ENVIRONMENTAL STUDIES

Improved daily $PM_{2.5}$ estimates in India reveal inequalities in recent enhancement of air quality

Ayako Kawano[1]*, Makoto Kelp[2], Minghao Qiu[3,4], Kirat Singh[1], Eeshan Chaturvedi[1], Sunil Dahiya[5], Inés Azevedo[6,7,8,9], Marshall Burke[10,11,12]

Poor ambient air quality poses a substantial global health threat. However, accurate measurement remains challenging, particularly in countries such as India where ground monitors are scarce despite high expected exposure and health burdens. This lack of precise measurements impedes understanding of changes in pollution exposure over time and across populations. Here, we develop open-source daily fine particulate matter ($PM_{2.5}$) datasets at a 10-kilometer resolution for India from 2005 to 2023 using a two-stage machine learning model validated on held-out monitor data. Analyzing long-term air quality trends, we find that $PM_{2.5}$ concentrations increased across most of the country until around 2016 and then declined partly due to favorable meteorology in southern India. Recent reductions in $PM_{2.5}$ were substantially larger in wealthier areas, highlighting the urgency of air quality control policies addressing all socioeconomic communities. To advance equitable air quality monitoring, we propose additional monitor locations in India and examine the adaptability of our method to other countries with scarce monitoring data.

INTRODUCTION

Exposure to ambient fine particulate matter ($PM_{2.5}$) is a recognized global health concern. India, with its large population and high average pollution levels, bears a substantial share of the global health burden from poor air quality (*1*). The health burden of air pollution often varies across individuals and groups with different socio-economic status due to differences in pollution concentrations and the increased sensitivity of health outcomes to pollution exposure in lower-income communities (*2*). While these disparities in air pollution exposure across wealth groups have been well documented in high-income countries (*2*–*4*), evidence from low- and middle-income countries remains limited, primarily due to sparse air quality monitoring networks, especially in rural areas, along with a lack of data on wealth at fine temporal and spatial resolutions (*3*). A population-scale understanding of trends and exposure to air pollution, including in wealthier and poorer areas of low- and middle-income countries such as India, is urgently needed to understand and address the impacts of air pollution exposure across diverse socioeconomic groups.

Despite recent efforts in expanding the air quality monitoring network in the country, India still faces a great challenge in comprehensive measurement of surface air quality. The Central Pollution Control Board (CPCB) initiated the National Ambient Air Quality Monitoring (NAAQM) Network in 1984, beginning with the installation of manual monitoring stations, where pollutants are subsequently analyzed in the laboratory (*5*). Continuous Ambient Air Quality Monitoring System (CAAQMS), which provides real-time data, was first introduced in Delhi in 2006 and expanded to other cities after 2016 (*6*). The CPCB manages 883 manual stations and 438 continuous monitoring stations as of February 2023 (*7*), and this relatively small number of monitors results in a much higher ratio of population to continuous air quality monitors in India (3.2 million people per monitor) compared to the United States [0.1 to 0.5 million people per monitor (*8*)], EU [0.1 million people per monitor (*9*)], and China [0.9 million people per monitor (*10*)]. Moreover, most current CAAQMS monitors are situated in populous urban areas where wealthier people reside (fig. S1). The government of India has committed to expanding the CAAQMS network up to 1000 monitors under the National Clean Air Program (NCAP), which started in 2019 (*11*); however, the placement of these additional continuous monitors remains an ongoing policy question, and it is uncertain whether environmental inequalities are considered or prioritized in determining their locations.

Previous population-based studies (*12*, *13*) have explored the disproportionate exposure to air pollution and associated health impacts in low- and middle-income countries, including India, by using publicly available modeled $PM_{2.5}$ estimates. One widely used dataset for these analyses is global monthly estimates of $PM_{2.5}$, integrating satellite retrievals of aerosol optical depth (AOD), atmospheric chemical transport models, and ground-based measurements (*14*). While this dataset has proven valuable on a global scale, substantial uncertainties persist in regions with limited ground monitoring stations, including India (*14*). Furthermore, it is likely that region-specific models could substantially outperform global models in measuring air pollution in a target region of interest, as has been found in the United States (*15*).

In addition to the uncertainty in estimating $PM_{2.5}$ concentrations, the coarse temporal resolution of existing datasets (i.e., monthly) hinders the assessment of short-term health effects of $PM_{2.5}$, such as effects on all-cause, respiratory, and cardiovascular mortality (*16*–*23*). Furthermore, monthly estimates fail to capture short-term spikes in $PM_{2.5}$ emissions at the local to regional scales, such as crop residue burning. Acknowledging the need for a dataset with finer temporal resolution, a growing number of studies (*24*–*26*) have worked on developing daily $PM_{2.5}$ estimates for India. However, their datasets are not publicly accessible, inhibiting their use by both researchers and a

[1]Emmett Interdisciplinary Program in Environment and Resources, Stanford University, Stanford, CA, USA. [2]Department of Earth System Science, Stanford University, Stanford, CA, USA. [3]School of Marine and Atmospheric Sciences, Stony Brook University, Stony Brook, NY, USA. [4]Program in Public Health, Stony Brook University, Stony Brook, NY, USA. [5]Centre for Research on Energy and Clean Air, New Delhi, Delhi, India. [6]Department of Energy Science and Engineering, Stanford University, Stanford, CA, USA. [7]Woods Institute for the Environment, Stanford University, Stanford, CA, USA. [8]Precout Institute for Energy, Stanford University, Stanford, CA, USA. [9]Visiting Professor, Nova School of Business and Economics, Lisbon, Portugal. [10]Doerr School of Sustainability, Stanford University, Stanford, CA, USA. [11]Center on Food Security and the Environment, Stanford University, Stanford, CA, USA. [12]National Bureau of Economic Research, Cambridge, MA, USA.
*Corresponding author. Email: akawano@stanford.edu

host of governmental and non-governmental actors. Using up-to-date data to comprehensively characterize temporal and spatial trends in exposure, and potential differences in exposure by socioeconomic status, is also critical for understanding how and why exposures are changing and who is being most affected.

Here, we develop an open-source daily $PM_{2.5}$ dataset at 10-km resolution for India over 2005–2023 by training a machine learning model to predict the limited available ground monitor data with abundant data from remote sensing. Previous studies have used machine learning algorithms—such as neural networks, random forest, and extreme gradient boosting (XGBoost)—to predict ambient $PM_{2.5}$ concentrations (*27*, *28*). Because of the theoretical relationship between satellite AOD and surface $PM_{2.5}$ (*27*, *29*), satellite-derived AOD has long been a key input feature used to train machine learning models for air pollution measurement, often in combination with data on meteorology, land cover, elevation, and population density (*27*, *28*). However, the substantial amount of missing AOD values due to clouds and bright surfaces has posed challenges in reducing predictive errors when estimating $PM_{2.5}$ concentrations through machine learning methods (*27*, *30*). To address this issue, a common recent approach involves imputing missing data in AOD observations using machine learning methods (*24*, *25*, *31*). In addition, new satellite sensors provide alternative input features for predicting $PM_{2.5}$ concentrations without relying on AOD, such as data from the Sentinel-5 Precursor (Sentinel-5P) mission's TROPOspheric Monitoring Instrument (TROPOMI) (*30*, *32*, *33*). However, it is still unknown how much the predictive performance differs between satellite-derived AOD and TROPOMI-based features in estimating $PM_{2.5}$ concentrations, and whether the two sources of data independently add value in predicting surface $PM_{2.5}$ concentrations.

To take advantage of both the longer available time series of AOD data and information from these newer sensors, we use available ground monitoring data to train two separate models, which we term as the "Full model" and the "AOD model." The Full model combines both moderate resolution imaging spectroradiometer (MODIS) AOD (*34*) and TROPOMI satellite inputs [nitrogen dioxide (NO_2) (*35*) and carbon monoxide (CO) (*36*) along with other inputs; fig. S2 and Materials and Methods] and is trained on data from July 2018, when TROPOMI features become available, through September 2023 (figs. S2 and S3). It produces daily $PM_{2.5}$ estimates across the country for the period corresponding to its training period. The AOD model retains all inputs except TROPOMI and trains on monitor data beginning January 2015 (figs. S2 and S3) and is used to generate daily predictions from January 2005 to September 2023. In both the Full and AOD models, we first fill the missing satellite observations in either AOD or TROPOMI using a separate machine learning model (see Materials and Methods). The gap-filled estimates are then combined with the other features in a second-stage model that predicts surface $PM_{2.5}$ concentrations as measured by CAAQMS monitors ($n = 435$) (fig. S4).

Critically, and in contrast to related work (*24*, *25*), our two-stage model is evaluated using spatial cross-validation (CV) (fig. S5)—i.e., we evaluate model predictions on entirely held-out monitoring stations—rather than conventional random CV, in which a given station can contribute data to both the training and test sets. This more challenging performance metric is meant to ensure that the model generalizes to locations where it has no data to train, which is most of India. Using spatial CV, we then calculate two performance metrics: (i) the total coefficient of determination (R^2) or the percent of variation in observed $PM_{2.5}$ explained by model predictions and (ii) the "within" R^2 (see Materials and Methods) or the percent of variation in observed $PM_{2.5}$ explained by predictions after accounting for both differences in average $PM_{2.5}$ across locations and seasonal differences in $PM_{2.5}$ within a given location. In essence, the within R^2 measures whether our model can predict daily $PM_{2.5}$ anomalies relative to location- and season-specific averages, rather than simply predict spatial or seasonal patterns correctly. This "within" variation is often exploited in studies of the impact of air pollution of health and related societal outcomes and thus is a highly relevant, if rarely reported, performance metric.

We use model-derived predictions to better characterize nationwide air quality trends, including inequalities in $PM_{2.5}$ exposure by wealth levels and to identify locations with extreme $PM_{2.5}$ concentrations. We then use our predictions, along with emissions inventories, meteorological data, and administrative data on national air quality programs to better understand why pollution concentrations are changing. This includes, to our knowledge, the first evaluation of whether the recently adopted NCAP is improving air quality in targeted areas relative to a comparable set of nontargeted areas. Subsequently, to address existing disparities in the air quality monitoring network in India, we use our $PM_{2.5}$ estimates and compressed sensing methods (see Materials and Methods) (*37*) to identify strategic locations for a new set of air quality monitors, aiming to maximize the network's ability to capture variability in surface air pollution across both wealthier and lower-wealth regions. Last, we explore how our approach could supplement limited monitor data in other low- and middle-income countries by investigating the number of air quality monitors required to achieve reasonable model performance using our machine learning approach.

RESULTS

Model performance

The spatial out-of-sample performance of the Full model, assessed across daily $PM_{2.5}$ observations in the held-out fold, yielded an R^2 of 0.67 (Fig. 1A). Our model effectively predicts local and temporal $PM_{2.5}$ variation rather than average differences in $PM_{2.5}$ concentrations between locations, months, or years (within $R^2 = 0.49$) (Fig. 1A). The AOD model demonstrated a comparable performance, with an R^2 of 0.64 and within R^2 of 0.45 (fig. S6). We further validated its performance using 376 daily $PM_{2.5}$ observations from five monitors between 1 January 2013 and 31 December 2014. The performance of the AOD model on this dataset resulted in an R^2 of 0.54 and a within R^2 of 0.35, indicating moderate predictive accuracy for the limited areas during this period. When aggregated at the monthly level, our model substantially outperforms the existing publicly available monthly $PM_{2.5}$ dataset (*14*) when evaluated on monitoring data, with an R^2 of 0.74 and within R^2 of 0.52 (fig. S7). Examining location-specific performance, the out of sample within R^2 for each 10-km grid suggests that the northwest, northeast, and south regions exhibit higher within R^2 (Fig. 1B) (for the performance of AOD model, fig. S8). We find that these differences in regional performance are substantially driven by differences in the level and variance of $PM_{2.5}$ concentrations at the station level and the number of air quality monitors within 100 km (fig. S9); performance is much higher in locations with higher and more variable $PM_{2.5}$ and with more stations nearby, consistent with the model having an easier time learning patterns in these higher signal-to-noise areas.

