution, machine learning

ACM Reference Format:

Jacquelyn Shelton, Przemyslaw Polewski, and Wei Yao. 2020. U-net for learning and inference of dense representation of multiple air pollutants from satellite imagery. In *10th International Conference on Climate Informatics (CI2020)*, September 22–25, 2020, virtual, United Kingdom. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3429309.3429328>

1 INTRODUCTION

In 2020 the United Nations Environment Programme (UNEP), United Nations Human Settlement Programme (UN-Habitat) and IQAir AG announced the release of the largest air quality data platform

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CI2020, September 22–25, 2020, virtual, United Kingdom

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8848-1/20/09...\$15.00

<https://doi.org/10.1145/3429309.3429328>

to date. While this initiative's announcement focuses on the health aspects of pollution the value of this data set are not to be understated. In 2019 with the outbreak of the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2), commonly referred to as *COVID-19* and the discovery of SARS-CoV-2 RNA on air pollution the value in this data and the potential to not just measure local but also gain insight into the movement of air pollution became of interest. While SARS-CoV-2 is primarily related to $PM_{2.5}$ and PM_{10} , Particulate Matter pollution of 2.5 and 10 micrometers and smaller in diameter respectively, other interests, health, environmental and agricultural, are also interested in other pollution and air quality measures, in particular NO_3 , Nitrate Pollution, and NH_4 , Atmospheric Ammonium.

Furthermore, it is important to note the connections between the lethality rate of *COVID-19* and high levels of atmospheric pollution. For example, in [8] these connections are explored in Lombardy and Emilia Romagna. In particular, the paper studies the correlation between pollution, which is a known instigator of chronic lung disease even in young and other wise healthy subjects, and the lethality of SARS-CoV-2. A similar result for the lethality in the United States [19] using county level fatality rate, and county level long term air pollution, shows that, after adjusting for other known factors, that there is a strong correlation between the concentration of particulate matter 2.5 micrometers or less in diameter, or $PM_{2.5}$, and the county level lethality. Specifically, a 1 $\frac{\mu g}{m^3}$ increase in $PM_{2.5}$ corresponds to an 8% increase in the fatality rate.

However, as observed in [7], there is also a correlation between $PM_{2.5}$ air pollution and the number of reported cases. This suggests that pollution-to-human may serve as another transmission dynamic for SARS-CoV-2. These two results suggest a two factor vulnerability to SARS-CoV-2 caused by increased particulate matter in the air: On one hand it increases the likelihood of having a more severe reaction to infection, and the other it serves a transmission vector.

Traditionally, $PM_{2.5}$ concentration data can be obtained from ground sensors and measurement stations. However, the spatial resolution of these measurements is greatly limited by the sparsity of sensor networks. Thus, detailed pollution maps have been developed that integrate heterogeneous data sources (see e.g. [18]). Namely, they proposed a $PM_{2.5}$ estimation of fusing aerosol optical depth information computed from satellite imagery with a chemical transport model and regional ground-based observations. With this, they published a database of estimated monthly pollutant concentrations over several countries. This data was used as

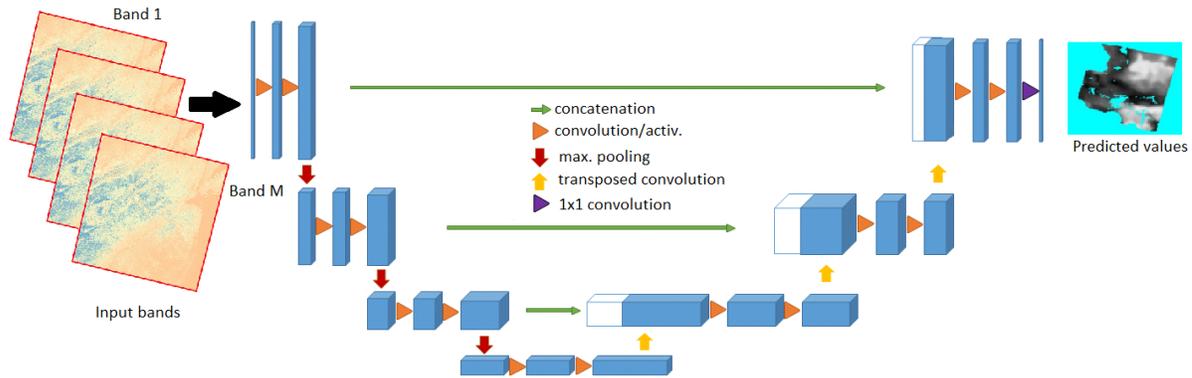


Figure 1: Network structure for predicting pollutant ($PM_{2.5}$, NO_3 , NH_4) concentrations from multispectral imagery. The network follows the classic U-net [15] architecture.

