Predicting Atmospheric Air Pollution: A Convolutional-Transformer Approach for Spatial and Temporal Analysis of PM2.5

Janmesh Kalra*, Pratyush Muthukumar*, Shaurya Pathak*, Kabir Nagrecha*, Hirad Hosseini*,
Dawn Comer**, Navid Amini*, Jeanne Holm**, and Mohammad Pourhomayoun*
* Department of Computer Science, California State University Los Angeles, Los Angeles CA,
[pmuthuk2, knagrec2, namini, mpourho]@calstatela.edu,
** City of Los Angeles, Los Angeles CA,
[dawn.comer, jeanne.holm]@lacity.org

***Abstract*— 6.7 million premature deaths occur annually due to household and ambient air pollution. Air pollution affects individuals globally and can be derived from a variety of factors including household cooking fuel, motor vehicles, industrial practices, and natural fires. To tackle this global crisis, our research focuses on understanding the spatial and temporal patterns between air pollutants to predict future levels of air pollutants. Our approach uses a novel deep learning methodology that involves a spatiotemporal Convolutional-Transformer architecture (ConvTransformer) We harnessed multiple data sources as inputs to our ConvTransformer, including several remote sensing instruments such as NASA's Terra+Aqua satellites with the Moderate Resolution Imaging Spectroradiometer (MODIS), and the European Space Agency's Sentinel-5P satellite with the Tropospheric Monitoring Instrument (TROPOMI), gridMET ground-based meteorological remote-sensing data from North American Land Data (NLDAS-2), US Census TIGER Roadways, and ground-level PM2.5 sensing data sourced from EPA AirNow. Our approach shows a 38.8% decrease in 5-frame average SSIM error compared to existing modern deep learning architectures utilizing remote-sensing data, satellite imagery, and ground-level data for PM2.5 prediction.**

I. Introduction

Industrialization and urbanization worldwide are causing an influx of air pollution, which in turn is leading to negative impacts on human health, deforestation, and habitat destruction. The trend to move towards urban cities and Central Business Districts, particularly in developing countries, is leading to record-breaking levels of PM2.5 year after year [Yang et al., 2018]. In fact, according to Zhou et al. [2017], over the past decade, the number of deaths caused by the harmful effects of ambient PM2.5 have increased by 23%. PM2.5's ability to travel deep into human respiratory tracts can cause premature death to those with pre-existing heart and lung conditions. Its effects on human health transcend across age groups and can worsen diseases such as COPD, asthma, pulmonary fibrosis, and pneumonia [Xing et al., 2016].

Our examination of PM2.5 at the hourly level allows us to understand the global issue at its source and predict future patterns to be able to form a solution before the situation aggravates. Our paper relies on a model architecture that utilizes deep learning methodology to predict PM2.5 both spatially and temporally (spatial-temporal). Using a model capable of identifying both spatial and temporal patterns, we can capture how PM2.5 moves across a geographic region as well as how it contrasts with data from years prior.

Our predictive performance is set apart from traditional techniques because we utilize multisource big data from ground-level sites, atmospheric remote sensing information, and robust satellite imagery of air pollution and meteorological features. Our complex deep learning architecture applies the state-of-the-art Transformer architecture with Multi Head Attention in conjunction with Convolution to learn time-invariant features across gridded data.

The ConvTransformer architecture utilizes various datasets of time-series images. The ConvTransformer is fed video-like inputs which represent various input features of our dataset, including satellite imagery, remote-sensing data, and ground-level interpolated data masks over time and space. Fundamentally, the ConvTransformer architecture extends the Transformer architecture with the Convolution layer. We rely on the following equation to perform 2D Convolution:

$$O = \sum_i \sum_j I(i,j) \cdot K(x-i, y-j) + B$$

where O is representative of an output feature map, I is the input feature map, K is the kernel, i is the height of the input channel, j is width of the input channel, and B is bias. We pass our output feature map into the activation function, Rectified Linear Unit (ReLU).

$$f(x) = \max(0, x)$$

ReLU is a piecewise activation function that has a linear relationship with the data when the values are positive, and always zero when the values are negative. RELU prevents issues with vanishing gradient descent and is computationally efficient, which suits it better than other activation functions such as the sigmoid or hyperbolic tangent.

After performing 2D Convolution on the data, we can generate 2D tensors of the input as layers in our model. We then implement the Transformer layer into our model which

979-8-3503-2759-5/23/$31.00 ©2023 IEEE
DOI 10.1109/CSCE60160.2023.00261

is primarily characterized by Multi Head Self Attention. The Transformer architecture is a retentive model which utilizes the Attention mechanism to learn patterns across time-series data effectively [Vaswani et al., 2017]. Transformers parse the data at once and have an attention mechanism that gives us context about each position in the data. One upside to the Transformer model compared to its predecessors is its ability to gather information quickly about each part of the sequence and perform operations in parallel due to the multi-headed nature of the Attention mechanism [Medsker and Jain, 2001, Hochreiter and Schmidhuber, 1997].