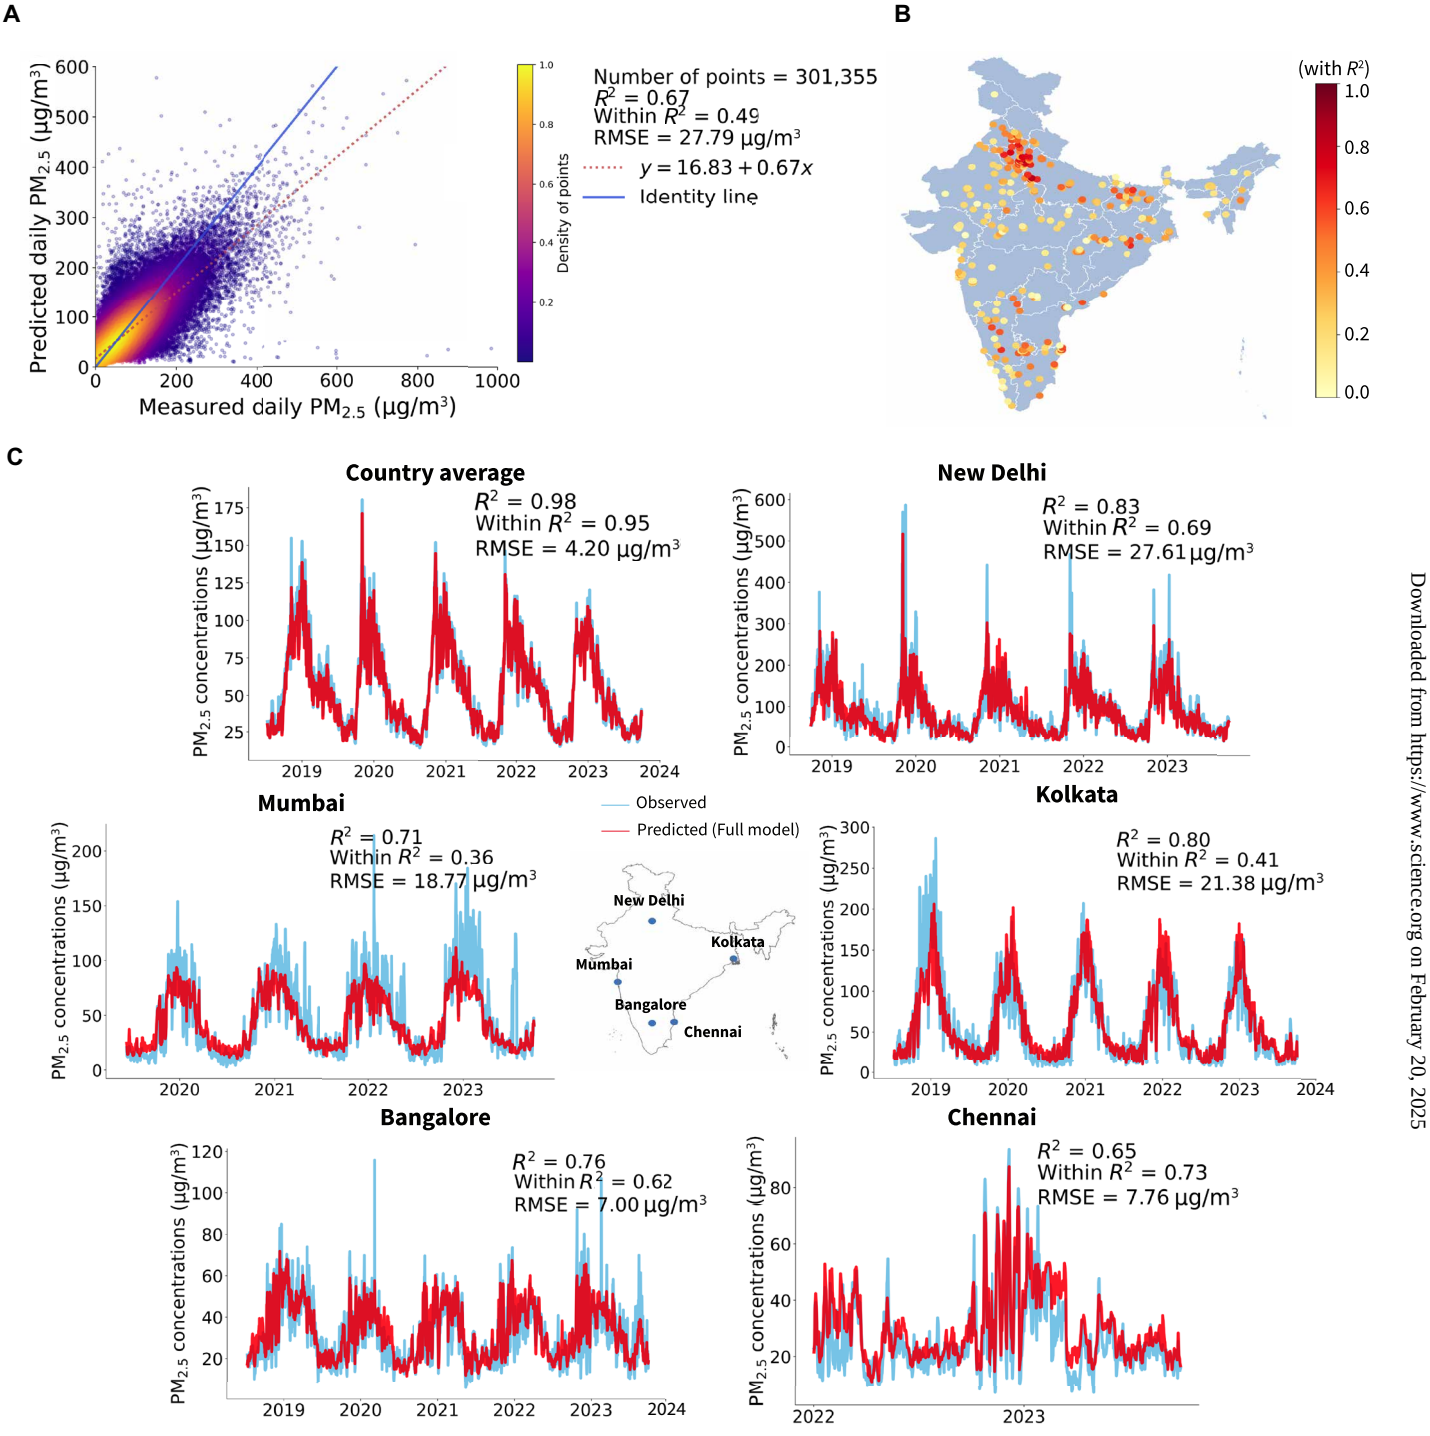

Fig. 1. Model performs well out-of-sample and across the range of observed $PM_{2.5}$ concentrations. (**A**) Overall performance of the Full model. Colors indicate the count of monitor observations, with observed $PM_{2.5}$ on the horizontal axis and predicted $PM_{2.5}$ concentrations on the vertical axis. The blue line represents the 1–1 line, indicating perfect match between predictions and observations. (**B**) Out-of-sample model performance in each grid cell with at least one monitor reporting on at least 5 days. Performance measured using "within R^2" after removing local seasonality and year trends. It is calculated over observed $PM_{2.5}$ using predicted $PM_{2.5}$ from the Full model, in which that station was out of sample, with month of the year and year-fixed effects. (**C**) Out-of-sample performance of the Full model in predicting daily time series of observed $PM_{2.5}$ concentrations for the entire country, New Delhi, Mumbai, Bangalore, Chennai, and Kolkata. Different *y*-axis and *x*-axis scales are used in the figures to accommodate variations in $PM_{2.5}$ concentrations and monitor availability across the cities.

To further evaluate the out-of-sample spatiotemporal performance of the Full model in different locations, the daily variations in the observed and predicted $PM_{2.5}$ are compared in five mega-cities in India: New Delhi, Mumbai, Bangalore, Chennai, and Kolkata (Fig. 1C) (for the AOD model, fig. S10). These five mega-cities are selected on the basis of their population size, with each being among the most populated cities in India (*38*). Among these cities, New Delhi shows the highest performance, R^2 of 0.83 and within R^2 of 0.69 (Fig. 1C). This aligns with our prior analysis on model performance because New Delhi demonstrates dense CAAQMS network and substantial variance in $PM_{2.5}$ concentrations (fig. S11), primarily driven by distinct seasonal patterns caused by meteorological conditions restricting pollutant dispersion and the concurrent operation of brick manufacturing around Delhi during winter (*39*, *40*).

We observe variations in model performance driven by distinct seasonality throughout the year across the winter (December to February), spring (March), summer (April to May), monsoon (June to September), and post-monsoon (October to November). Both the Full and AOD models demonstrate strong performance in dry seasons (winter and post-monsoon) (fig. S12). However, their performance declines during spring, summer, and monsoon periods (fig. S12), likely the result of clouds introducing noise in the remotely sensed input features during these wetter periods.

We find consistent results in sensitivity analysis, yielding an R^2 of 0.62 and within R^2 of 0.44 when the Full model underwent training and testing on significantly larger blocks based on latitude (figs. S13 and S14), which helps rule out the possibility that our high model performance is simply being driven by auto-correlation between nearby train and test locations. When evaluated using 10-fold random CV rather than spatial CV, our model showed notably higher performance (R^2 of 0.85 and within R^2 of 0.72) (fig. S15), highlighting the potential of random CV to overstate model performance on critical real-world applications (i.e., accurately predicting variation in locations where the model was not trained).

Furthermore, as a result of model comparison using the same sets of training and test data from July 2018 to September 2023, the Full model reported the highest performance, achieving an R^2 of 0.67 and within R^2 of 0.55 (fig. S16). The TROPOMI model, which excludes AOD but incorporates other features present in the Full model, slightly outperformed the AOD model by 0.01 in R^2 and within R^2 (fig. S16). This suggests that, at least in our setting, TROPOMI data can be a substitute for AOD in predicting $PM_{2.5}$ concentrations, which is perhaps appealing given their lower amount of missing data compared to AOD. Nevertheless, the combined use of both TROPOMI and AOD features provides the strongest predictive power (fig. S16).

The better performance of the TROPOMI model compared to the AOD model may seem to contradict the mechanistic expectation of a stronger relationship between $PM_{2.5}$ concentrations and AOD than between $PM_{2.5}$ and CO or NO_2. However, in our setting, observed TROPOMI NO_2 is as strongly correlated with $PM_{2.5}$ monitor data as observed AOD, and observed TROPOMI CO is more strongly correlated than AOD, with correlation coefficients of 0.35, 0.35, and 0.47, respectively, potentially due to the large amount of missing data in observed AOD. Furthermore, the stronger correlation between $PM_{2.5}$ and CO, a by-product of carbon-containing fuel combustion, aligns with the substantial contribution of residential combustion and crop residue burning to air pollution in India (*41*).

Long-term spatiotemporal trends in predicted $PM_{2.5}$ concentrations

To better understand longer-term shifts in $PM_{2.5}$ concentrations, we calculate the changes in average $PM_{2.5}$ concentrations between 6-year blocks, beginning in 2005–2010 (Fig. 2A) and ending in either 2011–2016 (Fig. 2B) or 2017–2022 (Fig. 2C). We find that much of India experienced substantial increases in $PM_{2.5}$ concentrations between 2011 and 2016 compared to 2005–2010, except for regions such as Jammu and Kashmir, Punjab, and Rajasthan states. However, the pace of increase moderated in more recent years, and during 2017–2022, a higher percentage of locations across the country showed decreases in $PM_{2.5}$ concentrations (45% of grid cells at a 10-km resolution) compared to 2011–2016 (16% of grid cells), with notable decreases observed in Jammu and Kashmir, Punjab, Haryana, Delhi, and Rajasthan states and union territories.

In addition, given the noise in our daily, grid-level estimates [root mean square error (RMSE) = 27.79 $\mu g/m^3$], we conducted further analysis to understand whether our predictions had enough signal to detect changes in $PM_{2.5}$ concentrations in temporally aggregated data. Specifically, we performed simulations by generating data that had the daily RMSE of our modeled data but that had a "true" reduction over time in $PM_{2.5}$ ranging from 0.5 to 5 $\mu g/m^3$ and studied whether we were able to recover this reduction (see Materials and Methods). The results demonstrated that the mean of the simulated changes in both the monitor data and predictions is closely aligned with the "true change" values (fig. S17), highlighting that the $PM_{2.5}$ predictions generated by our machine learning model are robust enough to detect long-term small changes at the 10-km grid cell level.

To characterize trends in $PM_{2.5}$ concentrations, we quantify population-weighted annual average $PM_{2.5}$ concentrations for the country overall and five mega-cities from 2005 to 2022 (Fig. 2D) by combining our 10-km $PM_{2.5}$ estimates with gridded population data. Among the five mega-cities, the average New Delhi resident has consistently faced the highest population-weighted average of $PM_{2.5}$ concentrations, with 88.67 $\mu g/m^3$ in 2022, more than double India's annual national air quality guideline of 40 $\mu g/m^3$ (Fig. 2D) (*42*). Similarly, residents in India overall, as well as those in Kolkata and Mumbai, have consistently experienced $PM_{2.5}$ levels exceeding the national annual threshold (Fig. 2D). Notably, we estimate that residents in all mega-cities have experienced a moderate decline in $PM_{2.5}$ exposure since 2016–2018 (Fig. 2D), as assessed by computing the 3-year rolling averages of population-weighted $PM_{2.5}$ concentrations (Fig. 2E). Mumbai exhibited the most substantial decline of 10%, followed by New Delhi with 8% in 2020–2022 (Fig. 2E).

We further examine whether the observed declining trend since 2016–2018 is attributable to favorable meteorology, such as increasing trends in precipitation and relative humidity (*43*, *44*) observed in 70 and 90% of 10-km grid cells, respectively, from 2005–2015 to 2016–2022 (fig. S18). Using trend analysis (see Materials and Methods), we compare the observed average annual trend for 2005–2015 and 2016–2022 with the meteorology-corrected trend for the same periods. Our analysis revealed that the declining trend in $PM_{2.5}$ concentrations from 2016 to 2022 was influenced by meteorological variability in the southern regions, but not in the northern regions such as Delhi, Haryana, Rajasthan, Uttar Pradesh, and Bihar states and union territories (fig. S19). This suggests that without meteorological influence, the southern regions would have experienced fewer decreases in $PM_{2.5}$ concentrations