the basis for the aforementioned study by [19] linking $PM_{2.5}$ concentrations to COVID-19 mortality rates. Although this historical data helped establish this causality relationship, up-to-date/live pollution information is still needed to monitor pollutants such as $PM_{2.5}$, NO_3 , and NH_4 (i.e. pollutant concentration data past the range 2001-2017).

Recently, there has been a growing interest in retrieving spatio-temporal air quality products based on high resolution optical satellite missions such as *Landsat 8* [5], utilizing top-of-atmosphere image reflectances as a basis for directly predicting pollution parameters. To the best of our knowledge, dense (per-pixel) pollution level prediction with fully convolutional neural networks has not yet been described in literature. On the other hand, a number of authors perform ‘sparse’ (i.e. patch or region granularity) estimation using classical convolutional neural network architectures. Li et al. [11] utilized deep belief networks to predict $PM_{2.5}$ levels in China from MODIS (Moderate Resolution Imaging Spectroradiometer)-derived aerosol optical depth (AOD) fused with meteorological parameters. Shen et al. [17] utilized a similar approach, but regressed ground-level $PM_{2.5}$ concentrations directly from top-of-atmosphere reflectances instead of AOD layers. Recently [23], a convolutional neural network was applied to high-resolution imagery from commercial micro-satellites in order to extract informative features for random-forest based regression of $PM_{2.5}$ concentrations. A multi-layer perceptron network was used for predicting particulate matter concentrations from *Landsat 8* imagery [21]. Moreover, several contributions deal with estimating particulate matter levels from street (terrestrial perspective) imagery by means of deep learning. In [22], an ensemble of deep neural networks is used as a basis for classification of $PM_{2.5}$ and PM_{10} with a granularity of 6 categories. Another contribution [14] proposes an ensemble of 3 deep learning networks aggregated by a feed-forward network in the role of the meta-learner for regression of the pollution levels.

The primary goal of this work is to provide a machine learning model capable of estimating the pollutants’ $PM_{2.5}$, NO_3 , and NH_4 concentrations with high granularity comparable to the data gathered by [18] (details in Sec.3) using plentiful, publicly available satellite imagery. Unlike previous work, the fully convolutional

neural network U-net we employ can predict these pollutant values at a pixel-level high-resolution instead of predicting only a single value for an entire region. The paper is organized as follows: Sec. 2 introduces the proposed U-net model, Sec. 3 presents the data, Sec. 4 describes the experiments and results, and finally Sec. 5 provides a summary of the work and outlook.

2 U-NET MODEL FOR POLLUTANT AGENTS

We build upon the well known U-net architecture [15] to predict dense (per-pixel) pollutant concentrations from multispectral satellite imagery. Since its introduction in 2015, the U-net has been successfully applied to various semantic segmentation tasks, E.g. in medical imaging [9] and astronomy [6]. As a fully convolutional network (and hence devoid of fully connected layers), the U-net can be applied to an image of arbitrary size in a sliding window/tile based fashion. The U-net consists of two symmetrical parts. The encoder path downsamples the original image into meaningful features by means of convolutional filters and pooling operations, whereas the upsampling path aims at decoding these features into a full-sized output map using transposed convolution operations, driven by an appropriate loss function. Moreover, upsampling layers are augmented with feature maps from the downsampling path at the corresponding resolution, to provide more context information. The original U-net was meant for classification and featured a softmax layer after the top-level upsampling layer’s output, resulting in per-class posterior probabilities meant to be optimized with using a cross-entropy type objective with discrete ground-truth labels. However, the U-net has also been used directly for dense regression, by removing the softmax layer and optimizing the squared difference between the upsampled output and a continuous target variable [13, 20]. In this work, we adopt a similar approach. Let $S = \{(x_i, y_i)\}, 1 \leq i \leq N$ denote a set of predictor variable vectors x and matching continuous response variables y (i.e. the target pollutant concentration), while $f(x|\theta)$ indicates the prediction function (i.e. the U-net) parameterized by θ . In the case of fully convolutional networks, the parameter θ comprises the set of trainable convolutional and transposed convolutional filter parameters, i.e. the weight and bias matrices. Although most state-of-the-art