The Multi-Head Attention Mechanism relies on several key equations to be able to operate and quickly understand our input sequence with the "attention" identifiers. The layer begins by taking our input sequence and multiplying specifically learned weight matrices to obtain the Queries (Q), Keys (K), and Values (V) vectors. We then split up these vectors into different heads, where each head contains a Query, Key, and Value and the model contains n number of heads. Each head in our Multi-Head Attention layer is assigned an attention score which is a representation of how similar the query and key in the head are. We use the following equation to compute Attention scores:

$$A_i = \text{softmax}\left(\frac{Q_i K_i^T}{\sqrt{d_k}}\right),$$

where A_i is the attention score for query-key pair i, Q_i denotes query i, K_i denotes key i, and d_k denotes the dimensionality of the key vector.

We use this attention scores to generate a weighted sum between the attention scores and the projected values of each head. It is denoted using the following equation:

$$O_i = A_i V_i$$

We finally then model the linear relationship between each head, and multiply the weighted sum with a matrix of model weight and add a bias term. This completed sum is the Multi Head Attention output:

$$O = [O_1, O_2, \ldots, O_h]\, W_O + b_O$$

II. Methodology

To predict PM2.5 at the hourly level for the urban area of Los Angeles, we rely on various remote sensing instruments and data sets sourced by NASA, the European Space Agency, and the Environmental Protection Agency. We utilize remote-sensing air pollution, roadway outlines, ground-based air pollution, atmospheric air pollution satellite imagery, and ground-based meteorological gridded data in our ConvTransformer architecture.

Our Los Angeles County study area falls within a 70 km by 70 km area, between the latitudes of $33.5°N - 34.5°N$ and longitudes of $117.5°W - 118.75°W$. Figure 1 describes the geographic bounds of LA County along with the locations of the ground-based PM2.5 and meteorological sensors used as ground-truth labels for prediction.

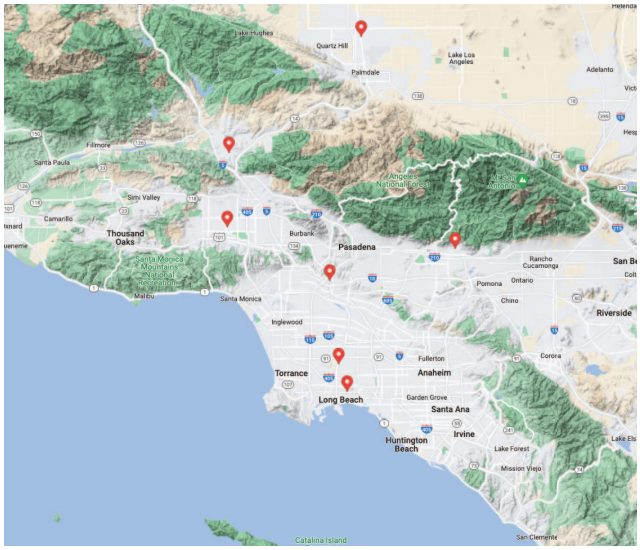

Fig. 1: Ground-level Meteorological and Air Pollution Site Locations shown as red markers.

Los Angeles typically demonstrates a Mediterranean climate pattern, but due to recent cycles of rainfall and dry spells, the climate has adapted to become a large source of wildfire. In particular, multiple parts of California experienced heavy wildfire seasons during the year 2020. Wildfire smoke data demonstrates a direct correlation with PM2.5, making it very useful as an incorporation to our model. One-third of total atmospheric emissions are attributed to wildfire smoke emissions, and approximately 90% of the emissions from wildfires are fine particles like PM2.5 [Vedal and Dutton, 2006].

By combining direct PM2.5 historical data and indirect data-characterizing factors that are associated with PM2.5, our remote-sensing and ground-based data forms a comprehensive approach to the hourly prediction of PM2.5. Air pollution cannot be effectively linearly predicted using a statistical approach because of the way it evolves and changes from various environmental factors [Rosenlund et al., 2008]. Instead, we provide the ConvTransformer architecture with data to deduce how the environment and topographical features affect patterns of PM2.5 spatiotemporally in Los Angeles county. The various remote-sensing instruments are recorded at differing sea levels, urban structures, and topographies, thus providing our model with robust multifaceted information.

A. Dataset

Our remote sensing and ground-based inputs were geographically bounded to the 70km x 70km area of Los Angeles County. We utilized three years of hourly data from January 1, 2018 to December 31, 2020. We collected 26304 samples corresponding to 24 hourly samples for the 1096 days of data from January 1, 2018 to December 31, 2020.

To add definition to our prediction model, we utilize the TIGER roadways shapefile which outlines all of the major highways and roads in LA County. The US Census provides the TIGER dataset in shapefiles for users to visualize traffic- and ground-based data. We use it in our model to explore the dispersion and transition of PM2.5 across LA County [McCormack and Blaine, 1990].

For the ground-based meteorological data input to our model, we utilized the gridMET meteorological data source which provides climate variables that are mapped onto a geographic grid with climatically-aided interpolation [Abatzoglou, 2013]. gridMET has several bands applicable to our prediction model including precipitation accumulation, which can help us understand how dry the climate is and model the relationship between PM2.5 before and after rainfall. Another important variable obtained from gridMET is wind speed, which is useful when understanding how a wildfire moves and spreads across our geographic region. We also use the wind direction band from the gridMET data source, which provides information regarding the directionality and movement of future PM2.5 patterns. A visualization of these input features is described in Figure 2.

