Environmental Research 271 (2025) 120948

Contents lists available at ScienceDirect

Environmental Research

journal homepage: www.elsevier.com/locate/envres

XIS-$PM_{2.5}$: A daily spatiotemporal machine-learning model for $PM_{2.5}$ in the contiguous United States

Allan C. Just [a,b,c,*], Kodi B. Arfer [a,c], Johnathan Rush [c], Alexei Lyapustin [d], Itai Kloog [c,e]

[a] Department of Epidemiology, Brown University School of Public Health, Providence, RI, USA
[b] Institute at Brown for Environment and Society, Brown University, Providence, RI, USA
[c] Department of Environmental Medicine and Public Health, Icahn School of Medicine at Mount Sinai, New York, NY, USA
[d] NASA Goddard Space Flight Center, Greenbelt, MD, USA
[e] The Department of Geography and Environmental Development, Ben-Gurion University of the Negev, Beer Sheva, Israel

ARTICLE INFO

Keywords:
XGBoost
Aerosol optical depth
Absolute error
Air pollution and health

ABSTRACT

Air-pollution monitoring is sparse across most of the United States, so geostatistical models are important for reconstructing concentrations of fine particulate air pollution ($PM_{2.5}$) for use in health studies. We present XGBoost-IDW Synthesis (XIS), a daily high-resolution $PM_{2.5}$ machine-learning model covering the contiguous US from 2003 through 2023. XIS uses aerosol optical depth from satellites and a parsimonious set of additional predictors to make predictions at arbitrary points, capturing near-roadway gradients and allowing the estimation of address-level exposures. We built XIS with a computationally tractable workflow for extensibility to future years, and we used weighted evaluation to fairly assess performance in sparsely monitored regions. Averaging across all years in site-level cross-validation, the weighted mean absolute error of predictions (MAE) was 2.09 $\mu g/m^3$, a substantial improvement over the mean absolute deviation from the median, which was 4.15 $\mu g/m^3$. Comparing XIS to a leading product from the US Environmental Protection Agency, the Fused Air Quality Surface Using Downscaling (FAQSD), we obtained a 17% reduction in MAE. We also found a stronger relationship between $PM_{2.5}$ and social vulnerability with XIS than with the FAQSD. Thus, XIS has potential for reconstructing environmental exposures, and its predictions have applications in environmental justice and human health.

1. Introduction

Particulate-matter air pollution comprises a mixture of solid and liquid particles that are suspended in the air. Concentrations of fine particulate matter ($PM_{2.5}$; having a diameter of less than 2.5 μm) are widely monitored and studied due to associations of short- and long-term exposure to $PM_{2.5}$ with disease (US Environmental Protection Agency, 2022a). Long-term ambient $PM_{2.5}$ was the leading environmental risk factor and ranked sixth among all modifiable risk factors in the 2019 Global Burden of Disease Study (Health Effects Institute, 2020), with additional health impacts attributable to short-term $PM_{2.5}$ exposure.

A primary difficulty for such health studies is the geographically sparse monitoring of $PM_{2.5}$, especially over large areas with complex emissions patterns, such as the contiguous United States (CONUS). Regulatory monitoring networks can provide high temporal resolution, with hourly or daily samples (US Environmental Protection Agency, 2022b), but sparse spatial coverage can lead to substantial measurement error on the exposure side of epidemiological analyses (Zeger et al., 2000). Hence, in place of merely assigning cases the $PM_{2.5}$ concentration measured at the nearest monitor, researchers increasingly use a variety of methods to model and interpolate $PM_{2.5}$, from chemical-transport models (van Donkelaar et al., 2010) to land-use regression (Hoek et al., 2008) to machine-learning approaches (reviewed in Diao et al. (2019)). Even then, however, $PM_{2.5}$ epidemiology has focused on urban areas, where monitoring is most intense (Hoek et al., 2013; Pope and Dockery, 2006). This is a critical limitation, since many people across CONUS live in rural or suburban areas, which have a different source profile of air pollution. Especially necessary are evaluation metrics that give rural areas appropriate weight, instead of mostly reflecting performance in intensely monitored areas (Carrión et al., 2021; Just et al., 2020a).

Machine-learning models for $PM_{2.5}$ typically use predictors of sources (e.g., roadways), topography, and meteorological conditions that

* Corresponding author. Box G-S121-2, Providence, RI, 02912, USA.
E-mail address: allan_just@brown.edu (A.C. Just).

https://doi.org/10.1016/j.envres.2025.120948
Received 4 October 2024; Received in revised form 20 December 2024; Accepted 23 January 2025
Available online 27 January 2025
0013-9351/

relate to $PM_{2.5}$ concentrations. The advent of remotely sensed Earth observations have led many to include satellite aerosol optical depth (AOD) into modeling efforts, including those from our group (Kloog et al., 2011). Since AOD is a quantitative estimate of the amount of light absorbed or scattered by suspended particles along the vertical atmospheric column, it is a useful proxy for surface $PM_{2.5}$ concentrations (Just et al., 2018). We have developed multiple AOD-based models at various geographical scales (state, region, and country) integrating satellite data with land use regression predictors using linear mixed models to calibrate the satellite-to-surface relationship (Hough et al., 2021; Just et al., 2015; Stafoggia et al., 2017).

There is a clear trend for recent $PM_{2.5}$ models to use methods from machine learning, including random forests (Hu et al., 2017; Shtein et al., 2019; Stafoggia et al., 2017), gradient boosting (Just et al., 2020a; Reid et al., 2015), neural networks (Di et al., 2016), and heterogeneous ensembles (Di et al., 2019). While these more flexible models can improve predictive accuracy, we (and others) have shown that without adequate care for the structure of the data, they are prone to overfitting, and data leakage (inadvertent inclusion of test data in model training) can optimistically bias assessments of model performance (Just et al., 2020a). Importantly, a flexible model will appear to be much more accurate when evaluated without consideration of the spatial structure of the underlying phenomenon, and this effect on apparent performance is evidence of overfitting (Just et al., 2020a). Another issue with machine learning is the temptation to build huge models that demand extensive computational resources, run slowly, and use hundreds of predictors that would all need to be updated for future years. Such models sound very impressive, but are difficult to scale, update, and reproduce, and are not necessarily more accurate than more conservative models.

We aimed to develop a new computationally efficient national model based on integrating a machine-learning method, namely extreme gradient boosting (XGBoost)(Chen and Guestrin, 2016; Just et al., 2020a), with inverse-distance weighting (IDW). Predictors included a parsimonious set of satellite, land-use, meteorological, and topographical variables. We call our general modeling framework, which can be used to model not only $PM_{2.5}$ but also other environmental variables such as air temperature, XGBoost-IDW Synthesis (XIS); the adaptation of XIS for $PM_{2.5}$ is named XIS-$PM_{2.5}$. Our model covers CONUS for each day of 2003 through 2023, with planned updates as new data become available. The strengths of XIS include an ungridded core that can make predictions at arbitrary points in the study region; good computational efficiency, with a geospatial data pipeline that dramatically cuts down on time and processing requirements; recency, with pre-planned ability to update the model for new years; a weighted evaluation scheme that fairly considers model performance across the entire region; and interpretability of results. We provide detailed evaluation of predictive performance in site-wise cross-validation, with stratification by year and climate region. We use interpretable machine-learning approaches to better understand the predictor-to-$PM_{2.5}$ relationships learned by our model. We also compare our model with the EPA FAQSD daily tract-level $PM_{2.5}$ predictions, highlighting the value of point-based predictions. An environmental-justice application shows that XIS, versus EPA FAQSD, produces substantially different estimates of exposure disparities across CONUS.

2. Method

2.1. Study area and time period

We modeled $PM_{2.5}$ for each year from 2003 to 2023. Our study area was CONUS excluding some large water bodies, as defined by the US Census's state-wise cartographic boundary files (US Census, 2019); it covers 7,798,188 km^2 in all. For some analyses, we show results divided into the 9 standard NOAA climate regions (Karl and Koscielny, 1982; NOAA, 2013), for which we treat Washington, DC, as part of the Northeast region.