activation functions, such as the rectified linear unit (ReLU), do not contribute any parameters, there are also emerging approaches which attempt to learn the optimal activation function together with the filter parameters [12]. The optimization objective can be formulated as:

$$\operatorname{argmin}_{\theta} \sum_i w_i |y_i - f(x_i|\theta)|^2 \quad (1)$$

It is well known that least squares regression estimates the conditional mean of the response variable given the predictors, and is therefore sensitive to outliers. To alleviate this, we eliminate outliers from the training set by discarding pairs (x_i, y_i) such that y_i is below the 0.01-th or above the 0.99-th quantile of its distribution. This is achieved by setting the weights w_i corresponding to such pixels within the training images to 0. Otherwise, the respective weights are set to unity.

3 DATA

In order to address this problem, we use satellite imagery as a source of predictor variables, and approximate pollutant concentrations published by [18] for the time period of 2000-2017 in the role of ground-truth.

Satellite data. This study was based on Landsat 8 satellite imagery published by the United States Geological Survey (USGS) [5]. Landsat 8 is the latest installment in a series of Earth observation missions dating back to 1972 (Landsat 1), which provide planet-wide imagery from multispectral, spaceborne optical sensors. Landsat uses its own coordinate system [2], which assigns coordinates to an image frame based on nominal satellite orbital tracks and the frame’s latitudinal center line, resulting in a 2-tuple of (path,row) coordinates. The ground coverage cycle length is 16 days, which means that new data for an image frame is published approximately twice per month. Landsat 8 images contain a total of 11 spectral bands, ranging in wavelengths from 0.435 μm (band 1 - coastal/aerosol) to 12.51 μm (band 11 - thermal infrared). The spatial resolution (i.e. ground sampling distance per pixel) of the images ranges from 15 m (panchromatic band 8) through 30 m (bands 1-7, 9), up to 100 m (bands 10-11). We used the USGS Earth Explorer [4] as well as the Landsat dataset mirror on Amazon AWS [1] to download the imagery.

Pollution data. We downloaded monthly total $PM_{2.5}$, NO_3 , and NH_4 concentration maps for North America from the FTP server of Dalhousie University, Canada, published by Aaron van Donkelaar and collaborators [3]. These maps are made available at a resolution of 0.01 degrees per pixel and feature a standard WGS84 coordinate reference system. The pollutant particulate concentration was obtained via a hybrid approach which simulates geophysical relationships between aerosol optical depth (estimated from multi-source satellite imagery) and $PM_{2.5}$ concentration based on the GEOS-Chem chemical transport model, and fuses the estimates with ground measurements obtained from the US Environmental Protection Agency’s Air Quality System as well as Canada’s National Air Pollution Surveillance [18]. Moreover, estimates of various component pollutants contributing to the total $PM_{2.5}$ are derived from simulated relative composition. The authors report good long term agreement of predicted pollutant agent concentrations with ground-truth validation sites (R^2 coefficients of 0.59-0.90).

We considered only the period after the Landsat 8 mission became active, resulting in a time interval from March 2013 to December 2017.

Preprocessing. The Landsat imagery was first reprojected to the WGS84 coordinate system, and downsampled to the ground-truth resolution of 0.01 degrees (the panchromatic band was dropped). To match the temporal resolution of the ground-truth, we computed per-band average images grouped by month of acquisition. Next, we derived a per-pixel mask of regions within the image covered by cirrus clouds and hence not suitable for analysis. This was done on the basis of the Landsat band 9 (1.360-1.390 μm), which was designed to detect high-altitude clouds. We used the estimated cloud cover percentage [10], available as part of image metadata, to remove the pixels locations having reflectance values in band 9 within the corresponding top quantile. Finally, we combined the cloud cover mask with the data availability mask from the ground-truth $PM_{2.5}$ maps, representing other features that do not have available data e.g. missing data over large bodies of water. Also, pixels corresponding to the top and bottom 1% of ground-truth values were masked out as outliers (see Fig. 5). All input Landsat imagery bands were normalized to the interval [0;1] individually per band.