In our predictive model, we incorporated the NASA MODIS Multi-Angle Implementation of Atmospheric Correction (MAIAC) which is a remote imaging spectrometer data source. From this source, we utilized the Aerosol Optical Depth (AOD), which quantifies how atmospheric pollutants and air pollutants like PM2.5 block the transmission of light. It explains how particles move around in the atmosphere, and how their dispersion impacts solar radiation. AOD is measured using wavelengths of light, and it has demonstrated a very high and direct correlation with ambient air pollution. The AOD used in our ConvTransformer model has a wavelength of 0.47 μm and it provides a spatial resolution of 1 km/pixel for an area of 1200km by 1200km.

We also utilized the Sentinel 5-P satellite from the European Space Agency (ESA) which has the TROPOMI instrument. The satellite operates using a heliosynchronous orbit, which means that upon each orbit the satellite passes through the same location at the same local time. The Sentinel-5P Satellite transmits all of its data to the ground-based station in real-time through X-band frequency, and it is made publicly available through the Copernicus program. The TROPOMI instrument of the Sentinel 5-P satellite gives us insight into many different atmospheric measurements such as ozone, nitrogen dioxide, sulfur dioxide, methane, formaldehyde, carbon monoxide, aerosols, and ultraviolet radiation. These atmospheric measurements typically affect the creation of PM2.5 and they usually serve as the precursor to the air PM2.5. We find that products such as nitrogen dioxide, methane, carbon monoxide, etc. can combust to form the components of PM2.5. TROPOMI has a high spatial resolution of 7km x 3.5km per pixel.

To understand wildfires and heat we use two remote sensing instruments: NASA MODIS and MERRA-2. NASA MODIS provides us with the Fire Radiative Power (FRP) band which measures the thermal energy released by fires. As a remote

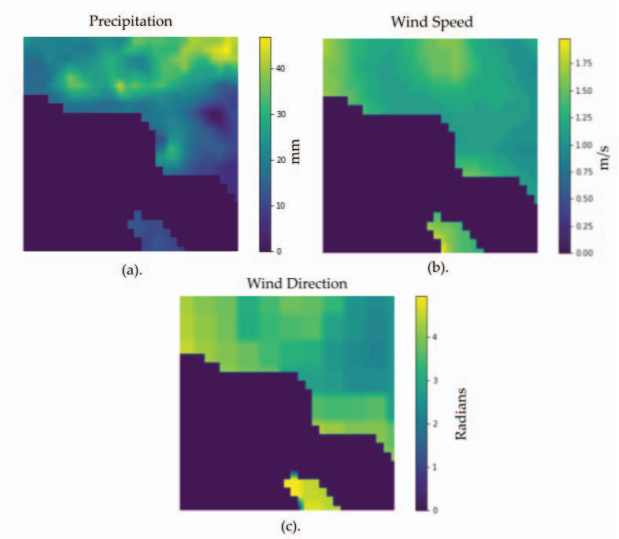

Fig. 2: Sample hourly gridMET ground-level gridded data in Los Angeles County. (a). Wind direction, (b). Wind speed, (c). Precipitation

sensing instrument, the FRP band from MODIS tells us exactly how intense a wildfire is and its geographic area covered. FRP obtains a metric that cannot be measured using a ground-based sensor, and it can be also useful in the prediction of future wildfires. Once we can detect where an active wildfire is in progress we can conclude that the levels of PM2.5 are going to change and spread. Wildfire smoke emissions include the incomplete combustion of vegetation, trees, and organic matter which are forms of PM2.5. Smoke acts as a transportation system to allow the PM2.5 to spread over an area, through an analysis of where a fire is (FRP) and the wind velocity (gridMET) we gain a complex understanding of PM2.5 behavior. FRP is measured at a resolution of 1km x 1km per pixel and is useful for at looking how a fire moves on a large scale. MODIS FRP has a frequent orbit and is especially useful when it comes to something as dynamic as wildfires. FRP measures in the thermal infrared wavelength spectrum and it uses light in the range of 2070 μm to 3200 μm.

Aside from remote sensing images of wildfires, we can utilize the NASA MERRA-2 remote sensing satellite to observe heat features that are correlated with wildfires and PM2.5. From MERRA-2 we found that the Planetary Boundary Layer (PBL) height, surface air temperature, and surface exchange coefficient for heat would aide in our analysis. Images of the PBL Height give us a metric to quantify the distance between ground level and the lowest point of the atmosphere, and the lowest point of the atmosphere is constantly changing. The PBL is impacted by weather conditions such as solar radiation and wind patterns, and its size is depend on current climate patterns [Guo et al., 2021]. We know that changing climate patterns have a link with PM2.5, which fits PBL into our

study. PBL can change from rainfall, and it is responsible for transporting moisture from the air to the atmosphere. A high PBL Height indicates that rainfall is about to come, which we can use to explore PM2.5 before and after rainfall. PBL Height also changes based on surface features, such as the terrain being mountainous or flat and if the area is covered with dense or sparse vegetation. PBL Height can range from several kilometers to only 500 meters depending on conditions, and it does correlate with fire and heat. Surface exchange coefficient (C_H) defines the transfer of heat between land surface and the overarching atmosphere. Learning about how heat energy is transferred between these two areas is an auxiliary to wildfires and radiation, and we can observe that higher C_H result in wildfire occurrences which have correlations with PM2.5 [Chen and Zhang, 2009].