Although XIS represents time as discrete days, in Central Standard Time (UTC−6), it does not discretize space into a grid. Rather, it represents locations as floating-point longitude–latitude pairs.

2.2. Data

2.2.1. Particulate matter

The outcome variable for XIS-$PM_{2.5}$, fine particulate matter mass concentration, is represented as $PM_{2.5}$ measurements recorded in the Environmental Protection Agency's Air Quality System (AQS) (US Environmental Protection Agency, 2022b). We included monitors using the federal reference method or federal equivalent methods (parameter code 88101) as well as other monitors reporting "acceptable $PM_{2.5}$" (parameter code 88502), including mass concentrations from speciation networks. Overall, 36% of the observations we used came from "acceptable" instruments, and 18% of monitoring sites only became available for this analysis because they include such an instrument. We filtered the outcome as follows.

1. Select observations inside the study area that are based on 24-h measures.
2. Drop observations with an event type of "Excluded"; use the corresponding "Included" observations instead.
3. Group observations into sites on the basis of longitude and latitude (disregarding AQS site identifiers).
4. Handle each observation of negative $PM_{2.5}$ by either setting it to 0 or discarding it entirely. A negative observation is discarded if it is more than 1 standard deviation (SD) away from the mean of other observations at the same site within the past 3 days and the next 3 days.
5. Reduce the data to at most one observation per site and day. Rank the observations as follows: prefer the designated primary monitor if there is one; prefer parameter code 88101 to 88502; prefer integrated 24-h measures over 24-h block averages; prefer lower AQSIDs; and prefer lower parameter occurrence codes.
6. Buddy-check the observations: for each site and day, compute the mean of the observations at all other sites within 30 km, weighted by the inverse square of the distance. If there is at least one such buddy and the interpolation differs from the interpolated-to observation by 20 $\mu g/m^3$ or more, flag the observation for removal. Once all observations have been checked, remove the flagged observations.

Across all years, the result comprised 6,192,276 station-days of observations, 13,669 of which (1 in 453) had the value 0 $\mu g/m^3$.

2.2.2. Predictors

XIS-$PM_{2.5}$ uses the following 22 variables as predictors. Each list item is headed with a bracketed pair showing the underlying [temporal resolution, spatial resolution] of the described predictors; floating-point variables are represented as 64-bit IEEE 754 floats. Further details of predictor computation are given below.

- [—, floating-point degree] Longitude and latitude
- [1 day, —] The integer day of the year
- [—, floating-point meter] An IDW feature, which is an interpolation of same-day particulate measurements at sites within 500 km, weighted by the cube of the distance (thus, the IDW exponent is 3)
- [1 day, 927 m] Daily AOD at 470 nm, from the MAIAC algorithm for Terra and Aqua (one variable per satellite) (Lyapustin, Alexei and Wang, Yujie, 2022)
- [1 h, 0.625° × 0.5°] Daily modeled surface $PM_{2.5}$ concentrations, from the Modern-Era Retrospective analysis for Research and Applications, Version 2 (MERRA-2) (Global Modeling And Assimilation Office, 2015)
- [1 month, 927 m] Monthly vegetation, quantified as the enhanced vegetation index from Aqua (Didan, 2021)

- [1 h, 0.125°] The daily height of the planetary boundary layer (PBL) from the fifth-generation reanalysis of the global climate dataset (ERA5) (Hersbach et al., 2018)
- [1 min, floating-point meter] The distance from the nearest fire on the same day, using fire locations from Aqua and Terra (Giglio et al., 2016)
- [—, floating-point meter] The distance from the nearest primary road, using the US Census's 2019 national road geodatabase (US Census Bureau, Geography Division, 2019)
- [irregular, 30 m] Two variables for surface imperviousness (from the National Land Cover Database (Dewitz, 2021)): one for the imperviousness at a single 30-m grid cell and one for the Gaussian-filtered imperviousness in a 1-km square around the query point
- [5 years, 30 arc-seconds] Population density, from the Gridded Population of the World (Center For International Earth Science Information Network-CIESIN-Columbia University, 2018)
- [—, 300 m] Elevation, from the US Geological Survey's 3D Elevation Program (US Geological Survey, 2017)
- [—, 300 m] Hilliness, or local relative topography, quantified as the multi-scale topographic dissection index computed from elevation (Oyler et al., 2015)
- [1 h, 0.125°] 6 meteorological variables from the North American Land Data Assimilation System-2 (NLDAS-2): temperature, specific humidity, air pressure, zonal wind speed, meridional wind speed, and precipitation (NASA, 2022)

The IDW interpolation for each case is computed as $(w_1{*}v_1 + w_2{*}v_2 + \cdots + w_n{*}v_n)/(w_1 + w_2 + \cdots + w_n)$, where the w_is are weights and the v_is are particulate concentrations.

The population-density product we used is available in intervals of 5 years, and the imperviousness product is available for irregularly spaced years. For each year of our data, we used the latest update of these products that was not in the future; for example, there is one Gridded Population of the World dataset for 2000 and another for 2005, so we used the 2000 dataset for our 2004 analyses and the 2005 dataset for our 2005 analyses.

ERA5 PBL height from the European Centre for Medium-Range Weather Forecasts (ECMWF) was downloaded from the Copernicus Climate Change Service (C3S) Climate Data Store (Hersbach et al., 2018). PBL height, MERRA-2 surface $PM_{2.5}$ concentrations (calculated using the formula from Global Modeling and Assimilation Office (2022)), and NLDAS-2 meteorological variables are available on an hourly basis, so we computed the mean for each day in Central Standard Time.

Elevation data for CONUS were aggregated from 1-arcsecond (~30 m) to 300-m resolution for computational speed. Hilliness was constructed from this aggregated raster as a multi-scale topographic dissection index (using window sizes of 3, 6, 9, 12, and 15 km) following the topoclimatic temperature modeling of Oyler et al. (2015).

Where multiple overpasses from the same MODIS instrument occurred over the same 1-km grid cell per day, we selected the non-missing AOD with the lowest associated theoretical AOD uncertainty (based on surface brightness) and restricted to "best quality" or "AOD within +-2 km from the coastline" or "Land, research quality" based on the field AOD_QA (Lyapustin et al., 2018).

2.3. Models

XIS uses a machine-learning algorithm called extreme gradient boosting (XGBoost) (Chen and Guestrin, 2016). XGBoost grows a forest of regression trees, fitting each tree to the error of prior trees, and applies several kinds of regularization, allowing it to strike a balance between flexibility, avoidance of overfitting, and computation time. XGBoost also automatically handles missing values: for each split, it chooses a default direction to use when the split variable is missing. Hence, XIS does not need to separately impute missing predictors.

Although the dependent variable for XIS-$PM_{2.5}$ is fine particulate concentration, we don't provide this directly to XGBoost. Instead, we compute the IDW interpolations first, and XGBoost models the observed concentration minus the IDW. For prediction, the IDW is added back. This strategy serves two purposes: it allows XIS to benefit from our prior knowledge that IDW is by far the most important feature, and it produces smoother predictions than raw XGBoost output, which comes in discrete steps because of the use of trees. The IDW is still provided to XGBoost as a feature in case XGBoost can benefit by altering the prediction based on its value. Lastly, all negative predictions (which are extremely rarely produced) are set to 0, because real particulate concentrations are nonnegative.

To tune XGBoost's hyperparameters, which control the learning process and model complexity, we first chose 50 hyperparameter vectors at random with a maximin Latin-hypercube algorithm (Carnell, 2020) to ensure broad coverage of the hyperparameter space. We conducted cross-validation (see below) on each of two years of our data, 2004 and 2018, with each of these hyperparameter vectors. We used only two years to ensure we had plenty of data to test the models on that we had not already used for tuning hyperparameters. On the basis of computation time and performance in both years and across regions, we chose these hyperparameter values: nrounds 500, max_depth 6, colsample_bytree 0.5, eta 0.073, gamma 0.0093, lambda 130, alpha 0.0012.