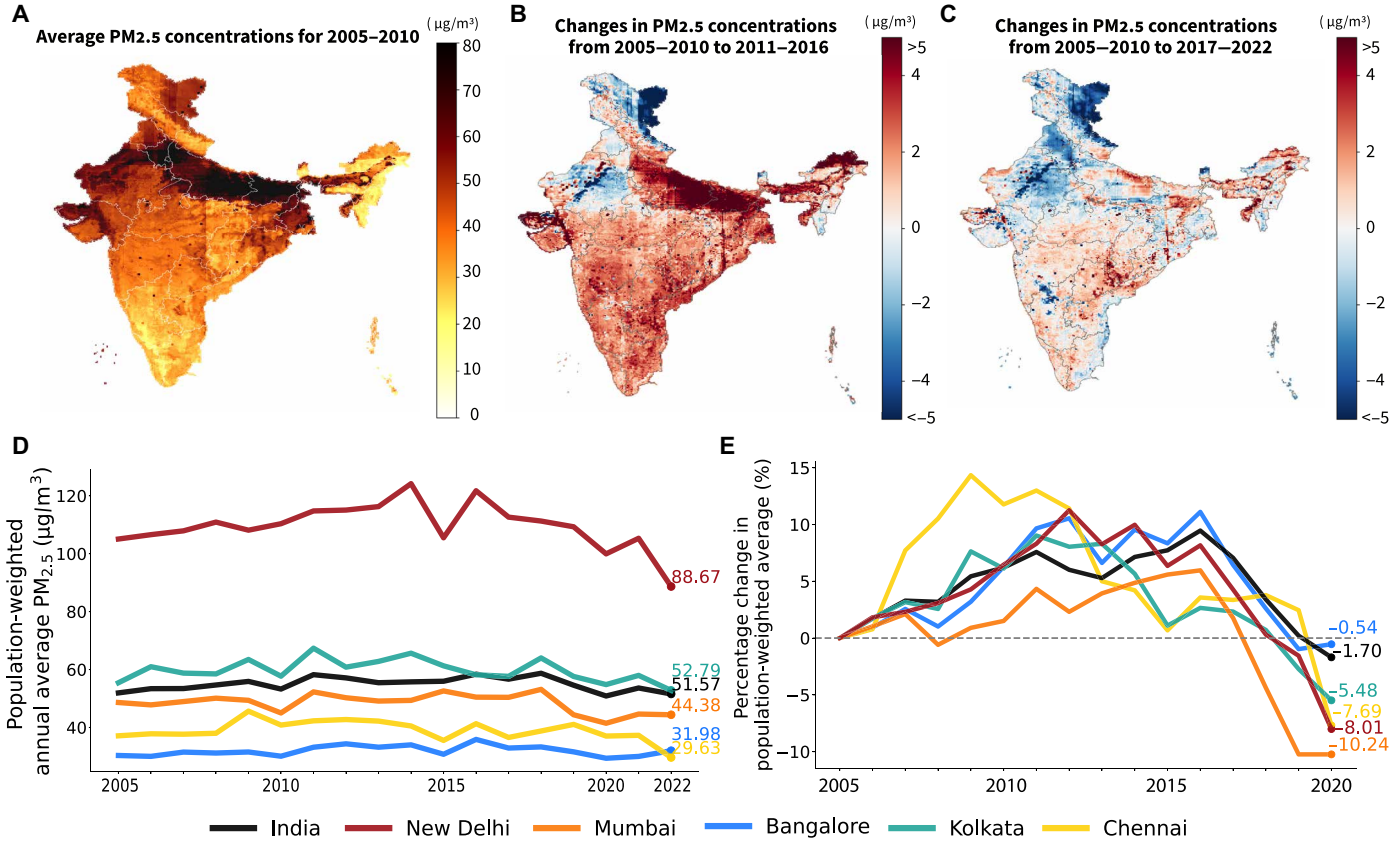

Fig. 2. $PM_{2.5}$ concentrations increased through 2016 but then declined through much of India thereafter. (**A**) Six-year average $PM_{2.5}$ concentrations for 2005–2010 computed as the average over all days in each grid cell for 6 years. (**B**) Changes in $PM_{2.5}$ concentrations between 2005–2010 and 2011–2016 show increases across most of the country. (**C**) Changes in $PM_{2.5}$ concentrations between 2005–2010 and 2017–2022 show a mix of increases and declines. (**D**) Population-weighted annual average $PM_{2.5}$ concentrations from 2005 to 2022 for all of India and selected mega-cities. (**E**) Percentage changes in 3-year population-weighted averages relative to the 2005–2007 average for all of India and selected mega-cities. The *x*-axis label represents the running means of years from 2005–2007 to 2020–2022.

from 2016 to 2022. In contrast, we find little evidence that the increasing trend from 2005 to 2015 was driven by changing meteorology (fig. S19), suggesting that these increases could be attributable to increased anthropogenic activities. We then confirm these trends using emissions data obtained from the Emissions Database for Global Atmospheric Research (*45*, *46*), focusing particularly on $PM_{2.5}$ and black carbon (BC) emissions. The northern, western, and southern regions experienced decreases in $PM_{2.5}$ and BC emissions from 2020 to 2022 compared to 2016–2018, while substantial increases were observed across the country between 2005–2007 and 2013–2015 (fig. S20).

As declines in $PM_{2.5}$ since 2016 are not attributed to favorable meteorology, we examine whether India's air quality control policies, particularly the NCAP, have contributed to recent reductions in ambient $PM_{2.5}$ concentrations. Acknowledging the need for improved air quality to reduce health and societal burdens in the country, the government of India has recently proposed and implemented a number of air quality control measures, such as NCAP, the Bharat Stage-VI (BS-VI) emission standards for vehicles that mandated vehicles to adhere to PM emission limits as strict as European standards, which went into effect in Delhi in 2018 and other parts of India in 2020 (*47*), and the closure of multiple coal-fired power plants located near Delhi (*48*). However, the specific contribution of each policy to improving nationwide and regional air quality remains uncertain.

NCAP was initiated in 2019 with the goal of reducing key air pollutants, including $PM_{2.5}$, by 20 to 30% by 2024, using the pollution levels observed in 2017 as a baseline (*7*). Focusing on 131 non-attainment cities across 28 states and union territories, selected based on air quality data from 2015 to 2019 (*7*), one of the primary objectives of NCAP is to prompt each non-attainment city to prepare and implement a clean air action plan that details sector-specific interventions to improve air quality with predetermined timelines and an agency responsible for execution of each intervention (*7*, *11*). Under the NCAP, 102 of the 131 non-attainment cities submitted comprehensive city action plans, which were approved by the CPCB in July 2020 (*11*).

We use a difference-in-differences approach to assess whether the implementation of the NCAP has affected changes in ambient $PM_{2.5}$ concentrations to date. Our approach compares within-subdistrict changes in $PM_{2.5}$ over time, in targeted and nontargeted areas, before and after initiation of NCAP. Our analysis is constrained to the subdistrict level, rather than the city level, due to the limited availability of reliable city-level shapefiles in India. We denote "treated" subdistricts as the 102 non-attainment cities whose city clean action plans were approved in 2020 and select a set of corresponding control subdistricts

that were not targeted by NCAP using a propensity score method that matches pretreatment trends in air pollution and covariates between later-treated and never-treated subdistricts (see Materials and Methods). To account for spillover effects, we exclude subdistricts adjacent to treatment subdistricts and any others within a 50-km buffer. Consequently, 88 treatment subdistricts and 74 control subdistricts are included (Figure S21), and the effect of NCAP is estimated by comparing whether trends in air pollution diverged between treated and control units after 2020. Our analysis reveals that there is no evidence that NCAP contributed to reducing $PM_{2.5}$ concentrations both in 2021 and 2022 (Figure S22) in targeted cities.

The observed decreases in $PM_{2.5}$ concentrations, especially in the mega-cities in 2020, are instead more consistent with previous studies examining the impact of the COVID-19 pandemic on air quality in India (*49*–*53*), which highlighted a substantial decrease (43%) in $PM_{2.5}$ in 2020 compared to 2017–2019 in urban areas, after controlling for meteorological variability (*49*). Other various air quality control policies in India, including the implementation of BS-VI emission standards and the closure of power plants near Delhi, may have contributed to the regional $PM_{2.5}$ declines. Our $PM_{2.5}$ estimates could be used to evaluate the impact of these programs in future work.

Population exposure to $PM_{2.5}$ concentrations

To understand the population exposure to daily high levels of $PM_{2.5}$ concentrations, we calculate the average number of days that each grid cell exceeded the World Health Organization (WHO) guideline of 15 $\mu g/m^3$, the national guideline of 60 $\mu g/m^3$, and the extreme concentration of 100 $\mu g/m^3$ between 2018 and 2022. Notably, much of India experienced over 300 days above the WHO daily threshold, except for the northeastern region (Fig. 3A). Delhi, Rajasthan, Uttar Pradesh, and Bihar states and union territories encountered days exceeding the national guideline of 60 $\mu g/m^3$ for at least 250 days (Fig. 3A). Moreover, certain locations in Delhi, Uttar Pradesh, and Bihar observed extreme days with $PM_{2.5}$ concentrations exceeding 100 $\mu g/m^3$ for 100 to 150 days (Fig. 3A).

We also assess the proportion of the overall population exposed to elevated annual concentrations of $PM_{2.5}$ over our study period. Notably, the entire population in India has consistently faced exposure above the WHO annual threshold of 5 $\mu g/m^3$ over the 17 years (Fig. 3B). Although there was a decrease in 2020, approximately 60% of the population consistently experienced exposure exceeding the national annual guideline of 40 $\mu g/m^3$, and 10% experienced extreme levels of $PM_{2.5}$, with an annual average of 80 $\mu g/m^3$ (Fig. 3B). Further analysis of the spatial distribution of these exposed populations revealed that 63% of the locations exceeded the national guideline between 2018 and 2022 (Fig. 3C). Areas with annual average $PM_{2.5}$ concentrations exceeding 80 $\mu g/m^3$ were predominantly observed around Delhi and in Bihar state (Fig. 3C).

Inequalities in $PM_{2.5}$ exposure

Identifying disparities in $PM_{2.5}$ exposure by socio-economic status is essential to understanding pollution burdens and developing policy

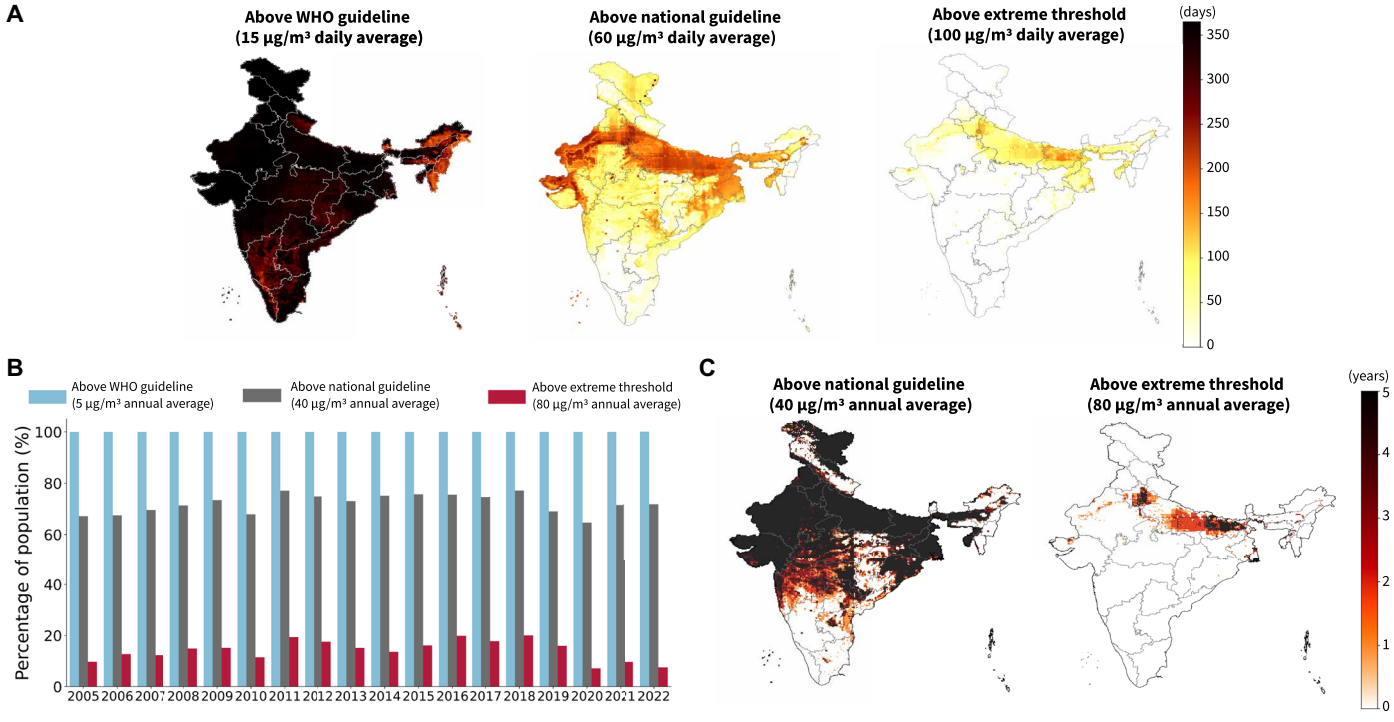

Fig. 3. Widespread exposure to elevated $PM_{2.5}$ concentrations across India, with daily extremes in the North. (**A**) Spatial distribution of grid cells with daily $PM_{2.5}$ concentrations above WHO guideline of 15 $\mu g/m^3$ (*74*) (left), national guideline of 60 $\mu g/m^3$ (middle), and daily extreme threshold of 100 $\mu g/m^3$ (right), demonstrating the average number of days for each 10-km grid cell for 2018–2022. (**B**) Proportion of populations exposed to annual averages of $PM_{2.5}$ concentrations exceeding WHO (5 $\mu g/m^3$) (*74*), national (40 $\mu g/m^3$), and extreme (80 $\mu g/m^3$) thresholds from 2005 to 2022. The entire population in India is consistently exposed to $PM_{2.5}$ concentrations above the WHO guideline during these years. (**C**) Locations highlighted in colors indicate areas where average $PM_{2.5}$ concentrations exceeded the national guideline of 40 $\mu g/m^3$ (left) and extreme threshold of 80 $\mu g/m^3$ (right) from 2018 to 2022, respectively. Colored gradients depict the number of years each grid cell exceeded each threshold during these years. The entire country would be shaded in colors on a map denoting the areas exposed to above the WHO threshold.

measures to alleviate them. To understand how pollution concentrations vary with socioeconomic status in India, we combine recent high-resolution estimates of local-level asset wealth, generated using the Demographic and Health Surveys (DHS) data from 2015 and 2019 (*54*) and a common proxy for socioeconomic status, with spatial and temporal variation in our $PM_{2.5}$ predictions at a 10-km resolution.

Using 5-year average $PM_{2.5}$ concentrations from 2015 to 2019 (see Materials and Methods), we find that a smaller proportion of the population in the wealthiest quintile is exposed to concentrations above the national annual guideline of 40 μg/m^3, while higher proportions are observed in other quintiles of the wealth distribution (Fig. 4A), although substantial majorities in all quintiles are exposed to levels above this guideline. However, higher proportions of people living in areas in the top two wealth quintiles are exposed to extreme $PM_{2.5}$ concentrations above 80 μg/m^3 annually (23.0 and 21.0%, respectively) than poorer areas (Fig. 4B and fig. S23), consistent with a previous study revealing that wealthy populations live in polluted urban centers in low- and middle-income countries (*3*). When stratified by urban and rural status based on the average level of urbanicity per grid cell as defined by the Urban and Rural Catchment Area data (*55*), higher proportions of people living in middle, poorer, and poorest communities within urban areas are exposed to higher PM2.5 concentrations above the national guideline and extreme threshold than people living in richer and richest communities (fig. S24).