4 NUMERICAL EXPERIMENTS

We selected an initial number $N = 114$ Landsat images, spanning March 2013 to December 2017, of 24 major cities from representative regions of the United States (see Fig. 2) and preprocessed them as described in Sec. 3. Next, we defined our ground-truth data as these images paired with corresponding pollutant images from which we use an 80:20% random split to create training/testing set of 91:23 images, illustrated in Fig. 2 with the green boxes. A U-net model based on adapting the implementation by [6] was trained separately for each pollutant ($PM_{2.5}$, NO_3 , NH_4) using an ADAM optimized over 500 epochs with minibatch size of 15 and 100 internal iterations, with the following parameterization: dropout ratio was 0.5, learning rate 0.00005. The network had a depth of 3 layers, with 32 filters at the top level. The convolutional kernel size was 3x3, and at each level 2x2 max-pooling operations were used. We applied the tiling strategy proposed by Ronneberger et al. [15] with a tile size of 200x200 pixels.

1: Sanity check with ground-truth.

At the end of training, the networks converged to root mean square error values of 1.24, 2.10, and 1.64 $\mu\text{g}/\text{m}^3$ respectively for $PM_{2.5}$, NO_3 , and NH_4 (see Fig. 3). The corresponding mean absolute errors of the predictions were 0.81, 1.44, and 1.14 $\mu\text{g}/\text{m}^3$. Considering the inter-quantile ranges between the 0.01-th and 0.99-th quantiles for the ground truth distribution of the three pollutants (see Fig.5), the mean absolute deviations constituted 4%, 3%, and 7% of the respective pollutant concentration ranges. Thus the U-net converged and was able to learn ground-truth pollutant concentrations with reasonably low error.

2: Generalizability to temporally novel data.

To assess the ability of our model to generalize in the temporal domain, we utilized 23 previously unseen images that overlapped with the training set spatially but not temporally. The test set contained images from 15 of the 24 cities. The mean absolute deviations

Pollutant concentration maps and regions defining training and testing data

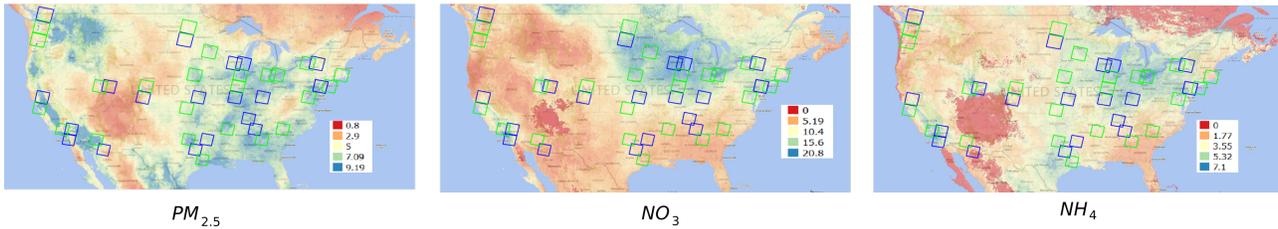


Figure 2: Locations of US cities chosen to reasonably represent the diverse US geography for our study. Rectangles indicate the bounding boxes of the Landsat images’ locations. Green boxes define regions used for training and validation, whereas blue boxes represent regions used exclusively for testing. The background color maps represents average concentrations, in micrograms per cubic meter, of $PM_{2.5}$, NO_3 , and NH_4 in 2017. The corresponding distributions of the training and test data for these pollutants are shown in Fig. 5

Convergence on training data

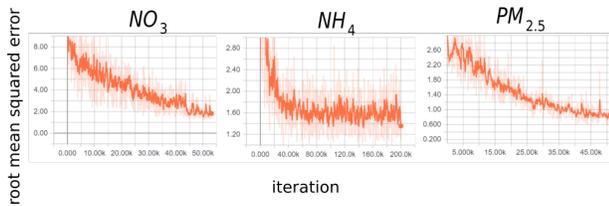


Figure 3: Convergence of training process measured by root mean squared error loss function on training set respectively for NO_3 , NH_4 , and $PM_{2.5}$.