2.4. Evaluation

Predictive models are often evaluated in terms of squared error, leading to the SD as a metric of variability and root mean square error (RMSE) as a metric of performance. However, $PM_{2.5}$ observations are strongly positively skewed: they are all nonnegative and most values are small, while a minority are very large. For example, among the 350,274 AQS observations from 2019, even after the aforementioned filtering steps, the quartiles were 3.9, 6, 9 $\mu g/m^3$, and the 0.95 quantile was 15.4 $\mu g/m^3$, but the 50 greatest values ranged from 54 to 113 $\mu g/m^3$. Squared error would emphasize performance for this handful of large values, incentivizing models to create sufficiently large predictions to support large values while overestimating the majority of values. Thus, we quantified performance with mean absolute error (MAE), computed as mean(abs(observed - predicted)), and baseline variability with mean absolute deviation from the median (MAD), computed as mean(abs(observed - median(observed))). In the same way that the SD equals the RMSE that is obtained by predicting every value with the mean, the MAD equals the MAE that is obtained by predicting every value with the median. For XGBoost's objective function, we used log-cosh error, an everywhere twice-differentiable approximation to absolute error, defined by ln(cosh(observed - predicted)).

We estimated MAE in new data with ten-fold site-wise cross-validation, as follows. In each year, we randomly partitioned all sites with at least one observation into ten equally sized folds. We separately fit XIS on each set of nine folds and made predictions to the held-out fold. During cross-validation, we computed IDW interpolations while holding out sites from both the test fold and the fold of the interpolated-to site.

Were we to take the raw cross-validated MAE as our measure of model performance, a problem would arise due to the spatial distribution of sites. AQS sites are spread unevenly throughout the CONUS, with high concentrations of sites in some places and only a few sites in others. Hence, the unweighted MAE would emphasize performance wherever there happens to be more sites. We weighted observations (in evaluation, but not in training) so as to give equal weight to each unit of spacetime covered by XIS. The method was, for each day and region, to draw a Voronoi diagram (Turner, 2020) for all sites with an observation, and use the areas of the Voronoi tiles, in km^2, as weights for the observations; thus, observations that were relatively isolated were assigned greater weight.

2.5. New predictions

To make predictions for new point-days, we fit XIS to all the training data we had for each year. The process of making new predictions is highly analogous to predicting held-out observations in CV, and uses much of the same code. Longitude–latitude–date triples are taken as input; the longitude and latitude are 64-bit floating-point numbers of degrees, whereas the date has a resolution of 1 day, as usual. Prediction supports simultaneous use of XIS variations with different dependent variables, such as temperature (Just et al., 2024), to save time constructing predictors when multiple dependent variables are desired. The IDW logic differs from CV in that no stations need to be excluded on account of folds; for computational reasons, however, if a prediction point happens to be at exactly the same longitude–latitude pair as a station, we exclude that station.

3. Results

3.1. Cross-validation

Table 1 shows the weighted MAD, MAE, and bias (mean signed error) for each year of cross-validated predictions. Each MAE is one or more $\mu g/m^3$ lower than its corresponding MAD, showing that XIS has meaningful predictive ability. Averaging across all years, the mean MAD is 4.15 $\mu g/m^3$ and the mean MAE is 2.09 $\mu g/m^3$. The MAD tends to decrease over the years, and so does the MAE, albeit not as fast, so in later years, accuracy is better but the improvement over baseline is not as impressive. The consistent negative bias is likely due to the skewed distribution of observations, which has a minority of large values. As an example of how the CV looks without weighting, the unweighted results for 2010 are MAD 4.46 $\mu g/m^3$, MAE 1.94 $\mu g/m^3$, and bias -0.23 $\mu g/m^3$.

Fig. 1 shows the same weighted-MAD and weighted-MAE metrics, but computed separately for each NOAA region. Overall, MAEs are less variable between regions and over time than the MADs are. Averaging across all years, we have the lowest mean MAE in the South region (1.7 $\mu g/m^3$), and the greatest in the Northwest (2.8 $\mu g/m^3$) and West (3.17 $\mu g/m^3$), where there is also the greatest variation in $PM_{2.5}$.

Across all years, cross-validation produced 2740 predictions of 0 $\mu g/m^3$ (1 in 2260 out of all predictions).

See the Supporting Information for additional results including a stratification by season and the cross-validation performance at isolated sites.

Table 1
Results from each yearly cross-validation, in $\mu g/m^3$.

Year	Observations	Sites	MAD	MAE	Bias
2003	181,665	1309	5.20	2.28	−0.29
2004	197,176	1262	5.08	2.22	−0.16
2005	215,482	1326	5.49	2.28	−0.06
2006	217,865	1291	4.82	2.19	−0.13
2007	248,887	1282	5.14	2.30	−0.15
2008	256,187	1283	4.59	2.20	−0.08
2009	274,408	1309	4.13	2.09	−0.20
2010	288,843	1298	4.23	2.13	−0.17
2011	286,683	1255	4.31	2.25	−0.16
2012	296,204	1245	3.89	2.21	−0.32
2013	302,980	1248	3.93	2.13	−0.21
2014	308,937	1271	3.80	2.10	−0.28
2015	317,884	1288	3.76	2.05	−0.22
2016	321,957	1261	3.26	1.94	−0.24
2017	331,614	1273	3.72	2.08	−0.25
2018	342,149	1283	3.63	1.95	−0.27
2019	350,274	1281	3.11	1.70	−0.19
2020	352,516	1261	3.84	2.01	−0.27
2021	362,757	1272	4.10	2.06	−0.32
2022	365,647	1260	3.24	1.81	−0.23
2023	372,161	1267	3.90	1.93	−0.24

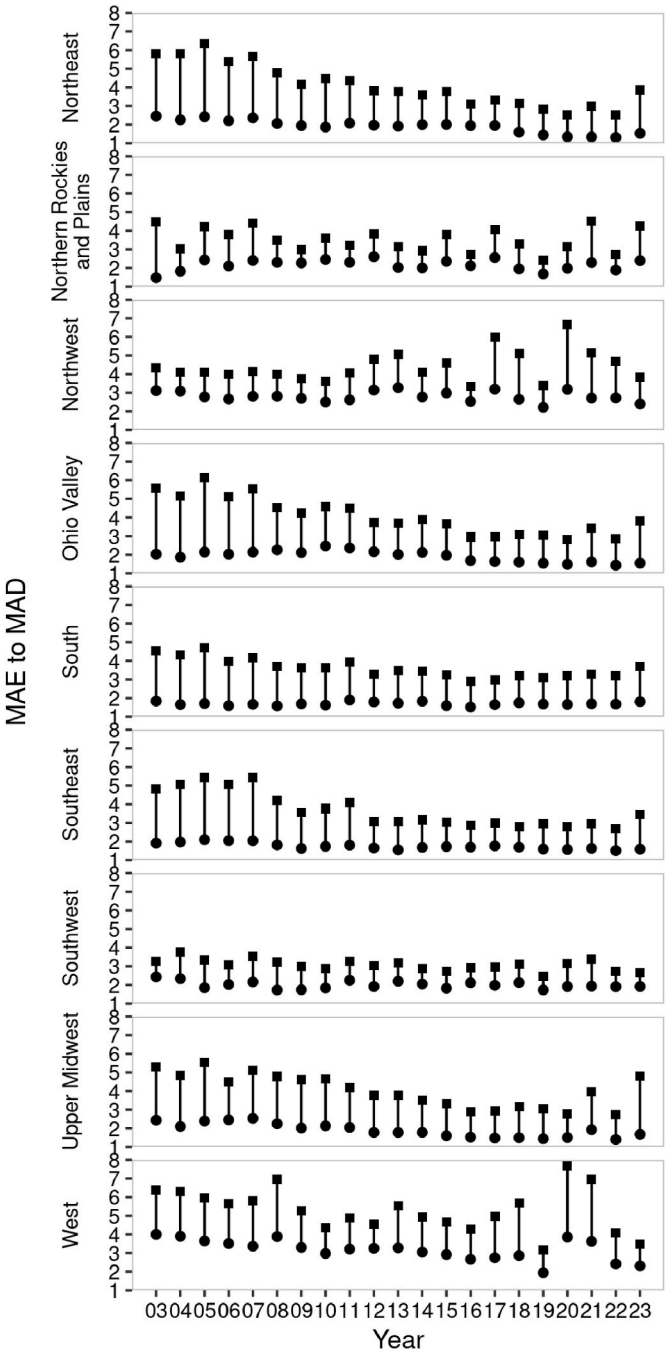

Fig. 1. Weighted MAD (squares) and MAE (circles) from cross-validation for each region and year.