Due to extreme exposures experienced by people living in wealthier communities, an average person at the 90th percentile of wealth in India has consistently faced higher $PM_{2.5}$ exposure than the average Indian or someone at the 10th percentile of wealth from 2005 to 2022, holding wealth constant across years (Fig. 4C). Notably, since 2016, an average wealthy individual has also experienced faster declines in $PM_{2.5}$ concentrations than an average poor individual, particularly evident since the start of the COVID-19 pandemic yet notably apparent as early as the pre-COVID era of 2017–2019. (Fig. 4, C and D, and fig. S25). This has contributed to shrinking the wealth gap in average $PM_{2.5}$ over time. These $PM_{2.5}$ reductions disproportionately experienced by wealthier individuals highlight the urgent need of air quality mitigation policies that effectively and intentionally target both the poor and the rich.

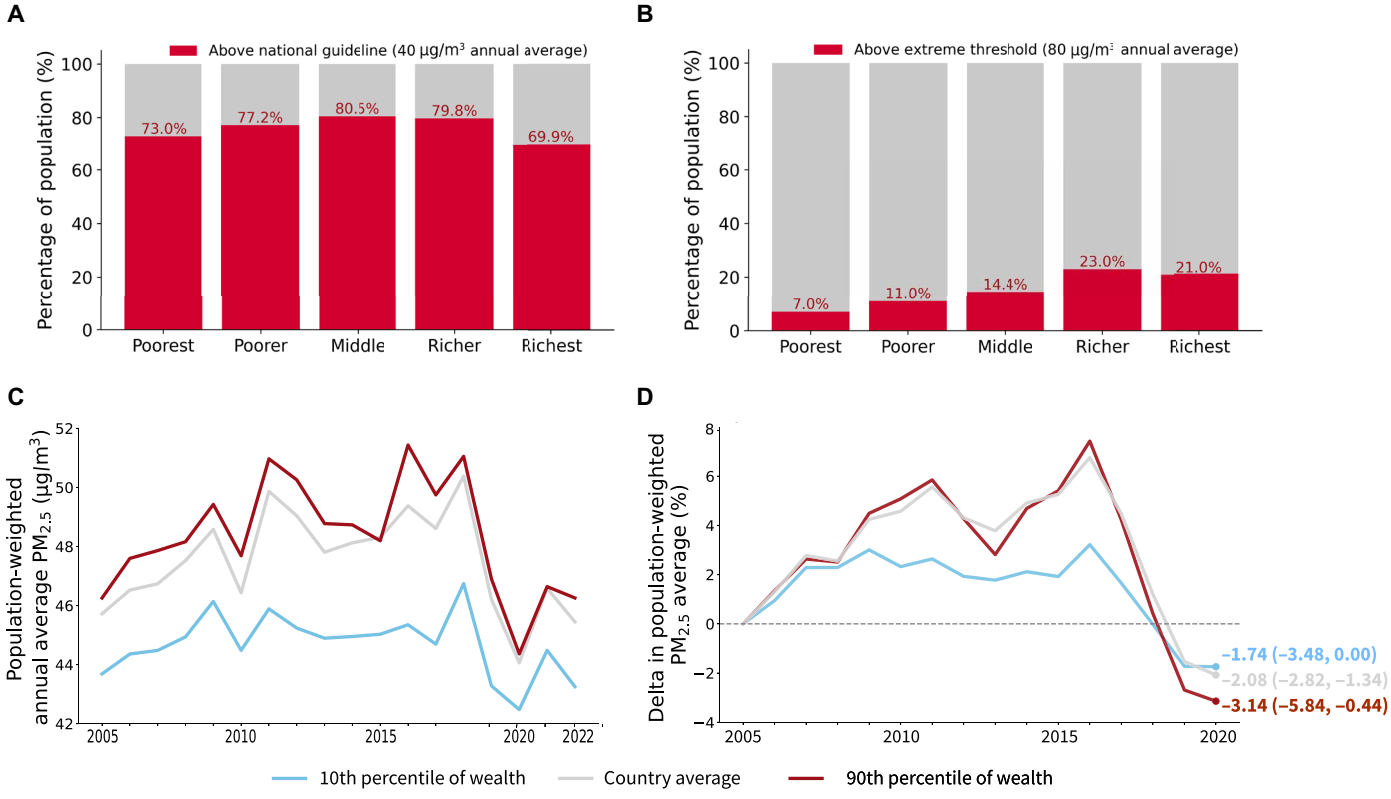

Fig. 4. Wealthier people experienced higher average $PM_{2.5}$ concentrations but faster recent declines. (**A**) Percentage of the population in each wealth category exposed to ambient $PM_{2.5}$ concentrations above the national guideline of 40 μg/m^3 for the years 2015–2019. These percentages were calculated using 5-year mean $PM_{2.5}$ concentrations from 2015 to 2019 because machine learning algorithms to generate wealth estimates (*54*) were trained on ground data from that period. (**B**) Percentage of the population in each wealth category exposed to concentrations above the extreme threshold of 80 μg/m^3 for the years 2015–2019. (**C**) Population-weighted annual average of $PM_{2.5}$ concentrations from 2005 to 2022 for the country average, locations with 90th percentile of wealth, and locations with 10th percentile of wealth. (**D**) Percentage changes in 3-year population-weighted averages relative to the 2005–2007 average. The *x*-axis label represents the running means of years from 2005–2007 to 2020–2022. Between 2020 and 2022, wealthier individuals experienced a 3.14% (5.84%, 0.44%) decline in $PM_{2.5}$ exposure compared to the period of 2005–2007, while poorer individuals saw a 1.74% (3.48%, 0.00%) reduction. When compared to the period of 2015–2017, during which much of India began to experience an overall declining trend in $PM_{2.5}$ concentrations, a wider disparity in reduction rates was observed, with the wealthiest experiencing an 8.12% (10.56%, 5.64%) reduction and the poorest experiencing a 3.62% (5.32%, 1.88%) reduction (fig. S25). The values in parentheses indicate the 95% confidence intervals obtained from bootstrapping to account for the uncertainty in the modeled estimates of $PM_{2.5}$ concentrations.

Assessing the optimal placement of air quality monitors ensuring equality

The current CAAQMS network is sparser in poorer communities, limiting understanding of disproportionate $PM_{2.5}$ exposure. While our predictions enable investigation of nationwide trends and exposure across wealth groups, ground monitor data would enhance precision of such monitoring, and ground monitor data will likely remain the basis for official evaluation of air quality trends and policy attainment. The government of India has committed to expanding the CAAQMS up to 1000 monitors under NCAP (*11*) to aid in more comprehensive air pollution monitoring. However, it is uncertain whether placement decisions account for the ability to accurately monitor pollution concentrations across the socioeconomic spectrum.

We use compressed sensing methods (*37*, *56*) to propose strategic placement of additional 565 CAAQMS monitors or the remainder of the 1000 total monitors proposed to be installed under NCAP using data from our AOD model to inform monitor placement (see Materials and Methods). Our approach identifies baseline national-scale long-term variability of $PM_{2.5}$ concentrations and then chooses locations of additional monitors that would optimally capture local and short-term anomaly spikes in $PM_{2.5}$ exceeding this baseline across the country. When placing monitors, we prioritize populous low-wealth locations to ensure that sudden spikes occurring in poorer communities are captured. The identified placement of monitors (Fig. 5A) highlights the need for an additional dense network in northern and northeastern India, particularly in Rajasthan, Delhi, Haryana, Uttar Pradesh, Bihar, Assam, and West Bengal states and union territories, as well as in the central region, such as Madhya Pradesh state. The proposed additional network would help promote equality in nationwide exposure assessment while also enabling the network to maximally capture spatial and temporal variation in surface $PM_{2.5}$ concentrations.

Examining applicability of our model to other low- and middle-income settings

Other low- and middle-income countries also face challenges in understanding overall levels and trends in population exposure to $PM_{2.5}$ as well as inequalities in these exposures due to limited ground-monitoring data. Our method could provide resource-efficient alternative to establishing extensive monitoring network by generating predictions for nonmonitored locations, but it relies on having at least some amount of ground monitoring data to train and validate predictions. To understand how the performance of our machine learning approach changes as the ground network becomes sparser, we vary the number of monitors our model is allowed to see in training and quantify the relationship between the number of monitors in training and model performance. To estimate uncertainty in performance, we repeat this experiment a thousand times, resampling a fixed number of stations for training each time and re-estimating model performance on a disjoint set of sampled test stations (for more details, see Materials and Methods).

We observed a nonlinear increase in model performance, as evaluated on held-out test monitors, ranging from an R^2 value of 0.29 with 5 monitors to 0.68 with 300 monitors (Fig. 5C). When evaluated at the monthly level using the same sets of daily predictions derived from this experiment, we achieved R^2 and within R^2 values comparable to those of the existing benchmark publicly-available monthly $PM_{2.5}$ dataset (fig. S7) (*14*) while training only on 15 and 50 monitors, respectively (Fig. 5D). Specifically, we achieved an R^2 of 0.57 for 15 monitors (219,150 km^2 per monitor), exceeding the benchmark of 0.54, and a within R^2 of 0.40 for 50 monitors (65,745 km^2 per monitor), exceeding the benchmark of 0.37. These results highlight the importance of investing in even a limited number of reference-grade air quality monitors. Such an investment not only enhances the geographic coverage and reliability of air quality assessments within a country but also improves model performance. By integrating ground-based data with satellite information, our method can predict $PM_{2.5}$ concentrations with performance that exceeds benchmark datasets commonly used for health impact analysis in low- and middle-income countries. While variations in $PM_{2.5}$ concentrations and country sizes differ, these findings offer valuable insights for other countries designing their monitoring network and implementing our machine learning method to understand nationwide trends and exposure to $PM_{2.5}$ across different populations.

DISCUSSION

Here, we generate daily $PM_{2.5}$ predictions at a spatial resolution of 10 km across India from 2005 to 2023. These daily estimates perform well over the range of observed monitor $PM_{2.5}$ measurements and accurately capture temporal variations in $PM_{2.5}$ concentrations, including daily peaks, within mega-cities in India. We find a declining trend in average $PM_{2.5}$ concentrations since 2016–2018, particularly in northern India, and confirm that these reductions are not attributable to meteorological variability nor to NCAP, a recently begun nationwide air quality improvement program; smaller declines in southern India are driven in part by favorable trends in meteorology. Our analysis also provides a comprehensive characterization of the spatial distribution of populations exposed to elevated levels of daily and annual $PM_{2.5}$, revealing that higher proportions of people living in wealthier areas are exposed to extreme $PM_{2.5}$ concentrations but that they have also experienced faster reductions in exposure in recent years. We propose the strategic placement of additional CAAQMS monitors to more effectively capture high $PM_{2.5}$ episodes occurring in both poorer and wealthier locations, and we study the applicability of our approach in settings where existing or proposed monitoring networks could be even sparser than in our study context, finding that only a relatively small number of monitors are needed to train a relatively accurate prediction model.

In comparison to many existing efforts to estimate $PM_{2.5}$ concentrations using machine learning (*27*, *28*), we incorporate data from multiple recent satellite sensors to estimate pollution concentrations across the country. In addition, we spatially validate predictions against time series of held-out monitor observations, which stands in contrast to the random CV used in many previous machine learning–based efforts (*24*, *25*). Last, our work complements recent machine learning–based studies to estimate $PM_{2.5}$ concentrations by providing insights into the predictive power of TROPOMI features in contrast to AOD.

Our $PM_{2.5}$ predictions could likely be further improved through improvements in both the monitor-based ground truth data and in the remotely sensed input features. Our method relies on ground $PM_{2.5}$ observations acquired from CAAQMS monitors for training; however, the management of these monitors and the quality of the collected data have not been verified by a third-party institution (*5*). For instance, in the United Kingdom, all regulatory air quality data collected as part of the automatic urban and rural network are validated by an independent agency (*5*). Using the quality-assured