for predictions on this test set were 2.23, 4.75, and 2.24 $\mu g/m^3$ respectively for $PM_{2.5}$, NO_3 , and NH_4 . This corresponded to 11.3%, 11.1%, and 14.1% of the inter-quantile distance. See Fig. 4 for an example of predicted pollutant concentration maps on previously unseen data. These maps illustrate that the network can reconstruct the ground-truth concentrations of each pollutant with which it can represent the spatial structure of these pollutants.

3: Cities with similar pollution profile.

We selected 24 new images at 20 additional locations within the United States (see Fig. 2, blue boxes), showing visually similar distributions of groundtruth pollutant concentrations within the Landsat image frames to the original training cities, in order to evaluate the performance of predicting the concentrations at locations unseen during training. The obtained mean absolute error was 2.46, 4.07, 2.27 $\mu g/m^3$ respectively for $PM_{2.5}$, NO_3 , and NH_4 , yielding percentages of the inter-quantile ranges of 12.4%, 9.5%, and 14.2%. This shows that the network was able to generalize equally well in the spatial and temporal domains.

4: Predicting pollutant $PM_{2.5}$ for COVID-19.

Following the outbreak of COVID-19, there has been a great deal of research into its understanding. Important results [7, 8, 19] suggest a two factor vulnerability to SARS-CoV-2 caused by increased particulate matter in the air: On one hand it increases the likelihood of having a more severe reaction to infection, and the other it serves a transmission vector. Furthermore, detailed pollution maps have been developed that integrate heterogeneous data sources (see e.g. [18]) which were used by [19] directly linking

$PM_{2.5}$ concentrations to COVID-19 mortality rates. Inferring air pollutants with reasonable accuracy, which we have demonstrated in Experiments 1-3 allows us to – essentially for free – highlight more dangerous regions wherein one is more likely to contract the virus.

The goal of the fourth experiment is to verify that our model can predict expected $PM_{2.5}$ concentration trends at time points before and after the government mandated lock-down in March 2020 intended to hinder the spread of SARS-CoV-2 (which consequently e.g. drastically reduced industrial emissions). We applied the U-net learned in Exp. 3 to satellite images of Los Angeles, CA from 2018, 2019, and early 2020. The results, shown in Fig. 6, illustrate that the U-net was able to learn $PM_{2.5}$ concentrations consistent with world events at the time. Namely, the pollution is inferred to be notably higher before the lock-down than after the lock-down, shown by the considerable shift of the predicted $PM_{2.5}$ distribution’s 0.9 quantile, i.e. its 99th percentile, from 13.2 in October 2019 to 9.7 in April 2020. Additionally, our approach was able to successfully isolate regions in LA that are the most densely populated by predicting higher pollutant $PM_{2.5}$ values. This suggests that our approach can generalize to new data and reliably predict pollutant $PM_{2.5}$ concentrations and their spatial structure, which informs on the presence and lethality of SARS-CoV-2.

5 DISCUSSION

The capability of measuring, monitoring, and predicting pollutant concentrations and their behavior is of crucial importance. Currently, the database established by [18] containing estimates of monthly pollutant concentrations over several countries is among the best sources for this information.

The goal of the present work was to learn and infer air pollutants, in particular $PM_{2.5}$, particulate matter 2.5, NO_3 , Nitrate Pollution, and NH_4 , Atmospheric Ammonium, at higher temporal and spatial resolution than has been done previously and to do so using readily available satellite imagery. For this, we proposed the use of a fully convolutional neural network known as a U-net that can, unlike previous work, predict these pollutant values at a pixel-level high-resolution instead of only predicting a single value for an entire region.