3.1.1. Feature contributions

To examine how individual predictors relate to XIS's predictions, we show SHapley Additive exPlanations (SHAPs) (Lundberg et al., 2019). SHAPs can be interpreted analogously to the terms of a linear-regression model: a SHAP of +2.5 for a given predictor and case means that the model attributes a +2.5 increase in its prediction for that case to that predictor. We generated SHAPs for predictions corresponding to each observation while it was held out in cross-validation.

Fig. 2 shows per-feature mean absolute SHAPs for one year, which can be interpreted as showing the typical impact of the feature on the prediction. Except for the IDW feature, which has a mean absolute SHAP of 9.3 $\mu g/m^3$ (when computed accounting for its initial inclusion outside of XGBoost), each feature has a relatively small contribution, effectively

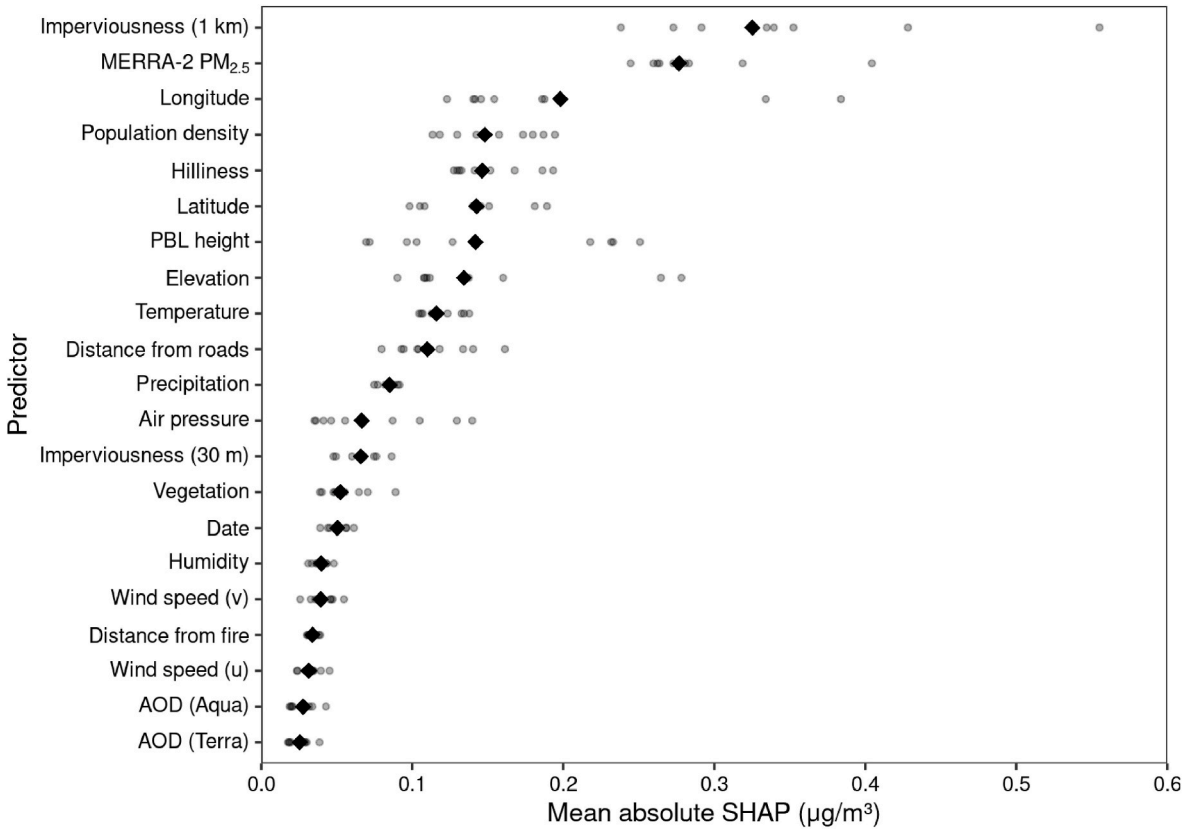

Fig. 2. Mean absolute SHAP of each predictor in 2010 (the IDW feature, which has much greater absolute SHAP than everything else, is omitted). Small dots show per-region means. Diamonds show overall means.

constituting fine adjustments to the IDW interpolation. About half of the features have a mean absolute SHAP greater than 0.1 μg/m^3.

Fig. 3 plots the mean SHAP of the road-distance feature for each AQS site. As one would expect, low distances are associated with a higher concentration of $PM_{2.5}$. SHAPs shrink towards 0 as the distance increases, indicating that our model finds the distance from the nearest road less predictively useful as it increases. See the Supporting Information for analyses demonstrating the relation of hilliness with SHAPs and the contribution of MERRA-2 as a function of site isolation.

3.2. *Comparison with a downscaler*

The EPA provides a $PM_{2.5}$ product with one prediction per day and US Census tract called the Fused Air Quality Surface Using Downscaling

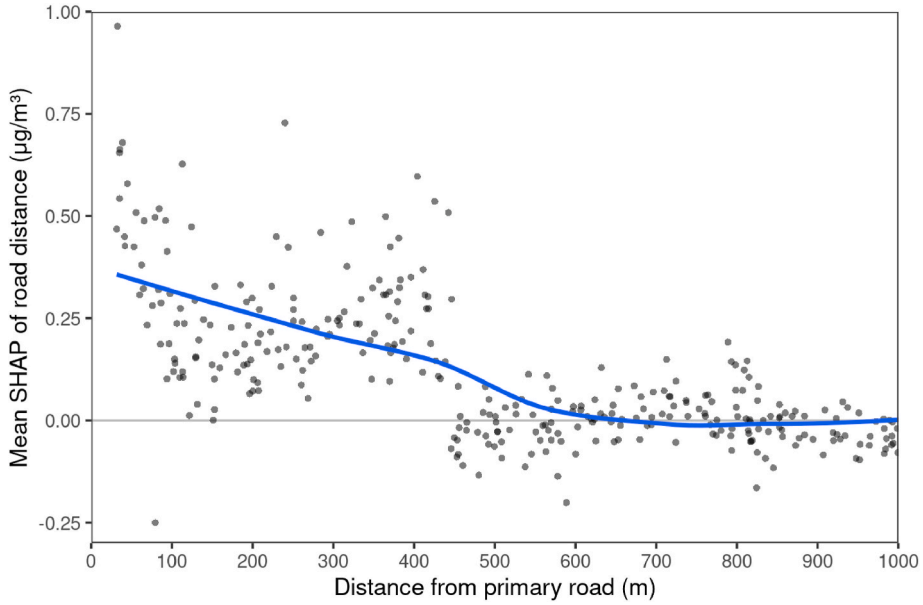

Fig. 3. SHAP of road distance as a function of road distance in 2010, showing only cases within 1 km of a primary road. A trendline is fit with locally estimated scatterplot smoothing (LOESS).

(FAQSD), which is currently available through 2021 (Adam, 2021; US EPA, 2015). The FAQSD uses a subset of AQS sites for a bias adjustment of surfaces generated from a 12-km resolution Community Multiscale Air Quality (CMAQ) model. We sought a direct comparison of XIS performance to the tract-level output of the FAQSD, taking advantage of the clear delineation of which AQS sites were and were not used by FAQSD (US EPA, 2015).