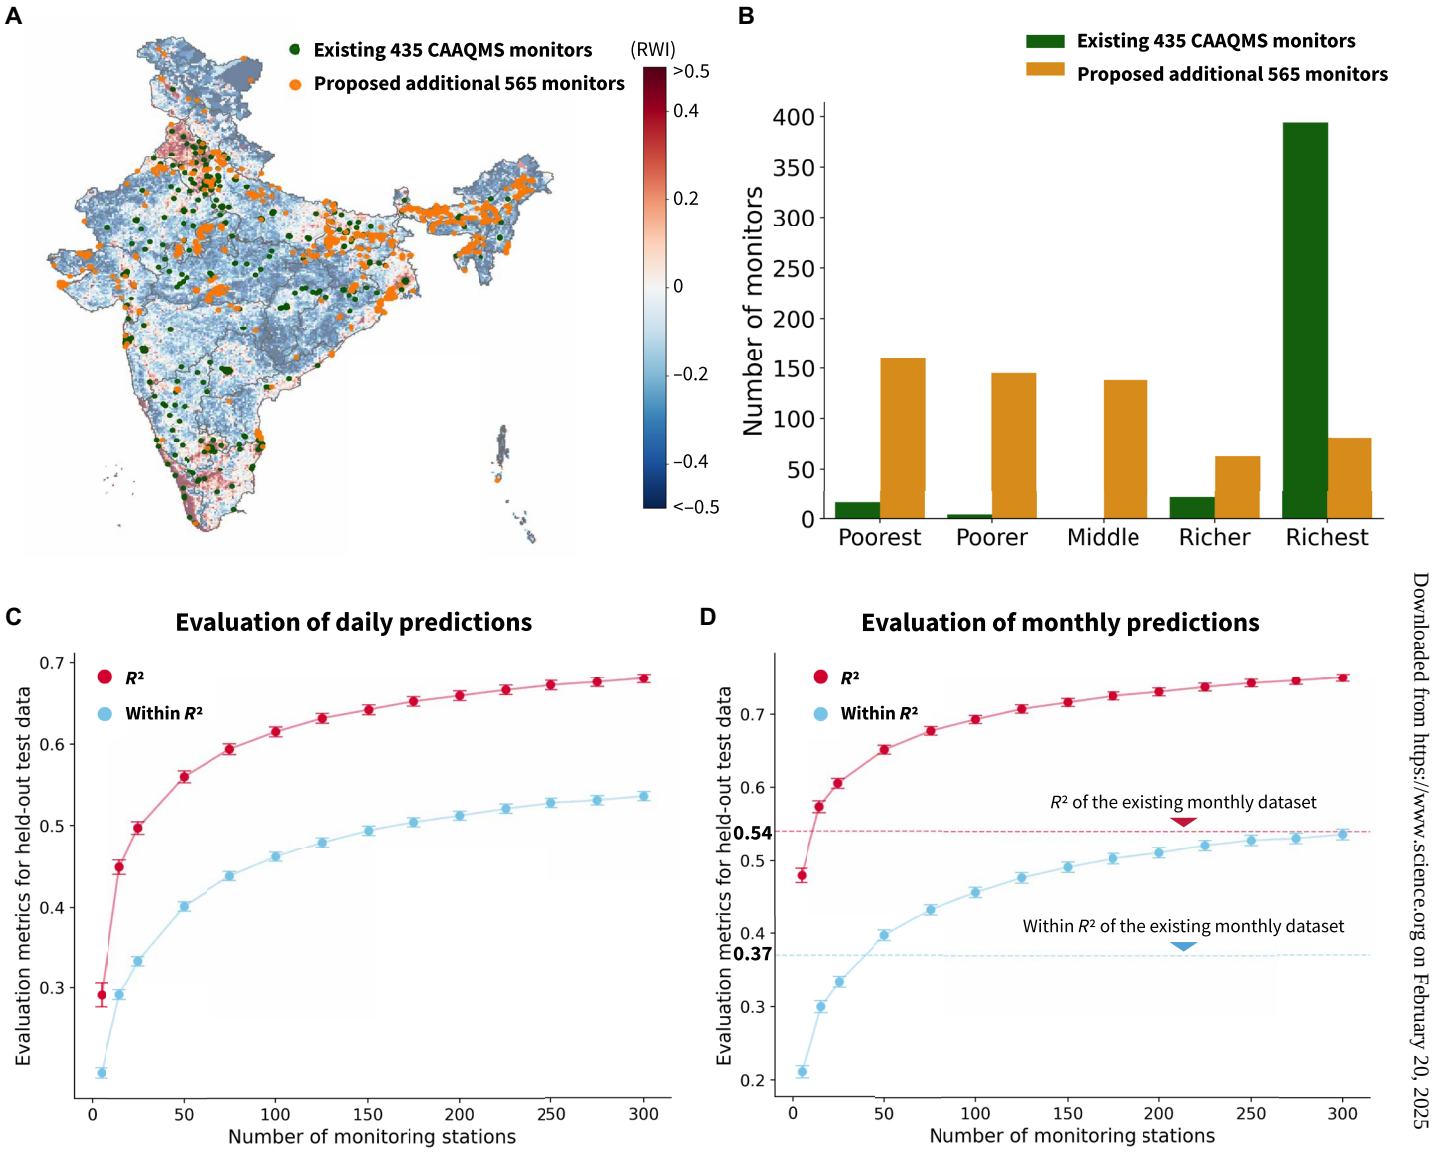

Fig. 5. Applications of our predictions for better air quality monitoring in India and other low- and middle-income countries. (**A**) Proposed locations of additional 565 CAAQMS identified by compressed sensing methods (see Materials and Methods), along with existing 435 monitor locations, ensuring increased equality in air quality monitoring across the country while capturing both baseline patterns and sudden anomaly spikes in $PM_{2.5}$ concentrations. Colored gradients indicate wealth estimates [Relative Wealth Indices (RWIs) (*54*)] in each grid cell. (**B**) Number of monitors in each wealth category based on quintiles of RWIs of each 10-km grid cell for existing CAAQMS monitors and proposed additional monitors. (**C**) Model performance at a daily level as a function of number of stations used in training. Performance is measured by R^2 and within R^2 as evaluated on a fixed number of held out stations (21 stations). Points represent the mean performance across 1000 experiments (randomly resampling training and test sets), and whiskers denote the corresponding 95% confidence intervals across experiments. Model performance increases with additional training data, but at a declining rate, with modest increases past 150 monitors. (**D**) Model performance at a monthly level, estimated by aggregating the daily predictions and monitor data on a monthly basis. The dashed lines represent the R^2 and within R^2 values of the existing publicly available monthly $PM_{2.5}$ dataset, evaluated against monitor data in India, which are 0.54 and 0.37, respectively.

ground measurements could help improve our predictions by reducing noise in both model training and validation. In addition, future model development could benefit from additional monitor data, including PurpleAir monitors and CPCB's manual monitoring stations. Calibration is imperative for these data as they are not considered as reference grade or regularly calibrated. Calibration of PurpleAir measurements to estimate $PM_{2.5}$ concentrations has been widely explored by previous studies (*57*–*61*). A recent study (*62*) highlights the significance of seasonally-optimized calibration for PurpleAir sensors in enhancing prediction performance, especially in India. Calibration of manual monitor measurements is also crucial since $PM_{2.5}$ samples are collected for 8 hours twice a week, providing only a snapshot of actual concentrations (*5*). Moreover, the remotely-sensed input features used in our model, particularly TROPOMI data, do not directly represent air pollution concentrations at the ground level. A recent study (*63*) revealed a mean relative and

absolute bias of approximately 10% between TROPOMI NO_2 products and ground-based observations, highlighting a tendency for frequent underestimation of elevated NO_2 levels on the ground (*63*). The robust predictive power of TROPOMI features indicates that the calibration of these data could lead to improved performance of machine learning–based estimations for $PM_{2.5}$ concentrations, although our machine learning–based approach is implicitly calibrating these satellite observations to ground data already. Last, uncertainty quantification from machine learning models is currently an active area of research. Subsequent enhancements to $PM_{2.5}$ estimates may involve more granular quantification of uncertainty.

Another limitation of our study is that we relied on cross-sectional wealth estimates for investigating temporal changes in wealth disparities in $PM_{2.5}$ exposure. The wealth data we used for India were derived from machine learning models trained on ground data from 2015 and 2019 and thus might not capture shifts in the wealth distribution in other years. While these wealth estimates represent the most comprehensive and up-to-date local-level estimates of income or wealth in India that we are aware of (*3*, *54*), future improvements to these data could further improve our understanding of spatial distribution of and temporal changes in income disparities in $PM_{2.5}$ exposure across the country.

Our publicly available $PM_{2.5}$ predictions serve as a platform for evaluating specific policies or interventions aimed at improving air quality. We used our daily $PM_{2.5}$ estimates in the initial evaluation of India's NCAP, finding a limited impact of the program to date in targeted cities. This null result could be because city-level clean air action plans did not yet have time to take effect, or because they are not effective, and our $PM_{2.5}$ data—which can be updated in future years—offer a platform for understanding which explanation is more likely true. They also offer the critical opportunity to evaluate other air quality control measures being rolled out across the country and an opportunity to identify the contribution to local $PM_{2.5}$ concentrations of emissions from specific sources such as brick kilns that exhibit distinct spatial and temporal patterns. Last, our data could also be used to better assess the health burden on air pollution in the country—a task often accomplished using existing monthly $PM_{2.5}$ datasets that are more temporally coarse and less accurate than the data we produce here.

MATERIALS AND METHODS

Model inputs

We collect daily $PM_{2.5}$ observations from 1 January 1 2013 to 30 September 2023, from 435 CAAQMS monitors with accessible geographic coordinate information, serving as ground truth for our machine learning model. We construct two main machine learning models to predict $PM_{2.5}$ concentrations: the AOD model and the Full model. The input features of the AOD model include MODIS AOD, meteorology, land cover, and elevation collected from the Google Earth Engine (GEE) platform, along with atmospheric reanalysis data retrieved from NASA's Earthdata portal. In addition to these features, the Full model incorporates Sentinel-5P mission's TROPOMI for NO_2 and CO, launched on 13 October 2017, by the European Space Agency. The AOD model is trained on ground-measured $PM_{2.5}$ from 1 January 2015 to 30 September 2023 and used to generate daily $PM_{2.5}$ predictions starting from 1 January 2005, which corresponds to the earliest available input feature for the model. The model was trained from 2015 due to the limited number of available observations during 2013–2014, when data from only five monitors were available. Conversely, the Full model is trained from 10 July 2018 to 30 September 2023 due to the limited availability of TROPOMI features and used to produce predictions for the corresponding period. Both $PM_{2.5}$ measurements and input features are merged to a consistent 10-km grid for model training and validation. The grids are constructed to cover country borders of the Republic of India as per survey of India records, resulting in a total of 33,074 grid cells.

From the collected $PM_{2.5}$ observations, we exclude values of 999.99 $\mu g/m^3$, the upper detection limit of CAAQMS monitors, as it does not accurately represent the actual concentration on the ground (*64*). In addition, we filter out $PM_{2.5}$ measurements if the difference between the rolling average of the preceding 5 days and the $PM_{2.5}$ concentration on the current day is less than 0.05 to ensure that air quality monitors exhibit valid variations in daily concentration changes (i.e., we remove observations when they were at an unrealistic constant level for over 5 days). As a result, 4213 observations (1.2%) were removed. Subsequently, for each 10-km grid, we exclude extreme outliers identified by an interquartile range (IQR) that fall below 15 times the first quartile or exceed 15 times the third quartile, resulting in the removal of 24 observations. The use of a threshold of 15 times allows us to identify $PM_{2.5}$ measurements that significantly deviate from the IQR within each 10-km grid and helps retain elevated observations that may reflect local variations, such as those caused by agricultural fires.

The complete list of input features for the AOD and Full models can be found in table S1. TROPOMI NO_2 (tropospheric vertical column of NO_2) (*35*) and CO (vertically integrated CO column density) (*36*) are derived from the Sentinel-5P–gridded level 3 product at a 1.11-km resolution. AOD is collected from the MODIS multi-angle implementation of atmospheric correction land-gridded level 2 product [(0.55 μm), produced daily at a 1-km resolution (*34*)]. Meteorological input features comprise the daily mean of temperature and dewpoint temperature at 2 m, wind speed in the eastward and northward directions, total precipitation, net thermal radiation at the surface, and surface pressure. These input features are drawn from the daily aggregate of European Centre for Medium-Range Weather Forecasts (ECMWF) Reanalysis 5th Generation Land (ERA5-Land) hourly assets at an 11.13-km resolution (*65*). In addition, the daily mean of wind direction and relative humidity per 10-km grid is calculated using wind speed in the eastward and northward directions, temperature, and dewpoint temperature. Furthermore, we incorporate MODIS land cover type data produced yearly at a 500-m resolution (*66*). In addition to MODIS land cover data, low and high vegetation indices obtained from ECMWF ERA5-Land are included in the model. Elevation data are sourced from the Shuttle Radar Topography Mission at a resolution of approximately 30 m (*67*).

To account for missing observations in TROPOMI features and AOD, atmospheric reanalysis data such as the Modern-Era Retrospective Analysis for Research and Applications, version 2 (MERRA-2) aerosol optical thickness (AOT at 550 nm) (*68*) and CO (*69*), as well as Aura Ozone Monitoring Instrument (OMI) NO_2 (*70*), are included. Those missing data in TROPOMI features and AOD are imputed using machine learning methods before being incorporated into the main machine learning model, which predicts $PM_{2.5}$ concentrations. Furthermore, the missingness of TROPOMI NO_2, CO, and AOD is computed on the basis of the amount of missing observations for each day, and this information is included as model

input. Last, month and day of the year, a dummy variable indicating a monsoon season, and centroids of each 10-km grid are incorporated as model input features. Weekly rolling averages of TROPOMI NO_2 and CO, AOD, MERRA-2 variables, as well as meteorological variables are also calculated and included in the model to capture the potential time dependency between input features and $PM_{2.5}$ observations to improve the model predictability.

Imputation of TROPOMI and AOD

For the Full model, 26.0% of TROPOMI NO_2, 18.7% of TROPOMI CO, and 49.0% of AOD had missing observations from 10 July 2018 to 30 September 2023 across the country (fig. S26). In the AOD model, 41.8% of AOD data were missing from 1 January 2015 to 30 September 2023 (fig. S27). These missing observations are predicted using light gradient-boosting machine (LightGBM) and XGBoost with input features, including MERRA-2 AOT and CO, OMI NO_2, meteorology, land cover, elevation, month and day of the year, a dummy variable indicating a monsoon season, centroids of each 10-km grid, and weekly rolling averages and annual averages of MERRA-2 variables, as well as meteorological variables (table S2). We conduct pairwise correlation analysis for feature selection to ensure no variable is highly correlated ($R^2 > 0.9$) with each other, resulting in the removal of redundant features and increased efficiency in learning tasks (*71*).

Models are trained on TROPOMI NO_2, TROPOMI CO, and AOD observations for each machine learning task to predict and impute missing data. For each model, approximately 1 million observations were randomly selected as training data (2% for models predicting TROPOMI NO_2 and TROPOMI CO, and 3% for the model predicting AOD) based on 50-km grid cells, month of the year, and year (table S5). Stratified random sampling was used to ensure a representative dataset across different locations and times, preventing potential biases. For model selection and hyperparameter tuning, we conduct fivefold inner spatial CV using training data in one of the 10-fold outer spatial CV to prevent any final test data from leaking into training tasks during model selection and hyperparameter tuning (fig. S28). Both the inner and outer spatial CV are conducted on the basis of 50-km grid to account for spatial autocorrection within the input features, especially MERRA-2 reanalysis data, which has the coarsest spatial resolution of approximately 50 km (0.5° latitude × 0.625° longitude). Splitting data into training and test sets based on the 50-km grid, rather than the more conventional method of random splitting by observation is a more demanding prediction task because a given monitor can contribute data to both training and test sets in case of random split. Using coarser spatial blocks, rather than 10-km grid, further increases the difficulty of such a task; however, spatial CV is a more realistic test of how well the model would perform in predicting time series of missing observations in a new location with no training data. We fit LightGBM and XGBoost for each model predicting TROPOMI NO_2, TROPOMI CO, and AOD while implementing GridSearchCV to search over the LightGBM and XGBoost hyperparameter ranges (table S3) and identify the optimal combination of hyperparameters. As a result, LightGBM was used for predicting missing TROPOMI NO_2 and TROPOMI CO, and XGBoost was used for predicting missing AOD (table S4). The final predictions derived from held-out test data in each of the 10 folds are compared with observations using the evaluation metrics, such as overall R^2, within R^2, and RMSE. The within R^2 is calculated by regressing observed TROPOMI NO_2, TROPOMI CO, and AOD on their respective predicted values while incorporating fixed effects for locations with observations, month of the year, and year. Specifically, the within R^2 can heuristically be written as

$$R^2_{\text{within}} = \text{Corr}^2\left[\widehat{y_{it}} - \widehat{y_i}, y_{it} - \overline{y_i}\right]$$

where $\widehat{y_{it}}$ denotes the fitted value for individual i at time t, $\widehat{y_i}$ represents the fitted value for individual i across all time periods, y_{it} denotes the actual observed value for individual i at time t, and $\overline{y_i}$ represents the mean outcome value for individual i over the observed time periods.