Additionally, we have demonstrated an important practical application of our approach. SARS-CoV-2 lethality has been concretely

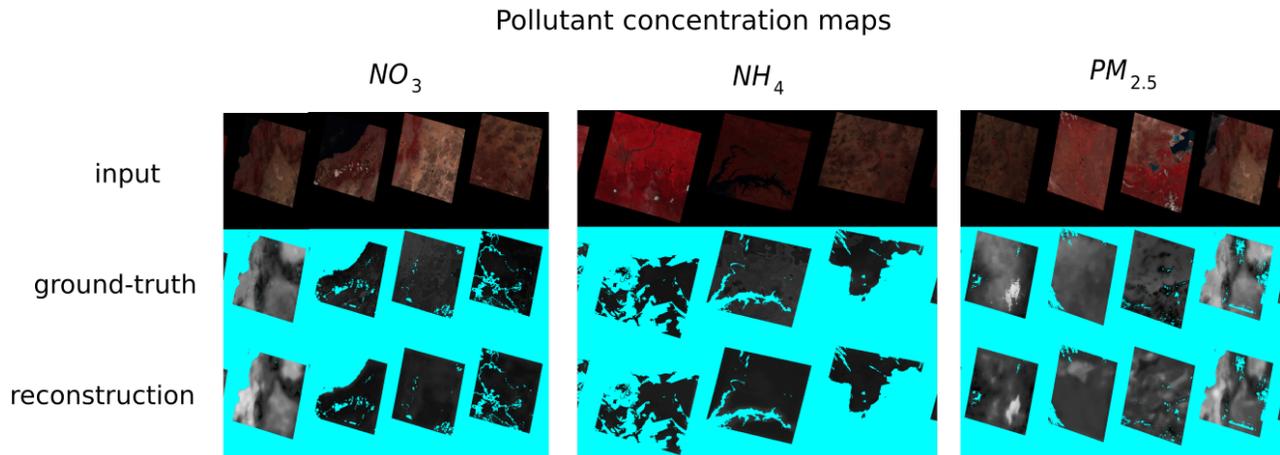


Figure 4: Examples of pollutant concentrations predicted by the U-net for new temporal data at locations used during training, respectively for NO_3 , NH_4 , and $PM_{2.5}$. The network is able to reconstruct the ground-truth concentrations of each pollutant, from which it follows, it can capture the spatial structure of these pollutants in the given satellite image.

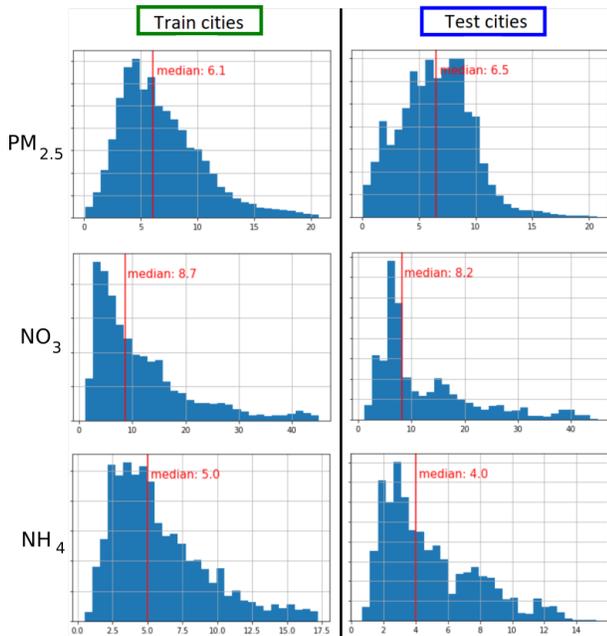


Figure 5: Distributions of the average pollutant concentration values of ground-truth training and test data, calculated from values within respectively the green and blue squares in Fig. 2 for $PM_{2.5}$, NO_3 , and NH_4 . The corresponding training and test ground-truth values show similar distributions for all 3 pollutants.

linked to the concentration of pollutant particulate matter, where a slight increase in $PM_{2.5}$ can drastically enhance the morbidity rate. Reliable predictions of the spatial structure and concentrations of $PM_{2.5}$ can help identify regions where SARS-CoV-2 may