For each year, we denoted all observations from sites that the FAQSD used for training (identified by AQS ID in the FAQSD input files) as the training set, and all remaining observations as the test set. To get an FAQSD prediction for each test observation, we used the nearest FAQSD value (based on tract centroids) on the same day. We retrained XIS yearly using only AQS data from the training set and made two sets of predictions for the test set: one at the same locations as the FAQSD values, and one at the true locations of the test observations. By looking at XIS at both the FAQSD locations (nearest tract centroid) and the true locations of test monitors, we can see how much of the improvement of XIS over FAQSD is due to the improvement in spatial resolution from using point-based predictions.

Table 2 shows for each year the weighted MAD among the test set and the weighted MAEs of the three kinds of predictions. XIS achieved substantially lower MAE than the FAQSD, at least in earlier years, despite how the FAQSD uses the CMAQ and XIS doesn't. There is substantial improvement when the model is changed from FAQSD to XIS without changing prediction locations, and another, generally smaller improvement when XIS makes predictions for the true locations. Averaging across years, XIS achieves a 14% reduction in MAE compared to FAQSD without changing prediction locations, and a 17% reduction for true locations. Similarly, also averaging across years, XIS achieves a 45% reduction in absolute yearly bias compared to FAQSD without changing prediction locations, and a 51% reduction for true locations.

3.3. New predictions

With XIS trained on all the data for each year, we drew maps of predictions. Fig. 4 shows the predicted $PM_{2.5}$ concentrations throughout the study region for one year. Fig. 5 shows one day in the greater New York City area; here, we plot approximately one prediction per 187 m.

3.3.1. Social vulnerability

We examined how predicted $PM_{2.5}$ at the centroids of 82,390 Census tracts in CONUS, averaged across all days of 2020, related to the CDC's per-tract Social Vulnerability Index (SVI) (Centers for Disease Control and Prevention, 2018). SVI scores range from 0 (least vulnerable) to 1 (most vulnerable). We fit two linear mixed models, one with $PM_{2.5}$ estimates from XIS as the outcome and one with estimates from FAQSD. The models had a fixed effect for vulnerability as well as per-county random intercepts and slopes of vulnerability (modeled as correlated), so the model formula (in the syntax of the R package lme4) was pm ~ vuln + (vuln | county); note that this implicitly includes a fixed intercept. The fixed effect of vulnerability was estimated as 0.063 μg/m^3 (95% CI [0.038, 0.089]) for FAQSD and 0.535 μg/m^3 (95% CI [0.493, 0.577]) for XIS.

4. Discussion

We modeled daily $PM_{2.5}$ across CONUS from 2003 through 2023. Our model, XIS, uses a streamlined geospatial processing and machine-learning pipeline to facilitate regular updates with new data. XIS is trained on a broad information base of ground-level concentrations from the AQS, comprising monitors using the federal reference method (or an equivalent) as well as other acceptable $PM_{2.5}$ monitors. While there are many potential uses for this exposure model, particularly in epidemiology and environmental justice, we focus our discussion on metrics of predictive accuracy, comparison with a leading EPA product (the FAQSD), and the interpretation of individual predictors.

We took special care in our evaluation of predictive accuracy. Our site-wise cross-validation scheme, in which we split data spatially before computing the IDW predictor or fitting XGBoost, can estimate accuracy without being optimistically biased by data leakage or overfitting. Thus our performance metrics were evaluated at sites for which the model has seen no prior ground truth, indicating what the performance would be when, for example, estimating exposure at a person's home where there isn't already a monitor. We used MAE as our primary metric, to suitably reflect overall accuracy in the presence of a small number of extreme $PM_{2.5}$ concentrations, and we used a log-cosh objective function to help minimize MAE. Given the goal of covering sparsely monitored suburban and rural regions, we weighted the MAE to reflect the land area covered rather than the number of monitors. The result was overall good accuracy, substantially better than the MAD. Accuracy increased, and the gap between the MAD and MAE shrank, in later years as ambient $PM_{2.5}$ decreased. We also examined unweighted MAE at isolated monitors to further assess XIS's performance in sparsely monitored regions. Here we again found good accuracy, comparable to that of our overall weighted MAEs.

Our workflow is reproducible and computationally tractable. To provide some performance examples, CV (including making IDW interpolations, training XGBoost, and making out-of-fold predictions) for

Table 2
Results of the comparison between XIS and the FAQSD, in μg/m^3. The observation and site counts include only the test set. FAQSD is only available through 2021.

Year	Observations	Sites	MAD	FAQSD		XIS, tract centroids		XIS, true locations	
				MAE	Bias	MAE	Bias	MAE	Bias
2003	28,342	208	5.10	3.14	0.78	2.53	0.09	2.45	0.20
2004	38,863	238	4.68	3.14	1.14	2.51	0.22	2.42	0.43
2005	54,221	283	4.70	3.38	1.60	2.44	0.44	2.31	0.62
2006	63,833	315	4.54	3.01	1.18	2.44	0.58	2.31	0.63
2007	76,641	338	4.70	2.97	0.72	2.52	0.32	2.44	0.36
2008	83,365	360	4.21	2.95	0.93	2.39	0.34	2.32	0.43
2009	95,946	405	3.88	3.03	1.07	2.40	0.29	2.39	0.43
2010	107,284	433	3.97	2.77	0.66	2.35	0.22	2.32	0.38
2011	121,352	486	4.16	2.88	0.28	2.45	−0.20	2.42	0.05
2012	129,135	496	3.79	2.68	0.17	2.41	−0.38	2.37	−0.19
2013	132,386	504	3.83	2.70	0.13	2.37	−0.02	2.37	0.15
2014	80,717	341	3.72	2.61	0.74	2.30	−0.15	2.24	0.05
2015	78,731	342	3.66	2.58	0.99	2.12	0.01	2.10	0.22
2016	76,456	319	3.19	2.38	0.65	2.12	−0.11	2.02	0.08
2017	76,276	316	4.06	2.78	0.17	2.51	−0.43	2.48	−0.18
2018	76,403	310	3.77	2.42	0.33	2.23	−0.55	2.16	−0.34
2019	77,944	317	3.09	2.40	0.90	2.21	0.02	2.11	0.20
2020	78,752	318	3.98	2.74	1.00	2.50	0.12	2.46	0.25
2021	76,596	305	4.25	3.03	0.88	2.84	0.18	2.80	0.30