In practice, we estimate the within R^2 by also accounting for seasonal variation in $PM_{2.5}$ in a given location and common variation over time in $PM_{2.5}$. This amounts to also subtracting off year-specific $PM_{2.5}$ averages and seasonal averages. The goal is to isolate within-location temporal variation in $PM_{2.5}$, after accounting for average seasonal differences in $PM_{2.5}$ and common differences across years (e.g., due to long-term trends in $PM_{2.5}$ or shocks due to COVID).

For the Full model, the predictions of missing TROPOMI NO_2 explained 52%, predictions of missing TROPOMI CO explained 92%, and predictions of missing AOD explained 82% of out-of-sample variation (table S5). For the AOD model, the predictions of AOD explained 82% of out-of-sample variation (table S5). We then predicted TROPOMI NO_2, TROPOMI CO, and AOD values for all 10-km grid cells over India and used these imputed variables as input features for the main machine learning model predicting $PM_{2.5}$ concentrations. Binary variables indicating whether each of TROPOMI NO_2, TOPOMI CO, and AOD is imputed (0 or 1) is also included. Last, we create additional weighted variables by assigning weights to the imputed TROPOMI NO_2, TROPOMI CO, and AOD values based on their respective imputation performances (1.0 for actual observations, 0.5 for imputed TROPOMI NO_2, 0.9 for imputed TROPOMI CO, and 0.8 for imputed AOD). The weighted variables are incorporated into the main machine learning model.

Model tuning and validation

The Full model is trained and validated on 301,355 daily ground-measured $PM_{2.5}$ concentrations from 10 July 2018 to 30 September 2023. The AOD model is trained on 345,559 $PM_{2.5}$ observations from 1 January 2015 to 30 September 2023. Similar to our models for the imputation of TROPOMI features and AOD, we conduct pairwise correlation analysis for feature selection to ensure no variable is highly correlated ($R^2 > 0.9$) with each other. For model selection and hyperparameter tuning, we conduct fivefold inner spatial CV using training data in one of the 10-fold outer spatial CV. Both the inner and outer spatial CV are constructed on the basis of 50-km grid. When splitting data into training and test sets through inner and outer spatial CV, we ensure that each test fold includes a nearly equal number of 50-km grid cells from each of three environmental regions (fig. S5), identified based on k-means clustering using imputed TROPOMI NO_2 and CO, imputed AOD, MERRA-2 AOT and CO, and OMI NO_2. This helps balance environmental characteristics, such as urban versus rural, within training and test data in each fold. For the inner fivefold CV, each 50-km block of 10-km grid cells goes into only one of the five test folds, and for the outer 10-fold CV, each 50-km block goes into only one of the 10 test folds. We fit LightGBM and XGBoost using the inner CV while implementing GridSearchCV to search over the hyperparameter ranges and identify the optimal combination of hyperparameters (table S6). On the

basis of the model performances using the inner CV, we selected XGBoost for both Full and AOD models (table S7). We apply XGBoost to the outer 10-fold spatial CV using the RMSE as the objective function.

We measure model performance by comparing observed $PM_{2.5}$ with model predictions derived from held-out test data in each of the 10 folds. For the model evaluation, we calculate the overall R^2, within R^2, and RMSE on the held-out test set for each of the 10 folds. Similar to the models for imputing missing observations in TROPOMI NO_2, TROPOMI CO, and AOD, the within R^2 is calculated by regressing observed $PM_{2.5}$ concentrations on predicted values while incorporating fixed effects for locations with observations, month of the year, and year.

Sensitivity analysis

As part of the sensitivity analysis, we aim to assess the robustness of the Full model by exposing it to a more spatially challenging task through the implementation of a 10-fold spatial CV with larger test blocks based on latitude (fig. S13). In addition, we conduct another iteration of the Full model using a 10-fold random CV where data are randomly split into 10 folds without specific consideration for spatial distribution. This approach enables us to assess and confirm the potential underestimation of spatial prediction errors and the optimistic overall results associated with random *k*-fold CV, which does not account for spatial auto-correction within spatial data.

Predictive power of AOD and TROPOMI features

To evaluate and identify the predictive power of TROPOMI features, we construct another model, the TROPOMI model, which excludes AOD but incorporates other features present in the Full model. We compare the performance of the Full, AOD, and TROPOMI models by fitting XGBoost using the same training and test sets from 10 July 2018 to 30 September 2023. Evaluation metrics include the overall R^2, within R^2, and RMSE.

Assessing long-term spatiotemporal trends in predicted $PM_{2.5}$ concentrations

Using the $PM_{2.5}$ predictions derived from the AOD model, we calculate the average $PM_{2.5}$ concentrations per 10-km grid between 2005 and 2010 and changes in average concentrations from 2005–2010 to 2011–2016 and to 2017–2022. Using 6-year windows helps control for meteorological variability between years and mitigates undue influence from extreme years, such as 2020 when India experienced a nationwide lockdown, similar to other countries (*50*). We identified that spatial discontinuities in the $PM_{2.5}$ predictions seen in the maps of Fig. 2, particularly north-south streaks across India, are driven by the coarse spatial resolution of MERRA-2 inputs, specifically AOT and CO. While attempts to smooth the predictions were made by taking the average predictions in the test fold within 25-, 50-, and 100-km radius, respectively, and by evaluating the smoothed predictions using CV, we chose to retain the unsmoothed predictions as they better capture substantial local variations in $PM_{2.5}$ concentrations and outperformed smoothed predictions. In addition, population-weighted annual average $PM_{2.5}$ concentrations are computed for India as a whole and five mega-cities from 2005 to 2022, combining population counts (*72*) with annual average $PM_{2.5}$ concentrations per 10-km grid. To elucidate the long-term temporal trend in population-weighted averages, we also calculate the percentage changes in population-weighted annual average $PM_{2.5}$ concentrations for each 3-year window from 2006–2008 to 2020–2022 relative to the 3-year average from 2005–2007.

To evaluate the ability of our model to detect small changes in $PM_{2.5}$ concentrations, we generated simulated data that had the daily RMSE of our modeled data, 27.79 $\mu g/m^3$ at the daily level, but that had a "true" reduction over time in $PM_{2.5}$ ranging from 0.5 to 5 $\mu g/m^3$ and studied whether we were able to recover this reduction. To do this, we randomly generated "true" $PM_{2.5}$ observations measured by ground monitors for each 10 km grid cell, using the mean and standard deviation of the monitor data for that grid cell over a period of 10 years. When generating these observations, the difference between the average of the first 5 years and the last 5 years was, in separate simulations, set to values ranging from 0.5 to 5 $\mu g/m^3$, which we refer to as the "true change." We then generated $PM_{2.5}$ "predictions" adding a noise level corresponding to an RMSE of 27.79 $\mu g/m^3$ to the true data. Next, we calculated the difference between average $PM_{2.5}$ concentrations for the first 5 years and the last 5 years based on the simulated predictions and the simulated ground truth. We repeated these steps 1000 times and computed the mean and 95% confidence intervals of the changes between the first and last 5 years for both the simulated monitor data and predictions.

Identifying the contribution of meteorological variability to long-term trends in predicted $PM_{2.5}$ concentrations

To identify the contribution of meteorological variability to the observed declining trend in $PM_{2.5}$ concentrations since 2016–2018, we model the $PM_{2.5}$ concentrations of each individual 10-km grid cell using an additive form of a trend component, a meteorology component, and time fixed effects (*43*). More specifically, we employ the following regression equation for each grid cell i

$$y_{it} = \beta_i^{obs} \times t + f_i(X_{it}) + \eta_{it} + \varepsilon_{it}$$

where y_{it} denotes the daily $PM_{2.5}$ concentration at grid cell i on day t. t is the time index (e.g., $t = 1$ for 1 January 2005, $t = 2$ for 2 January 2005, and $t = 32$ for 1 February 2005). $f_i(X_{it})$ represents a linear function that captures the effect of local meteorology on $PM_{2.5}$ concentration. X_{it} denotes the local meteorology variables in grid cell i on day t, including temperature and dewpoint temperature at 2 m, wind speed in the eastward and northward directions, total precipitation, and surface pressure. η_{it} is the month-of-year and day-of-month fixed effects to capture daily and monthly variability in pollutant concentrations that are not related to the meteorological variability (e.g., seasonal cycle in $PM_{2.5}$). ε_{it} is the normally distributed error term. β^{obs}represents the meteorology-corrected daily trend in $PM_{2.5}$ concentration for grid cell i estimated with the standard ordinary least-squares method. We apply the above regression equation to $PM_{2.5}$ daily predictions derived from the AOD model for 2005–2015 and 2016–2022, respectively, and obtain the meteorology-corrected annual trends of $PM_{2.5}$ concentrations for each grid cell for the corresponding periods by multiplying the estimated $\beta^{obs}$by 365 days.

In contrast, the observed annual trend is estimated by the following regression equation for each grid cell i

$$y_{it} = \beta_i^{obs} \times t + \eta_{it} + \varepsilon_{it}$$

where y_{it} denotes the daily $PM_{2.5}$ concentration at grid cell i on day t. t is the time index and η_{it} is the month-of-year and day-of-month

fixed effects. ε_{it} is the normally distributed error term. Using the β^{obs}, representing the daily trend in predicted $PM_{2.5}$ concentrations, annual trends for 2005–2015 and 2016–2022 are estimated.

Evaluation of NCAP

Our evaluation of NCAP is based on a matched difference-in-differences analysis, where we first identify a set of "control" urban areas that were not initially targeted by NCAP that look similar across a range of covariates to the treated cities that were targeted by the program, and then we compare trends in ambient $PM_{2.5}$ concentrations between treated and control areas, before and after program implementation. We implement our analysis at the subdistrict level due to the limited availability of city-level shapefiles in India. Specifically, we first identify the treatment subdistricts corresponding to the 102 non-attainment cities whose city clean action plans were approved in 2020. These non-attainment cities were selected under the NCAP based on air pollution levels of key pollutants, including $PM_{2.5}$, PM_{10}, and NO_2 (*7*). Next, we calculate the propensity score—the probability of becoming a city targeted by NCAP—based on factors that likely influenced a city's selection as a target under the NCAP, such as population count, proportion of urban area, proportion of forest area, and total area per subdistrict using logistic regression. Population count is crucial because higher populations are typically associated with increased emissions from sources such as transportation, industry, and domestic activities. The proportion of urban area is another critical factor, reflecting the higher pollution levels commonly found in densely urbanized areas due to the concentration of vehicles, industries, and construction activities. To account for natural mitigation factors, we also include the proportion of forest area, which plays a significant role in reducing pollution through carbon sequestration. The total area of each subdistrict is important because subdistricts vary significantly in size, and including this variable allows us to compare the effects of the NCAP on air quality improvements across areas of similar size, ensuring a more equitable assessment of the program's impact. We apply this to non-treated subdistricts that are not adjacent to the treatment subdistricts and are located outside a 50-km radius from them to account for potential spillover effects of NCAP on neighboring subdistricts. Last, we identify control subdistricts as those with propensity closest to those of the treatment subdistricts. Propensity score methods are commonly used to minimize selection bias in identifying control groups that share observed baseline characteristics with treatment groups as similar as possible (*73*). Consequently, our analysis includes 88 treatment subdistricts and 74 control subdistricts (fig. S20).

To establish the causal effect of NCAP, the trends in $PM_{2.5}$ concentrations in subdistricts that do not include NCAP's non-attainment cities must provide valid counterfactuals for the trends that we would have observed in the treatment subdistricts. We evaluate whether the parallel trends assumption may be reasonable in our case by plotting yearly average treatment effects on $PM_{2.5}$ concentrations before 2020, confirming that there was no systematic difference in treatment and control groups before city action plans were approved by the CPCB in 2020. More specifically, we use the following regression equation to estimate the average treatment effects for each year from 2005 to 2022

$$y_{it} = \sum_{j=t-15}^{t+2} \beta_j D_i{}^{*} 1\left(year = \mathrm{j}\right) + \alpha_i + \gamma_t + \theta x_{it} + \varepsilon_{it}$$

where y_{it} denotes the daily $PM_{2.5}$ concentration at subdistrict i on day t and the β_j coefficients estimate the difference between treated and control subdistricts in the 15 years prior to treatment ($t = 0$, which is 2020 in our data, the year in which city clean action plans were approved by the CPCB) and the two observed years following treatment (2021 and 2022). The estimate in $t = 0$ is omitted to avoid collinearity, and so the β_j coefficients are interpreted as the difference between treated and control concentrations relative to the reference year of 2020. α_i represents a vector of subdistrict fixed effects, and γ_t denotes vectors of month-of-year and year-fixed effects. θx_{it} denotes a vector of additional time-varying controls at the subdistrict level, including temperature and dewpoint temperature at 2 m, wind speed in the eastward and northward directions, total precipitation, and surface pressure. ε_{it} is the error term. Because our chosen treated units all adopted city clean action plans in the same year, our treatment is not staggered across time, and thus, our analysis avoids the inference issues that accompany difference-in-differences designs with staggered adoption that have been highlighted in recent literature.

Investigating population exposure to $PM_{2.5}$ concentrations

To identify locations with elevated levels of daily $PM_{2.5}$ concentrations, we use predictions from the AOD model to calculate the average number of days each 10-km grid cell exceeding the WHO daily guideline of 15 $\mu g/m^3$, national daily guideline of 60 $\mu g/m^3$, and extreme value of 100 $\mu g/m^3$ for 2018–2022.

Furthermore, we examine proportional changes in populations exposed to high levels of $PM_{2.5}$ annual average concentrations from 2005 to 2022. For this analysis, we first calculate annual average $PM_{2.5}$ concentrations per 10-km grid. Subsequently, we identify the 10-km grids above the thresholds of the WHO annual guideline (5 $\mu g/m^3$), national annual guideline (40 $\mu g/m^3$), and extreme annual concentration (80 $\mu g/m^3$). We then aggregate the population counts residing in those grid cells exceeding each threshold to calculate the percentage of whole population.