Figure 3 describes an overview of the various global remote-sensing data and satellite imagery used within this model.

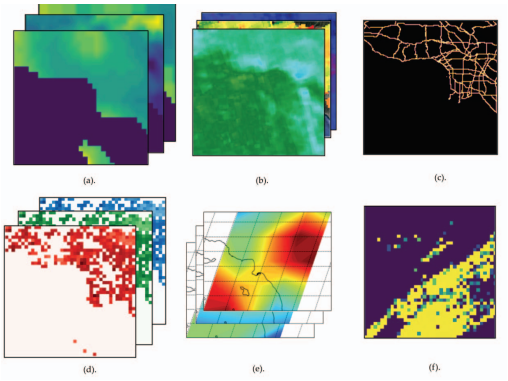

Fig. 3: ConvTransformer Input Features. (a). gridMET input features, (b). TROPOMI Satellite Imagery, (c). TIGER LA County Roadways Shapefile, (d). Sentinel-2 RGB Satellite Imagery, (e). MERRA-2 Gridded Fire Features, (f). NASA MAIAC MODIS AOD Remote-Sensing Data

Using the EPA Airnow API from the California Air Resources Board AQMIS2 portal we can collect PM2.5 data at the hourly level. There are 9 ground-based PM2.5 monitoring sites in our defined Los Angeles region, and we use this data to form hourly grids with the PM2.5 values demarcated on each location of the sensor in a 40 x 40 grid.

We apply an advanced elevation population interpolation technique described in our prior work to generate dense PM2.5 ground-level grids from sparse air pollution sensor measurements [Muthukumar et al., 2022b]. We apply inverse distance weighted interpolation with weights corresponding to values in a static elevation and population map. The elevation map is derived from digital elevation data sourced from the Sentinel Hub API [Braun, 2021]. The population map is sourced from the Global Human Settlement Layer accessed through the EO Browser via Sentinel Hub [Pesaresi et al., 2016].

To validate our model's performance, we evaluate our PM2.5 predictions against sites with over 95% of hourly validated data from the 3 years of our dataset length. In total, we find that all 9 sensors in Los Angeles can be used for validation purposes.

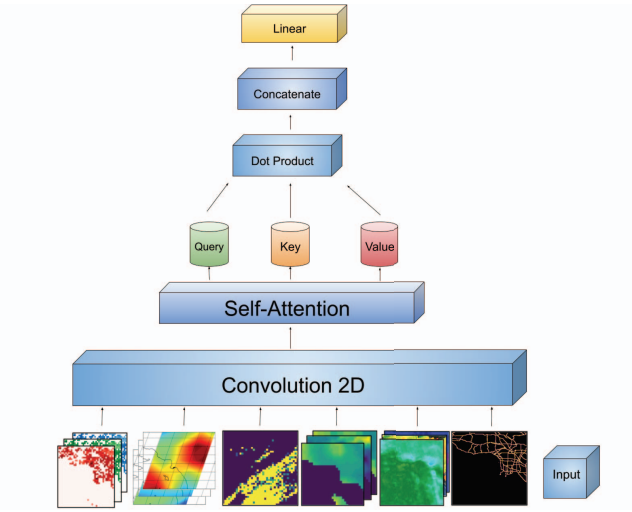

Fig. 4: Visualization of input filters for our ConvTransformer model and the bundling method applied. One sample consists of 24 frames, where the label sample is staggered 1 frame ahead of the input feature sample

B. Model Architecture

To use our ConvTransformer architecture, we shaped the input data into a five-dimensional tensor with dimensions (sample, frame, row, column, filter). A visualization of the ConvTransformer architecture is shown in Figure 4.

We downsampled the remote sensing satellite imagery data sets to a 40-by-40 pixel resolution, corresponding to a 40-row by the 40-column array to put into the 5D tensor. The row and column parameters of the 5D tensor for our remote-sensing imagery are defined as the 2D image, and the filter is the RGB value of the image.

Our ground-based air pollution sensor data was mapped onto a 40-by-40 pixel grid with the corresponding latitude & longitude of each monitoring site represented as a point on the grid.

For our 5D tensor, we bundled our input frames into smaller samples, where 24 consecutive frames were considered to be a part of one sample and each frame is representative of a single timestep in our hourly data frequency. One entire bundle represents an entire day's worth of data, and we staggered the hours used in our bundle so that the data remain continuous. For example, bundle 1 contains hours 1-24, but bundle 2 contains hours 2-25.