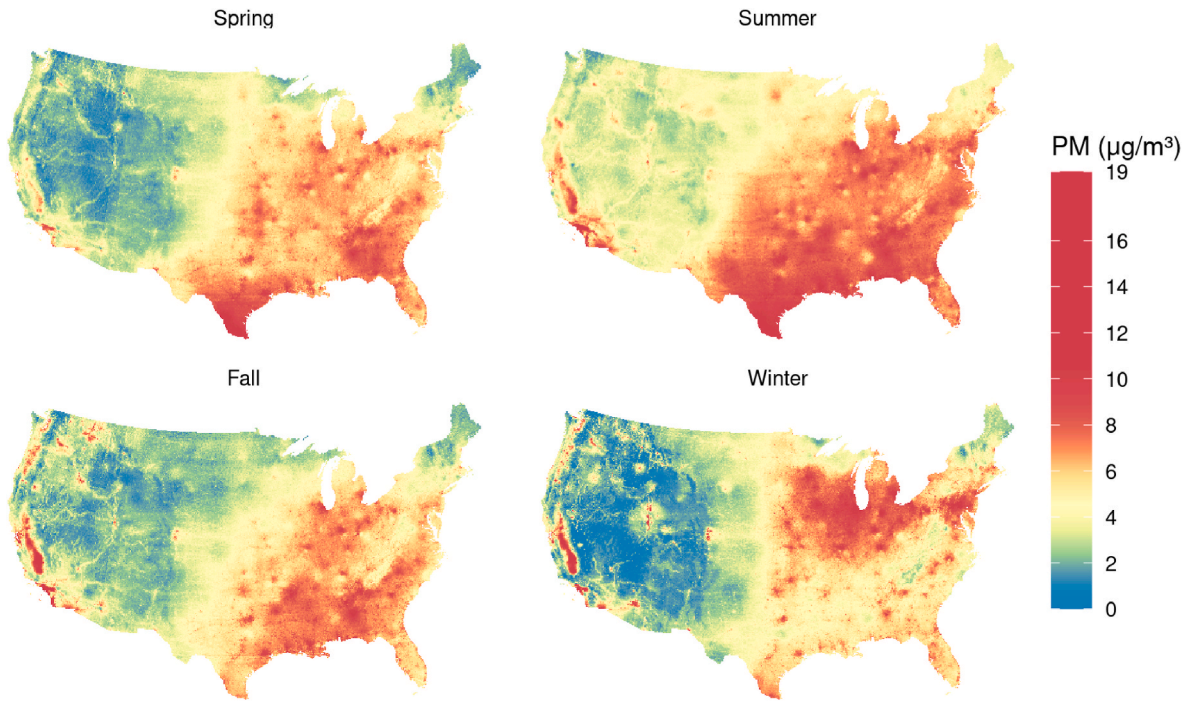

Fig. 4. Mean predicted $PM_{2.5}$ across the study area from Dec 1, 2018 through Nov 30, 2019, grouped by season. We use December from 2018 instead of 2019 so as to plot a contiguous winter. Colors are scaled according to quantiles, such that color changes more rapidly around values that are more common in the map. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

2023 (comprising 372,161 observations) took 0 h 8 min 44 s, while new predictions for the vulnerability analysis (comprising every day in 2020 at each of 83,163 tracts, for a total of 30,437,658 predictions) took 2 h 23 min 50 s. All computation took place on a single 24-core Xeon server built in 2018.

XIS's efficiency enables the incorporation of the most up-to-date data, including the EPA's 2024 extensive retroactive modification of $PM_{2.5}$ observations back through 2017 ("Update of PM2.5 Data from T640/T640X PM Mass Monitors," 2024). Efficiency also allows us to rerun XIS in order to interrogate performance. For example, we trained a version of XIS using only the AQS sites that the FAQSD is trained on, allowing us to fairly and comprehensively compare XIS to the FAQSD. This comparison demonstrated that XIS had a substantially lower MAE in every available year, even when fit without a substantial subset of ground monitors used in our full model. Had XIS required hundreds of predictors and thousands of processor cores, this kind of rerun would have been impractical. Rather, we chose our predictors parsimoniously, and we searched for hyperparameters in a way that was more efficient than exhaustive.

In XIS, as in our prior XGBoost-based models (Just et al., 2020a; Just et al., 2020b), we used SHAP to quantify the contribution of individual predictors. SHAPs show that the IDW interpolation of monitor observations is the greatest contributor to predictions at withheld monitors. The IDW interpolation serves as a base prediction (modified by XGBoost) and reuses distance matrices for speed. As an example of XIS's interpretability using SHAP, Fig. 3 shows that the positive effect of being near a primary road smoothly decreases with distance, particularly within 500 m. This figure also shows the value of XIS's point-based design, incorporating both raster data and continuous fields, in contrast to our previous 1-km gridded models (Just et al., 2020a). Such spatial precision was further justified in the comparison with the FAQSD when we found that using the exact locations of test monitors, compared to tract centroids, increased accuracy and decreased bias (Table 2).

For an environmental-justice application with relevance to human health, we examined the association between social vulnerability and annual $PM_{2.5}$ concentrations at Census tracts throughout CONUS. We see a striking difference between models in estimates of exposure disparities: XIS's estimate of the relationship between vulnerability and $PM_{2.5}$ is nearly an order of magnitude greater than the FAQSD's. Thus XIS reveals meaningful exposure disparities that would be greatly attenuated by using the FAQSD.

XIS starts in 2003, when MODIS instruments on two satellites (Aqua and Terra) became available, and is updated through 2023. Pairing this long duration (21 years) with such recency is important for health studies in a rapidly changing era. Few other national $PM_{2.5}$ models for use in health studies are regularly updated. The FAQSD is a notable example, although its dependence on the National Emissions Inventory, which is updated every three years, means that the FAQSD is typically several years out of date. We have developed our geospatial processing workflow to ingest new data (including new types of predictors) and make periodic updates so that we can continue to generate timely exposure estimates. Future developments may incorporate near-real time data streams such as EPA AirNOW and the near-real-time MAIAC AOD (MODIS Land Science Team, 2022), allowing us to generate predictions up to the past few days and model rapidly evolving air-quality and health conditions. Similarly, we have developed XIS to be adapted for other environmental exposures. Ongoing efforts are generating a suite of complementary exposure estimates (e.g., air temperature and humidity) that are synergistically useful with these $PM_{2.5}$ estimates for health studies.

A limitation of XIS-$PM_{2.5}$ is that it only models mass concentrations. The toxicity of particulates is also related to particulate composition, which XIS does not address. Furthermore, since XIS uses the AQS as ground truth, our ability to update XIS is limited by the AQS release schedule. For example, EPA first released complete 2023 AQS summary files on August 28, 2024. Finally, it remains possible that adding some of the complexity we have avoided, such as hundreds of additional predictors or elaborate ensembles of models, would increase XIS's accuracy. Diminishing returns, however, means that a great deal of XIS's agility and tractability could be lost for a marginal improvement in accuracy.

In summary, we present a new exposure model, XIS, which generates point-based daily $PM_{2.5}$ exposures across the conterminous United