Examining inequalities in $PM_{2.5}$ exposure

We combine estimates of relative wealth [Relative Wealth Indices (RWIs)] at a 2.4-km resolution (*54*) with our $PM_{2.5}$ predictions derived from the AOD model by averaging the wealth estimates for each 10-km grid. RWIs were generated using machine learning algorithms that used satellite data, mobile phone data, and data from Facebook users (*54*). For India, their machine learning algorithms were trained on ground measurements of wealth collected by DHS surveys in 2015 and 2019 (*54*). The gridded wealth data provide estimates for each grid cell of how wealthy that grid cell is relative to others in the same country and an estimated error for that estimate (*3*). We further combine population counts (*72*) for each grid, along with RWIs and daily $PM_{2.5}$ predictions from January 2005 to December 2022.

To understand the proportion of population exposed to elevated levels of $PM_{2.5}$ concentrations across wealth groups, we aggregate the population counts for 10-km grid cells whose average $PM_{2.5}$ concentrations for 2018–2022 are exceeding the national guideline of 40 $\mu g/m^3$ and extreme threshold of 80 $\mu g/m^3$ annually, respectively. We then calculate the percentage of population for each of the five wealth categories (poorest, poorer, middle, richer, and richest), which are created on the basis of quintiles of RWIs within the country. In addition, to understand the temporal changes in wealth disparities

in $PM_{2.5}$ exposure from 2005 to 2022, population-weighted annual average $PM_{2.5}$ concentrations are computed for all the wealth levels across the country, grid cells at 90th percentile of RWIs, and grid cells at 10th percentile of RWIs. To account for uncertainty in the estimates, 1000 bootstrap simulations were conducted for each of these cases. To elucidate the long-term temporal trend in population-weighted averages, we also calculate the percentage changes in population-weighted annual average $PM_{2.5}$ concentrations for each 3-year window from 2006–2008 to 2020–2022 relative to the 3-year average from 2005–2007.

To further understand disparities in urban and rural areas, we stratified the country into urban and rural regions using the URCA data (*55*), which provides pixel-level estimates of the travel time to the nearest urban area based on data collected in 2015. Pixels within urban areas were assigned values indicating their size, ranging from 1 (representing a city with at least 5 million people) to 30 (representing a rural hinterland). We defined urban areas as those grid cells with URCA values less than 10 and rural areas as those with values greater than or equal to 10. The same calculations for population exposure to elevated $PM_{2.5}$ levels were performed for both urban and rural areas, stratified by wealth index.

Assessing the optimal placement of air quality monitors ensuring equality

To calculate where additional ground monitors could be placed to optimally enhance the ability of the entire ground-monitoring network to capture spatial and temporal variation in ambient $PM_{2.5}$ concentrations, we use multiresolution dynamic mode decomposition (mrDMD), which recursively decomposes a dataset into low-rank spatial modes and their temporal Fourier dynamics. mrDMD has been shown to capture $PM_{2.5}$ concentrations spatially and temporally on short- (daily) and long-term (years to decade) timescales and to incorporate information from transient phenomena, such as wildfires and temperature inversions, which would otherwise be discarded or averaged out using similar data reduction techniques (*37*). The algorithm can thus capture a finer level of spatial and temporal variability in a dataset that would otherwise be averaged out using traditional mean or maximum $PM_{2.5}$ metrics. mrDMD is a dimensionality reduction algorithm, similar to principal components analysis (PCA), but mrDMD is more precise in capturing spatiotemporal variability than methods based on singular value decomposition such as PCA.

Here, we present an extension to the mrDMD framework that considers cost-constraining functions to optimize sensor placement based on wealth estimates. We excluded grid cells with no populations from wealth estimates to ensure that populous poor communities are prioritized. The mrDMD algorithm is based on a column-pivoted QR algorithm, where the pivot column is modified to balance (i) the decrease in accuracy of capturing the largest air pollution modal signals with (ii) the increase in capturing pollution exposure in grid cells with low RWI values, representing poorer communities. The cost function used here is a step function that penalizes placing sensors too far from poorer locations. For the cost function vector, all data are a binary of 0 or 1, with 0 representing wealthy locations and 1 representing poor locations. More details of the method can be found in (*56*).

We apply all mrDMD methods to the $PM_{2.5}$ estimates derived from the AOD model from 1 January 2005 to 30 September 2023. The resulting sensor network adds additional monitors to the existing ground network to better constrain variation in surface pollution, and we select a number of additional monitors such that the entire monitoring network would have 1000 ground monitors, the NCAP's target.

Examining applicability of our model to other low- and middle-income settings

To investigate the association between the quantity of ground truth data, represented by the number of 10-km grid cells with measured $PM_{2.5}$ observations, and the predictive performance of our machine learning model for $PM_{2.5}$ concentrations, we conduct an experiment involving multiple machine learning models trained with incremental training data. First, 21 of 321 grid cells with $PM_{2.5}$ observations across the country are randomly selected and held out until the final evaluation of predictions. Second, from the remaining 300 grid cells, we randomly selected grid cells, starting with 5 and 15, and then in increments of 25, from 25 up to 300, without replacement, to serve as training data for our Full model. This process results in the creation of 13 Full models with training data of different sizes. Each training dataset is then split into training and validation sets based on 50-km grid cells, creating a 10-fold spatial CV. For each model, XGBoost is applied with hyperparameter tuning, where the parameters include learning_rate (set to 0.01), lambda (1 or 10), max_depth (fixed at 10), and n_estimators (1000 or 1500), while keeping the remaining hyperparameters at their default values. The tuning process is performed using a 10-fold spatial CV. Last, XGBoost with the best hyperparameters is applied to predict $PM_{2.5}$ concentrations for the initially randomly selected 21 grid cells, and the final predictions are compared and evaluated with the test data. This entire process, from the random sampling without replacement of the 21 grid cells for the test dataset to making final predictions using the best hyperparameters for the 13 Full models, is repeated 1000 times to obtain the best possible unbiased estimates of prediction performances. The mean within R^2 is calculated of 1000 samples for each of the 13 Full models.

Supplementary Materials

This PDF file includes:

Figs. S1 to S28
Tables S1 to S7

REFERENCES AND NOTES

1. R. Burnett, H. Chen, M. Szyszkowicz, N. Fann, B. Hubbell, C. A. Pope, J. S. Apte, M. Brauer, A. Cohen, S. Weichenthal, J. Coggins, Q. Di, B. Brunekreef, J. Frostad, S. S. Lim, H. Kan, K. D. Walker, G. D. Thurston, R. B. Hayes, C. C. Lim, M. C. Turner, M. Jerrett, D. Krewski, S. M. Gapstur, W. R. Diver, B. Ostro, D. Goldberg, D. L. Crouse, R. V. Martin, P. Peters, L. Pinault, M. Tjepkema, A. van Donkelaar, P. J. Villeneuve, A. B. Miller, P. Yin, M. Zhou, L. Wang, N. A. H. Janssen, M. Marra, R. W. Atkinson, H. Tsang, T. Quoc Thach, J. B. Cannon, R. T. Allen, J. E. Hart, F. Laden, G. Cesaroni, F. Forastiere, G. Weinmayr, A. Jaensch, G. Nagel, H. Concin, J. V. Spadaro, Global estimates of mortality associated with long-term exposure to outdoor fine particulate matter. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9592–9597 (2018).
2. A. Hajat, C. Hsia, M. S. O'Neill, Socioeconomic disparities and air pollution exposure: A global review. *Curr. Environ. Health Rep.* **2**, 440–450 (2015).
3. A. P. Behrer, S. Heft-Neal, Higher air pollution in wealthy districts of most low- and middle-income countries. *Nat. Sustain.* **7**, 203–212 (2024).
4. J. Liu, L. P. Clark, M. J. Bechle, A. Hajat, S.-Y. Kim, A. L. Robinson, L. Sheppard, A. A. Szpiro, J. D. Marshall, Disparities in air pollution exposure in the United States by race/ethnicity and income, 1990–2010. *Environ. Health Perspect.* **129**, 127005 (2021).
5. P. Pant, R. M. Lal, S. K. Guttikunda, A. G. Russell, A. S. Nagpure, A. Ramaswami, R. E. Peltier, Monitoring particulate matter in India: Recent trends and future outlook. *Air Qual. Atmos. Health* **12**, 45–58 (2019).
6. A. Roychowdhury, A. Somvanshi, Breathing Space: How to track and report air pollution under the National Clean Air Programme, Center for Science and Environment (Centre for Science and Environment, 2020); www.cseindia.org/content/downloadreports/9923.