We have samples for 1096 days (3 years) and 24 data points each day, which gives us 26,304 images to use in our model. When we construct our 5D tensor, we obtain the shape of (26,304, 24, 40, 40, 10) where 26,304 is the amount of data, 24 is the size of each bundle, 40 x 40 is the size of each frame and 10 is the number of filters. The 10 filters we use in our tensor

include one filter for the MAIAC MODIS AOD imagery, three filters for the MERRA-2 fire features (PBL height, surface temperature, and surface exchange coefficient for heat), 1 filter for the MODIS FRP imagery, three filters for the TROPOMI data of air pollutants (nitrogen dioxide, carbon monoxide, and methane), one filter for gridMET meteorological data, and one filter for the ground-based PM2.5 sensor data

For the model implementation, we started by designing a custom ConvTransformer layer from the Keras library. The general structure is inherited from tf.keras.layers.Layer module, which is then initialized to perform the convolution operation on our input followed by Multi-Head Attention and normalization. We created two custom ConvTransformer layers with key dimensions of 15 and 30, and heads of 3 and 6 respectively. The Transformer layers help us identify broader interactions between different parts of the input. After that, we created two other layers and applied 3D convolution to the data, which are intended to capture local patterns in our data. Our Conv3D layers are directly imported from the Keras library, and the parameters are set to match our 10 filters. To recap we have the following layers in our model:

1. `ConvTransfomerLayer(k_dim = 15, head = 3)`
2. `ConvTransfomerLayer(k_dim = 30, head = 6)`
3. `Conv3D(filters= 10, activation='relu')`
4. `Conv3D(filters= 10, activation='relu')`

III. Results

Our model predicts hourly PM2.5 within LA County in units of micrograms per cubic meter ($\frac{\mu g}{m^3}$) using 24 prior frames to predict 24 future frames using multisource ground-level sensor data, satellite imagery, meteorological features, remote-sensing wildfire data, and atmospheric air pollution derived features such as Aerosol Optical Depth (AOD). We use 26304 samples of hourly data from January 1 2018 to November 30 2020 as training data and evaluate our prediction on a test dataset of 1320 samples (24 samples for 55 days) from November 6 2020 to December 31, 2020.

We used the Root Mean Square Error (RMSE) and Normalized Root Mean Square Error (NRMSE) error. RMSE and NRMSE are calculated as

$$\text{RMSE} = \sqrt{\sum_{i=1}^{n} \frac{(\hat{y}_i - y_i)^2}{n}};\ \text{NRMSE} = \frac{\text{RMSE}}{\bar{y}}$$

where n is the number of observations, $\hat{y}$ is the predicted value, y is the ground truth, and $\bar{y}$ is the mean of the test data.

We also used the Structural Similarity Index Measurement (SSIM) metrics to qualitatively assess the model performance. SSIM is an effective metric applied in image processing to quantify general similarities in visual features. SSIM values range from 0.0 to 1.0 where an SSIM score of 1.0 between two images implies they are visually identical. For a ground truth pixel value p and predicted value $\hat{p}$, the SSIM of a pixel is

$$\text{SSIM}(p, \hat{p}) = \frac{2\mu_p\mu_{\hat{p}} + c_1}{(\mu_p^2 + \mu_{\hat{p}}^2 + c_1)\sigma_p^2 + \sigma_{\hat{p}}^2 + c_2},$$

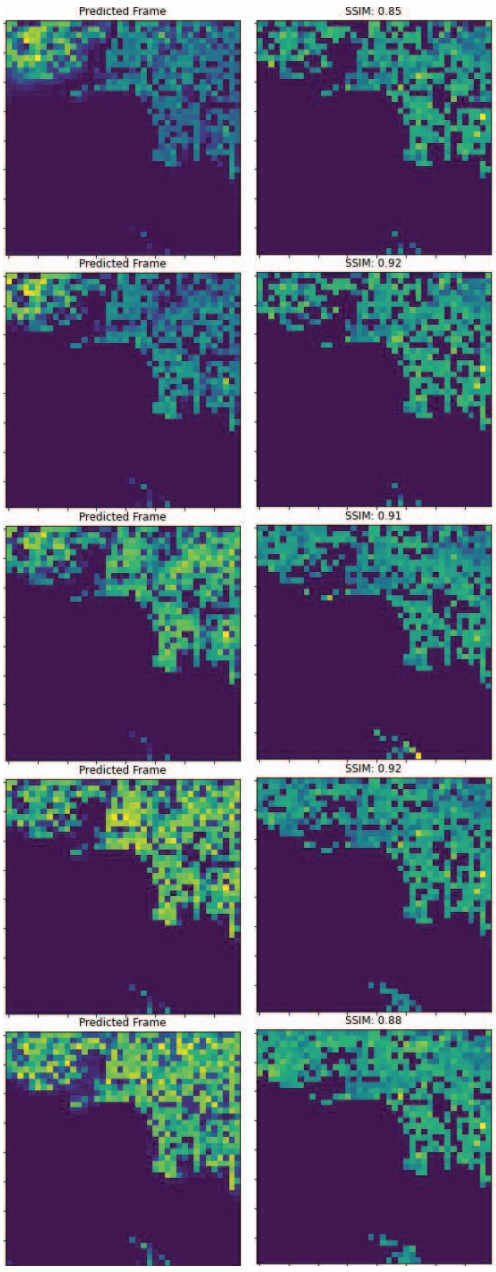

Fig. 5: Visualization of first five frames of ConvTransformer predictions. Left column is predicted PM2.5 frame, right column is ground-truth PM2.5 data. SSIM values per frame are reported above each row.