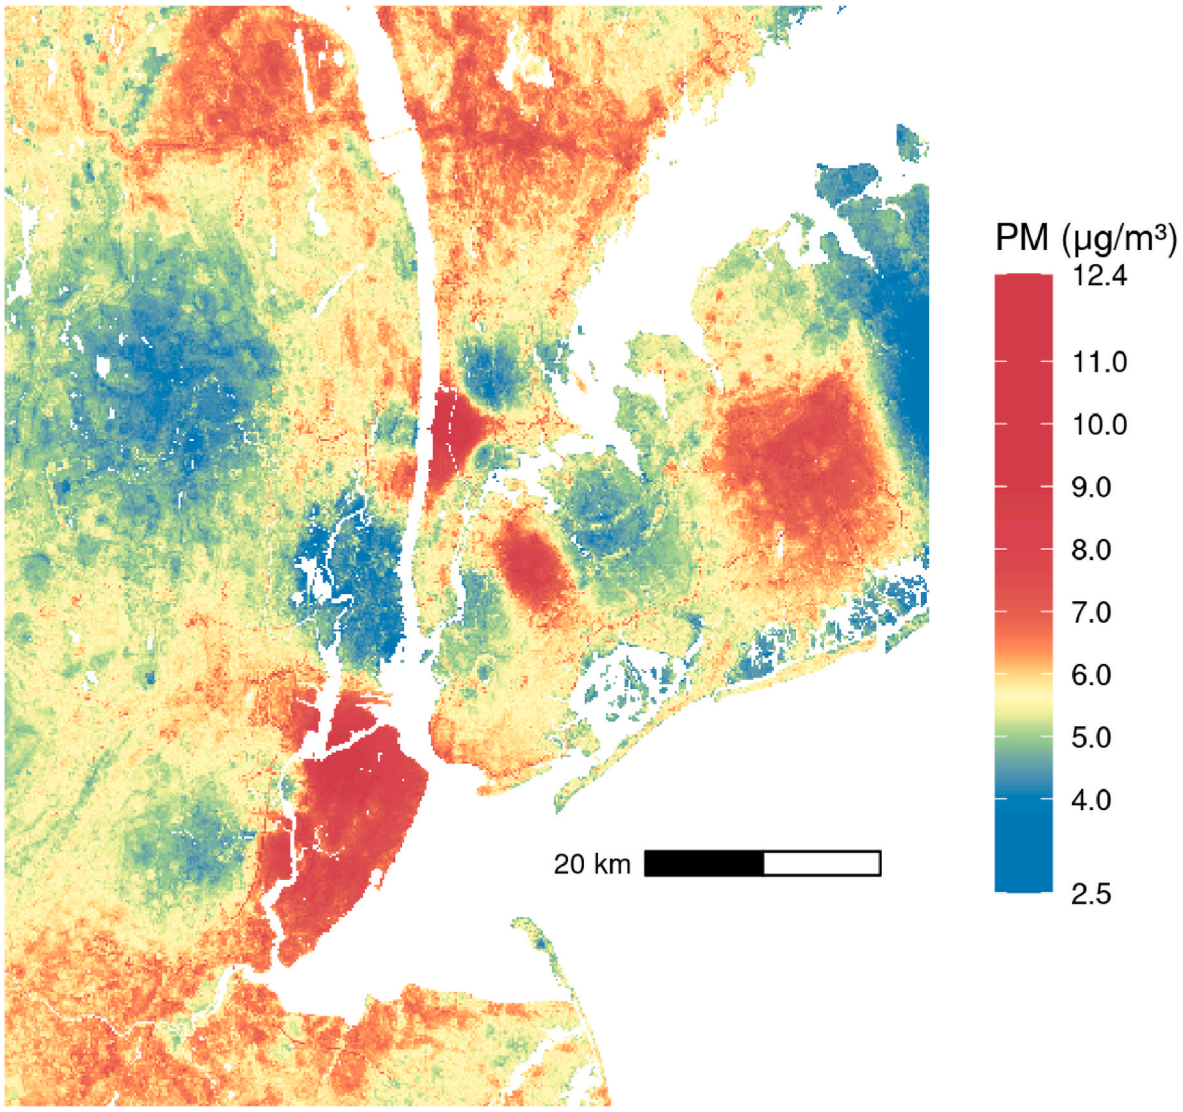

Fig. 5. Predicted $PM_{2.5}$ on Jan 24, 2023 in the New York City area. We chose this date by computing the mean $PM_{2.5}$ at all stations in this area for each day in 2023, then taking the median day.

States 2003–2023. This model is intended to generate ambient exposure estimates in cohort and registry-based epidemiological and exposure disparity studies in order to advance the evidence basis of the health impacts of chronic and acute exposure to $PM_{2.5}$.

CRediT authorship contribution statement

Allan C. Just: Writing – review & editing, Writing – original draft, Supervision, Project administration, Funding acquisition, Conceptualization. **Kodi B. Arfer:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation. **Johnathan Rush:** Writing – review & editing, Software. **Alexei Lyapustin:** Writing – review & editing. **Itai Kloog:** Writing – review & editing.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Allan Just reports financial support was provided by National Institutes of Health. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Research reported in this publication was supported by the Environmental influences on Child Health Outcomes (ECHO) program, Office of The Director, National Institutes of Health, under Award Numbers U2C OD023375, U24 OD023382, U24 OD023319, UH3 OD023337, and an ECHO Opportunities and Infrastructure Fund award to ACJ, as well as National Institutes of Health grants R01 ES031295, R01 DK127139, P20 AG089308, P30 ES023515, and UL1 TR004419.

The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.envres.2025.120948.

Data availability

Intermediate data files to reproduce our analyses are openly available on Zenodo at https://zenodo.org/doi/10.5281/zenodo.7331250. These data were derived from a combination of public-domain and restricted-access sources described in our Method.

References

Adam, Reff, 2021. Bayesian Space-time Downscaling Fusion Model (Downscaler) - Derived Estimates of Air Quality for 2018. EPA-454/R-21-003 135.

Carnell, R., 2020. maximinLHS (lhs: Latin Hypercube Samples) [WWW Document]. URL. https://CRAN.R-project.org/package=lhs. (Accessed 10 July 2020).

Carrión, D., Arfer, K.B., Rush, J., Dorman, M., Rowland, S.T., Kioumourtzoglou, M.-A., Kloog, I., Just, A.C., 2021. A 1-km hourly air-temperature model for 13 northeastern U.S. states using remotely sensed and ground-based measurements. Environ. Res. 200, 111477. https://doi.org/10.1016/j.envres.2021.111477.

Center For International Earth Science Information Network-CIESIN-Columbia University, 2018. Gridded population of the World, version 4 (GPWv4): population count, revision 11. Gridded population of the World. Version 4 (GPWv4). https://doi.org/10.7927/H4JW8BX5.

Centers for Disease Control and Prevention/Agency for Toxic Substances and Disease Registry/Geospatial Research, Analysis, and Services Program, 2018. CDC/ATSDR social vulnerability index (SVI) 2018 Database US [WWW Document]. URL. https://www.atsdr.cdc.gov/placeandhealth/svi/index.html. (Accessed 17 June 2022).

Chen, T., Guestrin, C., 2016. XGBoost: a scalable tree boosting system. https://doi.org/10.1145/2939672.2939785.

Dewitz, J., 2021. National land cover Database (NLCD) 2019 products. https://doi.org/10.5066/P9KZCM54.

Di, Q., Amini, H., Shi, L., Kloog, I., Silvern, R., Kelly, J., Sabath, M.B., Choirat, C., Koutrakis, P., Lyapustin, A., Wang, Y., Mickley, L.J., Schwartz, J., 2019. An ensemble-based model of PM2.5 concentration across the contiguous United States with high spatiotemporal resolution. Environ. Int. 130, 104909. https://doi.org/10.1016/j.envint.2019.104909.

Di, Q., Koutrakis, P., Schwartz, J., 2016. A hybrid prediction model for PM2.5 mass and components using a chemical transport model and land use regression. Atmos. Environ. 131, 390–399. https://doi.org/10.1016/j.atmosenv.2016.02.002.

Diao, M., Holloway, T., Choi, S., O'Neill, S.M., Al-Hamdan, M.Z., Van Donkelaar, A., Martin, R.V., Jin, X., Fiore, A.M., Henze, D.K., Lacey, F., Kinney, P.L., Freedman, F., Larkin, N.K., Zou, Y., Kelly, J.T., Vaidyanathan, A., 2019. Methods, availability, and applications of PM $_{2.5}$ exposure estimates derived from ground measurements, satellite, and atmospheric models. J. Air Waste Manag. Assoc. 69, 1391–1414. https://doi.org/10.1080/10962247.2019.1668498.

Didan, Kamel, 2021. MODIS/Aqua vegetation indices monthly L3 global 1km SIN grid V061. https://doi.org/10.5067/MODIS/MYD13A3.061.

Giglio, L., Schroeder, W., Justice, C.O., 2016. The collection 6 MODIS active fire detection algorithm and fire products. Remote Sens. Environ. 178, 31–41.

Global Modeling and Assimilation Office, 2022. MERRA-2 FAQ [WWW Document]. URL. https://web.archive.org/web/2021/. (Accessed 6 September 2022). https://gmao.gsfc.nasa.gov/reanalysis/MERRA-2/FAQ.

Global Modeling And Assimilation Office, 2015. MERRA-2 tavg1_2d_aer_Nx: 2d,1-Hourly,Time-averaged,Single-Level,Assimilation,Aerosol diagnostics V5.12.4. MERRA-2 tavg1_2d_aer_Nx. https://doi.org/10.5067/KLICLTZ8EM9D.

Health Effects Institute, 2020. State of Global Air 2020.

Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I., 2018. ERA5 hourly data on single levels from 1979 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS) 10. https://doi.org/10.24381/cds.adbb2d47.

Hoek, G., Beelen, R., De Hoogh, K., Vienneau, D., Gulliver, J., Fischer, P., Briggs, D., 2008. A review of land-use regression models to assess spatial variation of outdoor air pollution. Atmos. Environ. 42, 7561–7578.

Hoek, G., Krishnan, R.M., Beelen, R., Peters, A., Ostro, B., Brunekreef, B., Kaufman, J.D., 2013. Long-term air pollution exposure and cardio- respiratory mortality: a review. Environ. Health 12, 43. https://doi.org/10.1186/1476-069X-12-43.

Hough, I., Sarafian, R., Shtein, A., Zhou, B., Lepeule, J., Kloog, I., 2021. Gaussian Markov random fields improve ensemble predictions of daily 1 km PM2. 5 and PM10 across France. Atmos. Environ. 264, 118693.