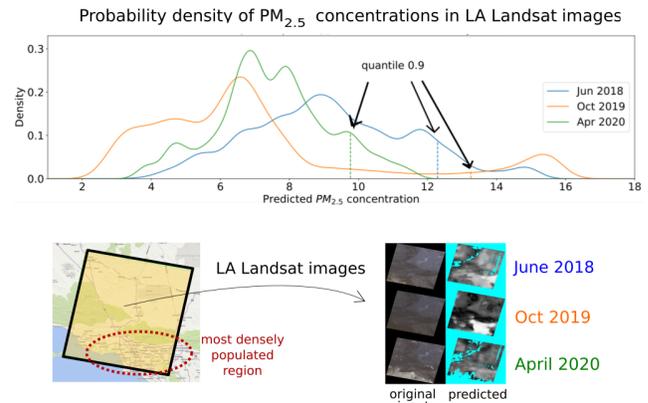


Figure 6: Density of predicted $PM_{2.5}$ concentrations in Los Angeles for time periods before and during the COVID-19 outbreak and subsequent lock-down (top) and corresponding Landsat images with their predicted $PM_{2.5}$ pollution maps (bottom). As expected, we see a significant reduction in the 0.9-quantile of the predicted $PM_{2.5}$ pollution distributions between April 2020 (immediately after the lock-down began) and the previous years. The predicted $PM_{2.5}$ maps match this trend, showing relatively less pollution in 2020 (darker pixels). Interestingly, yet unsurprisingly, the most densely populated regions of LA (encircled in red on the map) are predicted to have higher $PM_{2.5}$ concentrations for all of the years, as shown in the predicted $PM_{2.5}$ pollution maps.

be particularly transmissible and lethal can provide critical advice to implementable public health strategies.

In future work, we plan to include other types of pollutants in training the network. Furthermore, we intend to explore simultaneously learning multiple pollutant concentrations using a form

of multi-task learning [16]. Moreover, it would be interesting to explore synthetic data generated e.g. through simulation to gain better insights into the performance of our method for time intervals and spatial locations where true ground truth is missing.

ACKNOWLEDGMENTS

We would like to thank USGS for providing satellite imagery and Aaron Van Donkelaar, Randall Martin, and team for providing their data estimates of PM_{2.5} exposure. Additionally, we thank Drew Lipman for his inspirational discussions.