where c_1 and c_2 are constants relating to the relative noise of an image [Brunet et al., 2011]. For the entire image I, the SSIM score is defined as

$$SSIM(I) = \sum_{p \in I} \frac{2\mu_p\mu_{\hat{p}} + c_1}{(\mu_p^2 + \mu_{\hat{p}}^2 + c_1)\sigma_p^2 + \sigma_{\hat{p}}^2 + c_2}.$$

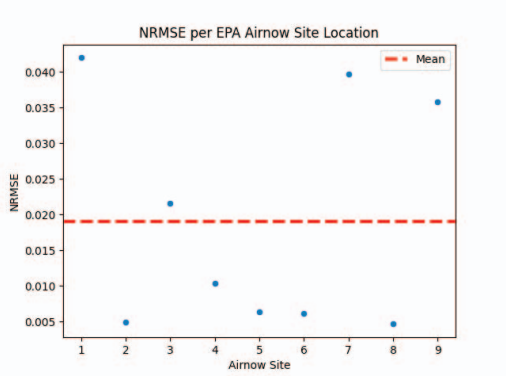

Fig. 6: Scatterplot visualizing NRMSE values mapped per EPA Airnow Site Monitoring Station along with NRMSE averaged across all sites

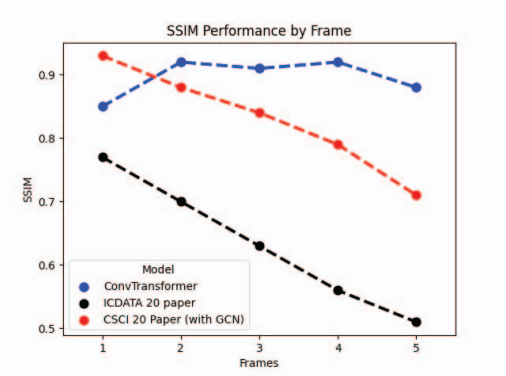

Fig. 7: SSIM values visualized per frame for ConvLSTM (black), MeteoGCN-ConvLSTM (red), and our ConvTransformer (blue)

To optimize the values of our model parameters we used the Adam Optimizer, a variant of Stochastic Gradient Descent. Adam performs well on image-classification tasks and is an adaptive gradient algorithm that tunes the learning rate per parameter.

A. *ConvTransformer PM2.5 Hourly Prediction Results*

In this section, we explore the results obtained from our ConvTransformer designed for hourly prediction of PM2.5 in Los Angeles County both spatially and temporally. Table I describes location-based NRMSE metrics for each ground-level PM2.5 site monitoring station. Figure 6 provides a visualization of per-site NRMSE values against the NRMSE calculate across all sites.

Sensor Location	NRMSE
Airnow Site 1	0.041967
Airnow Site 2	0.004857
Airnow Site 3	0.021533
Airnow Site 4	0.010322
Airnow Site 5	0.006353
Airnow Site 6	0.006146
Airnow Site 7	0.039644
Airnow Site 8	0.004643
Airnow Site 9	0.035833

TABLE I: NRMSE error values averaged over 24 frame bundles of test set for each sensor location in Los Angeles

To quantify a more precise similarity measure between our ground truth and predicted values, we have calculated the SSIM values of 5 predicted frames from our ConvTransformer. Figure 5 describes a visualization of the first five frames of the ConvTransformer test set, where the left column denotes the predictions and the right column denotes the ground truth values. Table II describes the first five frames' SSIM values between the predicted and ground truth data.

IV. Conclusion

In our spatiotemporal study, we leveraged modern deep learning methods to achieve precise hourly predictions of PM2.5 across our defined geographical region of Los Angeles County. Our approach contained a rich array of data sources including ground-based meteorological and sensor data, remote-sensing satellite instruments, and wildfire/heat data.

Our model brings in the best of both worlds by using a Convolutional Neural Network to get the most out of our spatial data and retain the information using the long-term dependency power of Transformers. We incorporated state-of-the-art machine learning techniques to abstract information from ground-based and remote data sources including US Census TIGER roadway outlines, ground-based air pollution data, gridded ground-based meteorological data, satellite images of Nitrogen Dioxide, Methane, and Carbon Monoxide; AOD, FRP, PBL Height, Surface Temperature, and Surface Exchange Coefficient of Heat from NASA remote-sensing instruments.