Hu, X., Belle, J.H., Meng, X., Wildani, A., Waller, L.A., Strickland, M.J., Liu, Y., 2017. Estimating PM $_{2.5}$ concentrations in the conterminous United States using the random forest approach. Environ. Sci. Technol. 51, 6936–6944. https://doi.org/10.1021/acs.est.7b01210.

Just, A.C., Arfer, K.B., Rush, J., Dorman, M., Shtein, A., Lyapustin, A., Kloog, I., 2020a. Advancing methodologies for applying machine learning and evaluating spatiotemporal models of fine particulate matter (PM2.5) using satellite data over large regions. Atmos. Environ. 239, 117649. https://doi.org/10.1016/j.atmosenv.2020.117649.

Just, A.C., Arfer, K.B., Rush, J., Kloog, I., 2024. XIS-Temperature: a daily spatiotemporal machine-learning model for air temperature in the contiguous United States. Earth Space Sci. Open Arch. https://doi.org/10.22541/essoar.167591086.68441300/v2.

Just, A.C., De Carli, M.M., Shtein, A., Dorman, M., Lyapustin, A., Kloog, I., 2018. Correcting measurement error in satellite aerosol optical depth with machine learning for modeling PM2.5 in the northeastern USA. Rem. Sens. 10, 803. https://doi.org/10.3390/rs10050803.

Just, A.C., Liu, Y., Sorek-Hamer, M., Rush, J., Dorman, M., Chatfield, R., Wang, Y., Lyapustin, A., Kloog, I., 2020b. Gradient boosting machine learning to improve satellite-derived column water vapor measurement error. Atmos. Meas. Tech. 13, 4669–4681.

Just, A.C., Wright, R.O., Schwartz, J., Coull, B.A., Baccarelli, A.A., Tellez-Rojo, M.M., Moody, E., Wang, Y., Lyapustin, A., Kloog, I., 2015. Using high-resolution satellite aerosol optical depth to estimate daily PM $_{2.5}$ geographical distribution in Mexico city. Environ. Sci. Technol. 49, 8576–8584. https://doi.org/10.1021/acs.est.5b00859.

Karl, T.R., Koscielny, A.J., 1982. Drought in the United States: 1895–1981. J. Climatol. 2, 313–329. https://doi.org/10.1002/joc.3370020402.

Kloog, I., Koutrakis, P., Coull, B.A., Lee, H.J., Schwartz, J., 2011. Assessing temporally and spatially resolved PM2.5 exposures for epidemiological studies using satellite aerosol optical depth measurements. Atmos. Environ. 45, 6267–6275. https://doi.org/10.1016/j.atmosenv.2011.08.066.

Lundberg, S.M., Erion, G.G., Lee, S.-I., 2019. Consistent individualized feature attribution for tree ensembles.

Lyapustin, Alexei, Wang, Yujie, 2022. MCD19A2 MODIS/Terra+Aqua land aerosol optical depth daily L2G global 1km SIN grid V061. https://doi.org/10.5067/MODIS/MCD19A2.061.

Lyapustin, A., Wang, Y., Korkin, S., Huang, D., 2018. MODIS Collection 6 MAIAC algorithm. Atmos. Meas. Tech. 11, 5741–5765. https://doi.org/10.5194/amt-11-5741-2018.

MODIS Land Science Team, 2022. MODIS/Terra+Aqua land aerosol optical depth daily L2G global 1km SIN grid. https://doi.org/10.5067/MODIS/MCD19A2N.NRT.061.

NASA, 2022. NLDAS-2 forcing dataset information [WWW Document]. URL. https://ldas.gsfc.nasa.gov/nldas/v2/forcing. (Accessed 23 June 2022).

NOAA, 2013. U.S. Climate Regions | Monitoring References. National Centers for Environmental Information (NCEI) [WWW Document]. URL. https://www.ncdc.noaa.gov/monitoring-references/maps/us-climate-regions.php. (Accessed 10 July 2020).

Oyler, J.W., Ballantyne, A., Jencso, K., Sweet, M., Running, S.W., 2015. Creating a topoclimatic daily air temperature dataset for the conterminous United States using homogenized station data and remotely sensed land skin temperature. Int. J. Climatol. 35, 2258–2279. https://doi.org/10.1002/joc.4127.

Pope, C.A., Dockery, D.W., 2006. Health effects of fine particulate air pollution: lines that connect. J. Air Waste Manag. Assoc. 56, 709–742. https://doi.org/10.1080/10473289.2006.10464485.

Reid, C.E., Jerrett, M., Petersen, M.L., Pfister, G.G., Morefield, P.E., Tager, I.B., Raffuse, S.M., Balmes, J.R., 2015. Spatiotemporal prediction of fine particulate matter during the 2008 Northern California wildfires using machine learning. Environ. Sci. Technol. 49, 3887–3896. https://doi.org/10.1021/es505846r.

Shtein, A., Kloog, I., Schwartz, J., Silibello, C., Michelozzi, P., Gariazzo, C., Viegi, G., Forastiere, F., Karnieli, A., Just, A.C., 2019. Estimating daily PM2. 5 and PM10 over Italy using an ensemble model. Environ. Sci. Technol. 54, 120–128.

Stafoggia, M., Schwartz, J., Badaloni, C., Bellander, T., Alessandrini, E., Cattani, G., de'Donato, F., Gaeta, A., Leone, G., Lyapustin, A., Sorek-Hamer, M., de Hoogh, K., Di, Q., Forastiere, F., Kloog, I., 2017. Estimation of daily PM10 concentrations in Italy (2006–2012) using finely resolved satellite data, land use variables and meteorology. Environ. Int. 99, 234–244. https://doi.org/10.1016/j.envint.2016.11.024.

Turner, R., 2020. deldir: delaunay triangulation and Dirichlet (Voronoi) tessellation [WWW Document]. URL. https://CRAN.R-project.org/package=deldir. (Accessed 1 April 2021).

Update of PM2.5 Data from T640/T640X PM Mass Monitors, 2024. WWW Document. https://www.regulations.gov/document/EPA-HQ-OAR-2023-0642-0031.

US Census, 2019. Cartographic boundary files [WWW Document]. URL. https://www2.census.gov/geo/tiger/GENZ2019/description.pdf. (Accessed 10 July 2020).

US Census Bureau, 2019. Geography division, 2019. TIGER/line shapefiles [WWW Document]. URL. https://www.census.gov/cgi-bin/geo/shapefiles/index.php?year=2019&layergroup=Roads. (Accessed 17 June 2022).

US Environmental Protection Agency, 2022a. Supplement to the 2019 Integrated Science Assessment for Particulate Matter, 2022. Final Report [WWW Document]. URL. https://cfpub.epa.gov/ncea/isa/recordisplay.cfm?deid=354490. (Accessed 23 June 2022).

US Environmental Protection Agency, 2022b. EPA AirData [WWW Document]. URL. https://aqs.epa.gov/aqsweb/airdata/download_files.html. (Accessed 23 June 2022).

US EPA, O, 2015. RSIG-related downloadable data files [WWW Document]. URL. https://www.epa.gov/hesc/rsig-related-downloadable-data-files. (Accessed 6 July 2022).

US Geological Survey, 2017. 1 arc-second digital elevation models (DEMs) - USGS national map 3DEP downloadable data collection [WWW Document]. URL. https://www.sciencebase.gov/catalog/item/4f70aa71e4b058caae3f8de1. (Accessed 15 June 2022).

van Donkelaar, A., Martin, R.V., Brauer, M., Kahn, R., Levy, R., Verduzco, C., Villeneuve, P.J., 2010. Global estimates of ambient fine particulate matter concentrations from satellite-based aerosol optical depth: development and application. Environ. Health Perspect. 118, 847–855. https://doi.org/10.1289/ehp.0901623.

Zeger, S.L., Thomas, D., Dominici, F., Samet, J.M., Schwartz, J., Dockery, D., Cohen, A., 2000. Exposure measurement error in time-series studies of air pollution: concepts and consequences. Environ. Health Perspect. 108, 419–426. https://doi.org/10.1289/ehp.00108419.