REFERENCES

- [1] [n.d.]. Amazon AWS Landsat dataset. <https://registry.opendata.aws/landsat-8/>. Accessed: 2020-06-06.
- [2] [n.d.]. Landsat World Reference System Description. <https://landsat.gsfc.nasa.gov/the-worldwide-reference-system/>. Accessed: 2020-06-06.
- [3] [n.d.]. PM_{2.5} concentrations published by Aaron van Donkelaar. <ftp://stetson.phys.dal.ca/Aaron/V4NA02.MAPLE/ASCII/Monthly/PM25/>. Accessed: 2020-06-06.
- [4] [n.d.]. USGS Earth Explorer. <https://earthexplorer.usgs.gov/>. Accessed: 2020-06-06.
- [5] 2016. *Landsat—Earth observation satellites* (version 1.0: originally posted november 25, 2015; version 1.1: august 22, 2016; version 1.2: april 8, 2020 ed.). Technical Report. Reston, VA. 4 pages. Report.
- [6] Joel Akeret, Chihway Chang, Aurelien Lucchi, and Alexandre Refregier. 2017. Radio frequency interference mitigation using deep convolutional neural networks. *Astronomy and Computing* 18 (2017), 35–39.
- [7] Mario Coccia. 2020. Two mechanisms for accelerated diffusion of COVID-19 outbreaks in regions with high intensity of population and polluting industrialization: the air pollution-to-human and human-to-human transmission dynamics. *medRxiv* (2020).
- [8] E. Conticini, B. Frediani, and D. Caro. 2020. Can atmospheric pollution be considered a co-factor in extremely high level of SARS-CoV-2 lethality in Northern Italy? *Environ. Pollut.* 261 (Jun 2020), 114465.
- [9] Hao Dong, Guang Yang, Fangde Liu, Yuanhan Mo, and Yike Guo. 2017. Automatic Brain Tumor Detection and Segmentation Using U-Net Based Fully Convolutional Networks. In *Medical Image Understanding and Analysis*, María Valdés Hernández and Víctor González-Castro (Eds.). Springer International Publishing, Cham, 506–517.
- [10] Steve Foga, Pat Scaramuzza, Song Guo, Zhe Zhu, Ronald Dilley, Tim Beckman, Gail Schmidt, John Dwyer, M Hughes, and Brady Laue. 2017. Cloud detection algorithm comparison and validation for operational Landsat data products. *Remote Sensing of Environment* 194 (2017), 379–390.
- [11] Tongwen Li, Huanfeng Shen, Qiangqiang Yuan, Xuechen Zhang, and Liangpei Zhang. 2017. Estimating Ground-Level PM_{2.5} by Fusing Satellite and Station Observations: A Geo-Intelligent Deep Learning Approach. *Geophysical Research Letters* 44, 23 (2017), 11,985–11,993. <https://doi.org/10.1002/2017GL075710>
- [12] Franco Manessi and Alessandro Rozza. 2018. Learning Combinations of Activation Functions. *CoRR* abs/1801.09403 (2018). arXiv:1801.09403 <http://arxiv.org/abs/1801.09403>
- [13] Christian Payer, Darko Štern, Horst Bischof, and Martin Urschler. 2016. Regressing Heatmaps for Multiple Landmark Localization Using CNNs. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*, Sebastien Ourselin, Leo Joskowicz, Mert R. Sabuncu, Gozde Unal, and William Wells (Eds.). Springer International Publishing, Cham, 230–238.
- [14] N. Rijal, R. T. Gutta, T. Cao, J. Lin, Q. Bo, and J. Zhang. 2018. Ensemble of Deep Neural Networks for Estimating Particulate Matter from Images. In *2018 IEEE 3rd International Conference on Image, Vision and Computing (ICIVC)*. 733–738.
- [15] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. *CoRR* abs/1505.04597 (2015). arXiv:1505.04597
- [16] Sebastian Ruder. 2017. An Overview of Multi-Task Learning in Deep Neural Networks. *CoRR* abs/1706.05098 (2017). arXiv:1706.05098 <http://arxiv.org/abs/1706.05098>
- [17] Huanfeng Shen, Tongwen Li, Qiangqiang Yuan, and Liangpei Zhang. 2018. Estimating Regional Ground-Level PM_{2.5} Directly From Satellite Top-Of-Atmosphere Reflectance Using Deep Belief Networks. *Journal of Geophysical Research: Atmospheres* 123, 24 (2018), 13,875–13,886. <https://doi.org/10.1029/2018JD028759>
- [18] Aaron van Donkelaar, Randall V. Martin, Chi Li, and Richard T. Burnett. 2019. Regional Estimates of Chemical Composition of Fine Particulate Matter Using a Combined Geoscience-Statistical Method with Information from Satellites, Models, and Monitors. *Environmental Science & Technology* 53, 5 (2019), 2595–2611.
- [19] Xiao Wu, Rachel C. Nethery, Benjamin M. Sabath, Danielle Braun, and Francesca Dominici. 2020. Exposure to air pollution and COVID-19 mortality in the United States: A nationwide cross-sectional study. *medRxiv* (2020).
- [20] Wei Yao, Zhigang Zeng, Cheng Lian, and Huiming Tang. 2018. Pixel-wise regression using U-Net and its application on pansharpening. *Neurocomputing* 312 (2018), 364 – 371.
- [21] Bo Zhang, Meng Zhang, Jian Kang, Danfeng Hong, Jian Xu, and Xiaoxiang Zhu. 2019. Estimation of PM_x Concentrations from Landsat 8 OLI Images Based on a Multilayer Perceptron Neural Network. *Remote Sensing* 11, 6 (2019). <https://doi.org/10.3390/rs11060646>
- [22] Chao Zhang, Junchi Yan, Changsheng Li, Hao Wu, and Rongfang Bie. 2018. End-to-End Learning for Image-Based Air Quality Level Estimation. *Mach. Vision Appl.* 29, 4 (May 2018), 601–615. <https://doi.org/10.1007/s00138-018-0919-x>
- [23] Tongshu Zheng, Michael H. Bergin, Shijia Hu, Joshua Miller, and David E. Carlson. 2020. Estimating ground-level PM_{2.5} using micro-satellite images by a convolutional neural network and random forest approach. *Atmospheric Environment* 230 (2020), 117451. <https://doi.org/10.1016/j.atmosenv.2020.117451>