Frame	SSIM
Frame 1	0.85
Frame 2	0.92
Frame 3	0.91
Frame 4	0.92
Frame 5	0.88

TABLE II: SSIM values over first five frames of test set for ConvTransformer predictions compared against ground truth

First Five Frame SSIM Scores					
	Frame				
Model	1	2	3	4	5
ConvTransformer	0.85	0.92	0.91	0.92	0.88
ConvLSTM (ICDATA '20)	0.77	0.7	0.63	0.56	0.51
MeteoGCN-ConvLSTM (CSCI '20)	0.93	0.88	0.84	0.79	0.71

TABLE III: Testing set first 5 frame SSIM comparison of ConvTransformer (ours), ConvLSTM (presented at ICDATA 20') and ConvLSTM with Meteorological GCN (presented at CSCI '20)

The ConvTransformer predicts PM2.5 hourly over Los Angeles County but can be altered to fit any geographic region across the globe. To evaluate our ConvTransformer, we compute the NRMSE and SSIM values of each sensor location/frame respectively. Upon comparison with our previous deep learning predictive algorithms namely the CNN-LSTM with GCN, described in Muthukumar et al. [2020c], we find that the ConvTransformer can demonstrate a sharp increase in performance. Initially, the ConvTransformer exhibits a pattern of requiring a greater amount of frames to effectively predict PM2.5 spatiotemporal patterns in the dataset, but its retention levels outperform our other models utilizing the ConvLSTM architecture by 0.164 SSIM on average [Cocom et al., 2020, Muthukumar et al., 2021b,a, 2020a, Nagrecha et al., 2020, Muthukumar et al., 2022a,b]. Table III describes the comparison of SSIM values of the first five frames from the testing set for the ConvTransformer model, our prior work on ConvLSTM models presented at ICDATA '20, and our prior research on ConvLSTM with Meteorological Graph Convolutional Network (GCN) presented at CSCI '20 [Muthukumar et al., 2020b,a]. Figure 7 describes a visualization of SSIM performance of ConvTransformer against our baseline models. Our results show that the ConvTransformer has a 71.6% decrease in 5-frame average SSIM error compared to the ConvLSTM model and a 38.8% decrease compared to the MeteoGCN-ConvLSTM model.

V. Future Work

We would like to expand the set of ground-level PM2.5 site monitoring stations by including community maintained sensors while accounting for data uncertainty under error. Additionally, we hope to scale our model to multiple cities worldwide to better understand ambient air pollution at the global scale.

Our current ConvTransformer model predicts PM2.5 frames at a 40-by-40 pixel resolution, and we hope to increase this resolution by upscaling input features in the near future. By training on a higher resolution of data, our model will command more granularity in its predictions which in turn allow us to better understand hourly PM2.5 patterns in Los Angeles and other regions around the world.

Ambient air pollution is affected by many factors and our data sets can be expanded upon to provide a more robust model of authentic geophysical patterns in our study area. We seek to model the effects of commercial and residential pollution on PM2.5 and identify potential sub regions of focus within our study area. Ground-level PM2.5 also plays a major role in many respiratory diseases such as COPD, emphysema, and asthma which can be included in our predictive model as a downstream task to forecast public health risks and outbreaks.

References

John T Abatzoglou. Development of gridded surface meteorological data for ecological applications and modelling. *International Journal of Climatology*, 33(1):121–131, 2013.

Andreas Braun. Retrieval of digital elevation models from sentinel-1 radar data–open applications, techniques, and limitations. *Open Geosciences*, 13(1):532–569, 2021.

Dominique Brunet, Edward R Vrscay, and Zhou Wang. On the mathematical properties of the structural similarity index. *IEEE Transactions on Image Processing*, 21(4):1488–1499, 2011.

Fei Chen and Ying Zhang. On the coupling strength between the land surface and the atmosphere: From viewpoint of surface exchange coefficients. *Geophysical Research Letters*, 36(10), 2009.

Emmanuel Cocom, Pratyush Muthukumar, Jeanne Holm, Dawn Comer, Anthony Lyons, Irene Burga, Christa A Hasenkopf, Chisato Calvert, and Mohammad Pourhomayoun. Particulate Matter Forecasting in Los Angeles County with Sparse Ground-based Sensor Data Analytics. In *AGU Fall Meeting Abstracts*, volume 2020, pages A061–0004, 2020.

Jianping Guo, Jian Zhang, Kun Yang, Hong Liao, Shaodong Zhang, Kaiming Huang, Yanmin Lv, Jia Shao, Tao Yu, Bing Tong, et al. Investigation of near-global daytime boundary layer height using high-resolution radiosondes: first results and comparison with era5, merra-2, jra-55, and ncep-2 reanalyses. *Atmospheric Chemistry and Physics*, 21 (22):17079–17097, 2021.

Sepp Hochreiter and Jürgen Schmidhuber. Long Short-term Memory. *Neural computation*, 9(8):1735–1780, 1997.

Edward W McCormack and Lawrence W Blaine. State level transportation applications using tiger. Technical report, 1990.

Larry R Medsker and LC Jain. Recurrent neural networks. *Design and Applications*, 5:64–67, 2001.

Pratyush Muthukumar, Emmanuel Cocom, Jeanne Holm, Dawn Comer, Anthony Lyons, Irene Burga, Christa Hasenkopf, and Mohammad Pourhomayoun. Real-time Spatiotemporal NO2 Air Pollution Prediction with Deep Convolutional LSTM through Satellite Image Analytics. In *AGU Fall Meeting Abstracts*, volume 2020, 2020a.