7. S. Guttikunda, N. Ka, T. Ganguly, P. Jawahar, Plugging the ambient air monitoring gaps in India's national clean air programme (NCAP) airsheds. *Atmos. Environ.* **301**, 119712 (2023).
8. S. Gulia, I. Khanna, K. Shukla, M. Khare, Ambient air pollutant monitoring and analysis protocol for low and middle income countries: An element of comprehensive urban air quality management framework. *Atmos. Environ.* **222**, 117120 (2020).
9. European Environment Agency, European Air Quality Index; https://airindex.eea.europa.eu/AQI/index.html.
10. G. Geng, Q. Xiao, S. Liu, X. Liu, J. Cheng, Y. Zheng, T. Xue, D. Tong, B. Zheng, Y. Peng, X. Huang, K. He, Q. Zhang, Tracking air pollution in China: Near real-time $PM_{2.5}$ retrievals from multisource data fusion. *Environ. Sci. Technol.* **55**, 12106–12115 (2021).
11. T. Ganguly, K. L. Selvaraj, S. K. Guttikunda, National Clean Air Programme (NCAP) for Indian cities: Review and outlook of clean air action plans. *Atmos. Environ. X* **8**, 100096 (2020).
12. P. N. deSouza, E. Chaudhary, S. Dey, S. Ko, J. Németh, S. Guttikunda, S. Chowdhury, P. Kinney, S. V. Subramanian, M. L. Bell, R. Kim, An environmental justice analysis of air pollution in India. *Sci. Rep.* **13**, 16690 (2023).
13. J. Chakraborty, P. Basu, Air quality and environmental injustice in India: Connecting particulate pollution to social disadvantages. *Int. J. Environ. Res. Public Health* **18**, 304 (2021).
14. A. van Donkelaar, M. S. Hammer, L. Bindle, M. Brauer, J. R. Brook, M. J. Garay, N. C. Hsu, O. V. Kalashnikova, R. A. Kahn, C. Lee, R. C. Levy, A. Lyapustin, A. M. Sayer, R. V. Martin, Monthly global estimates of fine particulate matter and their uncertainty. *Environ. Sci. Technol.* **55**, 15287–15300 (2021).
15. Q. Di, H. Amini, L. Shi, I. Kloog, R. Silvern, J. Kelly, M. B. Sabath, C. Choirat, P. Koutrakis, A. Lyapustin, Y. Wang, L. J. Mickley, J. Schwartz, An ensemble-based model of PM2.5 concentration across the contiguous United States with high spatiotemporal resolution. *Environ. Int.* **130**, 104909 (2019).
16. R. Rajak, A. Chattopadhyay, Short and long term exposure to ambient air pollution and impact on health in India: A systematic review. *Int. J. Environ. Health Res.* **30**, 593–617 (2020).
17. M. L. Bell, A. Zanobetti, F. Dominici, Evidence on vulnerability and susceptibility to health risks associated with short-term exposure to particulate matter: A systematic review and meta-analysis. *Am. J. Epidemiol.* **178**, 865–876 (2013).
18. Y. Wang, M. N. Eliot, G. A. Wellenius, Short-term changes in ambient particulate matter and risk of stroke: A systematic review and meta-analysis. *J. Am. Heart Assoc.* **3**, e000983 (2014).
19. H. Luo, Q. Zhang, Y. Niu, H. Kan, R. Chen, Fine particulate matter and cardiorespiratory health in China: A systematic review and meta-analysis of epidemiological studies. *J. Environ. Sci. (China)* **123**, 306–316 (2023).
20. M. Kowalska, K. Kocot, Short-term exposure to ambient fine particulate matter (PM2,5 and PM10) and the risk of heart rhythm abnormalities and stroke. *Postepy Hig. Med. Dosw. (Online)* **70**, 1017–1025 (2016).
21. R. Liang, B. Zhang, X. Zhao, Y. Ruan, H. Lian, Z. Fan, Effect of exposure to PM2.5 on blood pressure: A systematic review and meta-analysis. *J. Hypertens.* **32**, 2130–2141 (2014).
22. M.-H. Li, L.-C. Fan, B. Mao, J.-W. Yang, A. M. K. Choi, W.-J. Cao, J.-F. Xu, Short-term exposure to ambient fine particulate matter increases hospitalizations and mortality in COPD: A systematic review and meta-analysis. *Chest* **149**, 447–458 (2016).
23. J. Fan, S. Li, C. Fan, Z. Bai, K. Yang, The impact of PM2.5 on asthma emergency department visits: A systematic review and meta-analysis. *Environ. Sci. Pollut. Res.* **23**, 843–850 (2016).
24. V. Katoch, A. Kumar, F. Imam, D. Sarkar, L. D. Knibbs, Y. Liu, D. Ganguly, S. Dey, Addressing biases in ambient $PM_{2.5}$ exposure and associated health burden estimates by filling satellite AOD retrieval gaps over India. *Environ. Sci. Technol.* **57**, 19190–19201 (2023).
25. S. Dey, B. Purohit, P. Balyan, K. Dixit, K. Bali, A. Kumar, F. Imam, S. Chowdhury, D. Ganguly, P. Gargava, V. K. Shukla, A satellite-based high-resolution (1-km) ambient $PM_{2.5}$ database for India over two decades (2000–2019): Applications for air quality management. *Remote Sens.* **12**, 3872 (2020).
26. S. Mandal, A. Rajiva, I. Kloog, J. S. Menon, K. J. Lane, H. Amini, G. K. Walia, S. Dixit, A. Nori-Sarma, A. Dutta, P. Sharma, S. Jaganathan, K. K. Madhipatla, G. A. Wellenius, J. de Bont, C. Venkataraman, D. Prabhakaran, P. Prabhakaran, P. Ljungman, J. Schwartz, Nationwide estimation of daily ambient PM2.5 from 2008 to 2020 at 1 km2 in India using an ensemble approach. *PNAS Nexus* **3**, pgae088 (2024).
27. Z. Ma, S. Dey, S. Christopher, R. Liu, J. Bi, P. Balyan, Y. Liu, A review of statistical methods used for developing large-scale and long-term PM2.5 models from satellite data. *Remote Sens. Environ.* **269**, 112827 (2022).
28. Y. Rybarczyk, R. Zalakeviciute, Machine learning approaches for outdoor air quality modelling: A systematic review. *Appl. Sci.* **8**, 2570 (2018).
29. R. B. A. Koelemeijer, C. D. Homan, J. Matthijsen, Comparison of spatial and temporal variations of aerosol optical thickness and particulate matter over Europe. *Atmos. Environ.* **40**, 5304–5315 (2006).
30. R. Son, D. Stratoulias, H. C. Kim, J.-H. Yoon, Estimation of surface $PM_{2.5}$ concentrations from atmospheric gas species retrieved from TROPOMI using deep learning: Impacts of fire on air pollution over Thailand. *Atmos. Pollut. Res.* **14**, 101875 (2023).
31. M. L. Childs, J. Li, J. Wen, S. Heft-Neal, A. Driscoll, S. Wang, C. F. Gould, M. Qiu, J. Burney, M. Burke, Daily local-level estimates of ambient wildfire smoke $PM_{2.5}$ for the contiguous US. *Environ. Sci. Technol.* **56**, 13607–13621 (2022).
32. L. Mamić, M. Gašparović, G. Kaplan, Developing PM2.5 and PM10 prediction models on a national and regional scale using open-source remote sensing data. *Environ. Monit. Assess.* **195**, 644 (2023).
33. Y. Wang, Q. Yuan, T. Li, S. Tan, L. Zhang, Full-coverage spatiotemporal mapping of ambient PM2.5 and PM10 over China from Sentinel-5P and assimilated datasets: Considering the precursors and chemical compositions. *Sci. Total Environ.* **793**, 148535 (2021).
34. A. Lyapustin, Y. Wang, MODIS/Terra+Aqua Land Aerosol Optical Depth Daily L2G Global 1km SIN Grid V061 (NASA EOSDIS Land Processes Distributed Active Archive Center, 2022). https://doi.org/10.5067/MODIS/MCD19A2.061.
35. Copernicus Sentinel-5P (processed by ESA), TROPOMI Level 2 Nitrogen Dioxide total column products. version 02 (European Space Agency, 2021). https://doi.org/10.5270/S5P-9bnp8q8.
36. Copernicus Sentinel-5P (processed by ESA), TROPOMI Level 2 Carbon Monoxide total column products. version 02 (European Space Agency, 2021). https://doi.org/10.5270/S5P-bj3nry0.
37. M. M. Kelp, S. Lin, J. N. Kutz, L. J. Mickley, A new approach for determining optimal placement of $PM_{2.5}$ air quality sensors: Case study for the contiguous United States. *Environ. Res. Lett.* **17**, 034034 (2022).
38. United Nations, Department of Economic and Social Affairs, Population Division, World Population Prospects 2022, Data Sources (UN DESA/POP/2022/DC/NO. 9, United Nations, 2023).
39. R. Jat, B. R. Gurjar, D. Lowe, Regional pollution loading in winter months over India using high resolution WRF-Chem simulation. *Atmos. Res.* **249**, 105326 (2021).
40. B. M. Skinder, A. Q. Sheikh, A. K. Pandit, B. A. Ganai, Brick kiln emissions and its environmental impact: A Review. *JENE* **6**, 1–11 (2014).
41. D. Chatterjee, E. E. McDuffie, S. J. Smith, L. Bindle, A. van Donkelaar, M. S. Hammer, C. Venkataraman, M. Brauer, R. V. Martin, Source contributions to fine particulate matter and attributable mortality in India and the surrounding region. *Environ. Sci. Technol.* **57**, 10263–10275 (2023).
42. M. Kutlar Joss, M. Eeftens, E. Gintowt, R. Kappeler, N. Künzli, Time to harmonize national ambient air quality standards. *Int. J. Public Health* **62**, 453–462 (2017).
43. M. Qiu, C. Zigler, N. E. Selin, Statistical and machine learning methods for evaluating trends in air quality under changing meteorological conditions. *Atmos. Chem. Phys.* **22**, 10551–10566 (2022).
44. A. P. K. Tai, L. J. Mickley, D. J. Jacob, Correlations between fine particulate matter (PM2.5) and meteorological variables in the United States: Implications for the sensitivity of PM2.5 to climate change. *Atmos. Environ.* **44**, 3976–3984 (2010).
45. J.-P. Jalkanen, L. Johansson, J. Kukkonen, A. Brink, J. Kalli, T. Stipa, Extension of an assessment model of ship traffic exhaust emissions for particulate matter and carbon monoxide. *Atmos. Chem. Phys.* **12**, 2641–2659 (2012).
46. L. Johansson, J.-P. Jalkanen, J. Kukkonen, Global assessment of shipping emissions in 2015 on a high spatial and temporal resolution. *Atmos. Environ.* **167**, 403–415 (2017).
47. R. Lathia, S. Dadhaniya, Policy norms and proposed ways to achieve goals of Indian vehicle emission program. *J. Clean. Prod.* **208**, 1339–1346 (2019).
48. India temporarily shuts five coal-fired power plants around New Delhi (CNBC, 2021); www.cnbc.com/2021/11/17/india-temporarily-shuts-five-coal-fired-power-plants-around-new-delhi.html.
49. S. Sharma, M. Zhang, Anshika, J. Gao, H. Zhang, S. H. Kota, Effect of restricted emissions during COVID-19 on air quality in India. *Sci. Total Environ.* **728**, 138878 (2020).
50. M. Marwah, P. K. Agrawala, COVID-19 lockdown and environmental pollution: An Indian multi-state investigation. *Environ. Monit. Assess.* **194**, 49 (2022).
51. S. K. Sahu, P. Mangaraj, G. Beig, B. Tyagi, S. Tikle, V. Vinoj, Establishing a link between fine particulate matter (PM2.5) zones and COVID -19 over India based on anthropogenic emission sources and air quality data. *Urban Clim.* **38**, 100883 (2021).
52. R. L. Verma, J. S. Kamyotra, Impacts of COVID-19 on air quality in India. *Aerosol Air Qual. Res.* **21**, 200482 (2021).
53. S. Gulia, N. Goyal, S. Mendiratta, T. Biswas, S. K. Goyal, R. Kumar, COVID 19 lockdown - air quality reflections in Indian cities. *Aerosol Air Qual. Res.* **21**, 200308 (2021).
54. G. Chi, H. Fang, S. Chatterjee, J. E. Blumenstock, Microestimates of wealth for all low- and middle-income countries. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2113658119 (2022).
55. A. Cattaneo, A. Nelson, T. McMenomy, Global mapping of urban-rural catchment areas reveals unequal access to services. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2011990118 (2021).
56. M. M. Kelp, T. C. Fargiano, S. Lin, T. Liu, J. R. Turner, J. N. Kutz, L. J. Mickley, Data-driven placement of $PM_{2.5}$ air quality sensors in the United States: An approach to target urban environmental injustice. *Geohealth* **7**, e2023GH000834 (2023).
57. L. Wallace, J. Bi, W. R. Ott, J. Sarnat, Y. Liu, Calibration of low-cost PurpleAir outdoor monitors using an improved method of calculating PM. *Atmos. Environ.* **256**, 118432 (2021).
58. K. K. Barkjohn, B. Gantt, A. L. Clements, Development and application of a United States-wide correction for $PM_{2.5}$ data collected with the PurpleAir sensor. *Atmos. Meas. Tech.* **14**, 4617–4637 (2021).

59. W. W. Delp, B. C. Singer, Wildfire smoke adjustment factors for low-cost and professional $PM_{2.5}$ monitors with optical sensors. *Sensors* **20**, 3683 (2020).
60. J. Bi, A. Wildani, H. H. Chang, Y. Liu, Incorporating low-cost sensor measurements into high-resolution $PM_{2.5}$ modeling at a large spatial scale. *Environ. Sci. Technol.* **54**, 2152–2162 (2020).
61. B. Feenstra, V. Papapostolou, S. Hasheminassab, H. Zhang, B. D. Boghossian, D. Cocker, A. Polidori, Performance evaluation of twelve low-cost PM2.5 sensors at an ambient air monitoring site. *Atmos. Environ.* **216**, 116946 (2019).
62. M. J. Campmier, J. Gingrich, S. Singh, N. Baig, S. Gani, A. Upadhya, P. Agrawal, M. Kushwaha, H. R. Mishra, A. Pillarisetti, S. Vakacherla, R. K. Pathak, J. S. Apte, Seasonally optimized calibrations improve low-cost sensor performance: Long-term field evaluation of PurpleAir sensors in urban and rural India. *Atmos. Meas. Tech.* **16**, 4357–4374 (2023).
63. I. Ialongo, H. Virta, H. Eskes, J. Hovila, J. Douros, Comparison of TROPOMI/Sentinel-5 Precursor NO_2 observations with ground-based measurements in Helsinki. *Atmos. Meas. Tech.* **13**, 205–218 (2020).
64. D. Sharma, D. Mauzerall, Analysis of air pollution data in India between 2015 and 2019. *Aerosol Air Qual. Res.* **22**, 210204 (2022).
65. J. Muñoz-Sabater, E. Dutra, A. Agustí-Panareda, C. Albergel, G. Arduini, G. Balsamo, S. Boussetta, M. Choulga, S. Harrigan, H. Hersbach, B. Martens, D. G. Miralles, M. Piles, N. J. Rodríguez-Fernández, E. Zsoter, C. Buontempo, J.-N. Thépaut, ERA5-Land: A state-of-the-art global reanalysis dataset for land applications. *Earth Syst. Sci. Data* **13**, 4349–4383 (2021).
66. M. Friedl, D. Sulla-Menashe, MODIS/Terra+Aqua Land Cover Type Yearly L3 Global 500m SIN Grid V061 (NASA EOSDIS Land Processes Distributed Active Archive Center, 2022). https://doi.org/10.5067/MODIS/MCD12Q1.061.
67. T. G. Farr, P. A. Rosen, E. Caro, R. Crippen, R. Duren, S. Hensley, M. Kobrick, M. Paller, E. Rodriguez, L. Roth, D. Seal, S. Shaffer, J. Shimada, J. Umland, M. Werner, M. Oskin, D. Burbank, D. Alsdorf, The shuttle radar topography mission. *Rev. Geophys.* **45**, RG2004 (2007).
68. Global Modeling and Assimilation Office (GMAO), MERRA-2 tavg1_2d_lnd_Nx: 2d, 1- Hourly, Time-Averaged, Single-Level, Assimilation, Land Surface Diagnostics V5.12.4. (Goddard Earth Sciences Data and Information Services Center, 2015).
69. Global Modeling and Assimilation Office (GMAO), MERRA-2 inst3_3d_chm_Nv: 3d, 3- Hourly, Instantaneous, Model-Level, Assimilation, Carbon Monoxide and Ozone Mixing Ratio V5.12.4. (GES DISC, 2015).
70. L. N. Lamsal, N. A. Krotkov, A. Vasilkov, S. Marchenko, W. Qin, E.-S. Yang, Z. Fasnacht, J. Joiner, S. Choi, D. Haffner, W. H. Swartz, B. Fisher, E. Bucsela, Ozone Monitoring Instrument (OMI) Aura nitrogen dioxide standard product version 4.0 with improved surface and cloud treatments. *Atmos. Meas. Tech.* **14**, 455–479 (2021).
71. L. Yu, H. Liu, Feature selection for high-dimensional data: A fast correlation-based filter solution, in *Proceedings of the Twentieth International Conference on International Conference on Machine Learning* (AAAI Press, 2003). pp. 856–863.
72. Center for International Earth Science Information Network (CIESIN)–Columbia University), Gridded Population of the World, Version 4 (GPWv4): Population Count [NASA Socioeconomic Data and Applications Center (SEDAC), 2016]. http://dx.doi.org/10.7927/H4X63JVC.
73. E. A. Stuart, H. A. Huskamp, K. Duckworth, J. Simmons, Z. Song, M. E. Chernew, C. L. Barry, Using propensity scores in difference-in-differences models to estimate the effects of a policy change. *Health Serv. Outcomes Res. Methodol.* **14**, 166–182 (2014).
74. World Health Organization, *WHO global air quality guidelines: Particulate matter (PM2.5 and PM10), ozone, nitrogen dioxide, sulfur dioxide and carbon monoxide* (WHO, 2021); https://iris.who.int/bitstream/handle/10665/345329/9789240034228-eng.pdf.

Acknowledgments: We thank S. Luby, E. Bendavid, J. Li, C. Callahan, S. Heft-Neal, I. Higuera Mendieta, other members of Stanford ECHOLab, and seminar participants at IIT-Bombay for helpful comments. Some of the computing for this project was performed on the Stanford Sherlock cluster, and we would like to thank Stanford University and the Stanford Research Computing Center for providing computational resources and support that contributed to these research results. **Funding:** M.Q. is supported by Planetary Health Fellowship at Stanford's Center for Innovation in Global Health. **Author contributions:** Conceptualization: A.K., M.K., M.Q., and M.B. Methodology: A.K., M.K., M.Q., K.S., and M.B. Investigation: A.K., S.D., and M.K. Resources: A.K. Data curation: A.K. and M.B. Validation: A.K., S.D., and M.K. Formal analysis: A.K., and M.K. Software: A.K. and M.K. Visualization: A.K. Funding acquisition: M.B. Project administration: A.K., M.K., and M.B. Supervision: M.B. Writing—original draft: A.K. Writing—review and editing: A.K., S.D., M.K., M.Q., K.S., E.C., I.A., and M.B. Methodology: A.K., M.K., M.Q., K.S., and M.B. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** Our daily PM2.5 datasets derived from the Full and AOD models, as well as the code and source data needed to replicate the results, have been deposited in Zenodo (https://doi.org/10.5281/zenodo.13694585) and are publicly available. All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials.

Submitted 29 April 2024
Accepted 24 December 2024
Published 24 January 2025
10.1126/sciadv.adq1071