Pratyush Muthukumar, Emmanuel Cocom, Jeanne Holm, Dawn Comer, Anthony Lyons, Irene Burga, Christa Hasenkopf, and Mohammad Pourhomayoun. Real-time Spatiotemporal Air Pollution Prediction with Deep Convolutional LSTM through Satellite Image Analysis. In *16th International Conference on Data Science (ICDATA '20)*, pages 317–328. Springer Nature, 2020b.

Pratyush Muthukumar, Emmanuel Cocom, Kabir Nagrecha, Jeanne Holm, Dawn Comer, Anthony Lyons, Irene Burga, Chisato Fukuda Calvert, and Mohammad Pourhomayoun. Satellite Image Atmospheric Air Pollution Prediction through Meteorological Graph Convolutional Network with Deep Convolutional LSTM. In *7th Annual Conference on Computational Science and Computational Intelligence (CSCI-ISAI '20)*. IEEE CPS, 2020c.

Pratyush Muthukumar, Emmanuel Cocom, Kabir Nagrecha, Dawn Comer, Irene Burga, Jeremy Taub, Chisato Calvert, Jeanne Holm, and Mohammad Pourhomayoun. Predicting

PM2.5 Atmospheric Air Pollution using Deep Learning with Meteorological Data and Ground-based Observations and Remote-sensing Satellite Big Data. *Air Quality, Atmosphere & Health*, pages 1–14, 2021a.

Pratyush Muthukumar, Kabir Nagrecha, Emmanuel Cocom, Dawn Comer, Irene Burga, Jeremy Taub, Chisato Calvert, Jeanne Holm, and Mohammad Pourhomayoun. Predicting PM2.5 Air Pollution using Deep Learning with Multisource Satellite and Ground-based Observations and Meteorological and Wildfire Big Data. In *AGU Fall Meeting Abstracts*, volume 2021, 2021b.

Pratyush Muthukumar, Kabir Nagrecha, Dawn Comer, Chisato Fukuda Calvert, Navid Amini, Jeanne Holm, and Mohammad Pourhomayoun. PM2.5 Air Pollution Prediction through Deep Learning Using Multisource Meteorological, Wildfire, and Heat Data. *Atmosphere*, 13(5), 2022a. ISSN 2073-4433. doi: 10.3390/atmos13050822. URL https://www.mdpi.com/2073-4433/13/5/822.

Pratyush Muthukumar, Shaurya Pathak, Kabir Nagrecha, Hirad Hosseini, Dawn Comer, Navid Amini, Jeanne Holm, and Mohammad Pourhomayoun. Multi-Pollutant Ground-level Air Pollution Prediction through Deep MeteoGCN-ConvLSTM. In *9th Annual Conference on Computational Science and Computational Intelligence (CSCI-ISAI '22)*. IEEE CPS, 2022b.

Kabir Nagrecha, Pratyush Muthukumar, Emmanuel Cocom, Jeanne Holm, Dawn Comer, Irene Burga, and Mohammad Pourhomayoun. Sensor-Based Air Pollution Prediction using Deep CNN-LSTM. In *2020 International Conference on Computational Science and Computational Intelligence (CSCI)*, pages 694–696. IEEE, 2020.

Martino Pesaresi, Daniele Ehrlich, Stefano Ferri, Aneta Florczyk, Sergio Freire, Matina Halkia, Andreea Julea, Thomas Kemper, Pierre Soille, Vasileios Syrris, et al. Operating procedure for the production of the global human settlement layer from landsat data of the epochs 1975, 1990, 2000, and 2014. *Publications Office of the European Union*, pages 1–62, 2016.

Mats Rosenlund, Francesco Forastiere, Massimo Stafoggia, Daniela Porta, Mara Perucci, Andrea Ranzi, Fabio Nussio, and Carlo A Perucci. Comparison of regression models with land-use and emissions data to predict the spatial distribution of traffic-related air pollution in rome. *Journal of exposure science & environmental epidemiology*, 18(2): 192–199, 2008.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Sverre Vedal and Steven J Dutton. Wildfire air pollution and daily mortality in a large urban area. *Environmental research*, 102(1):29–35, 2006.

Yu-Fei Xing, Yue-Hua Xu, Min-Hua Shi, and Yi-Xin Lian. The impact of pm2. 5 on the human respiratory system. *Journal of thoracic disease*, 8(1):E69, 2016.

Dongyang Yang, Chao Ye, Xiaomin Wang, Debin Lu, Jianhua Xu, and Haiqing Yang. Global distribution and evolvement of urbanization and pm2. 5 (1998–2015). *Atmospheric Environment*, 182:171–178, 2018.

Wei Zhou, Dongdong Tian, Jun He, Li Zhang, Xiuli Tang, Lijun Zhang, Yimei Wang, Lizhong Li, Jun Zhao, Xiaoyan Yuan, et al. Exposure scenario: another important factor determining the toxic effects of pm2. 5 and possible mechanisms involved. *Environmental pollution*, 226:412–425, 2017.