ORBIT: A new, reduced-complexity air quality model for South Asia

Author: Sumil K. Thakrar[1,*]
[1]School of GeoSciences, University of Edinburgh, Edinburgh, Scotland, UK
*Corresponding Author: sumilthakrar@gmail.com or sumil.thakrar@ed.ac.uk

This manuscript is currently a non-peer reviewed preprint submitted to EarthArXiv. As such, this manuscript has not yet been formally accepted in a peer-reviewed journal. Subsequent versions of this manuscript may have slightly different content.

ORBIT: A new, reduced-complexity air quality model for South Asia

Sumil K. Thakrar
School of GeoSciences, University of Edinburgh

v.1.0.3 (October 2026)

Abstract

Exposure to outdoor air pollution is the most important environmental health risk worldwide, associated with over a million deaths each year in South Asia alone. To understand how to improve air quality, it is necessary to model how changes in emissions affect changes in pollutant concentrations and human health. However, air quality modeling typically demands substantial computation and expertise, limiting the capacity to inform policy. Reduced complexity air quality models (RCMs) have been developed, but for South Asia many models are unavailable or have poor performance. Here, I present ORBIT, a new reduced complexity model currently implemented for South Asia at $0.5° \times 0.625°$ resolution. ORBIT can natively report health damages, and has many features that distinguish it from other RCMs. Its predictions of annual $PM_{2.5}$ concentrations (R: 0.76, normalized mean error (NME): 21.6%, normalized mean bias (NMB): −4.2%) meet the goals for published benchmarks. I have used ORBIT to produce marginal health damage estimates per 1,000 kg pollution for each primary pollutant, month, and time of day, for each subdistrict in South Asia.

1 Introduction

Exposure to outdoor air pollution is the most important environmental health risk worldwide, associated with over a million deaths each year in South Asia alone (McDuffie et al., 2021). To understand how to improve air quality, it is necessary to model how changes in emissions affect changes in pollutant concentrations and human health. However, air quality modelling typically demands substantial computation and expertise, limiting the capacity to inform policy. Reduced complexity air quality models (RCMs) have been developed (Gilmore et al., 2019), but for South Asia many models are unavailable (Heo, Adams, & Gao, 2016; Muller, 2014) or have poor performance (Thakrar et al., 2022).

Here, I present ORBIT, a new reduced complexity model currently implemented for South Asia at $0.5° \times 0.625°$ resolution. Its predictions of annual $PM_{2.5}$ concentrations (R: 0.76, normalized mean error (NME): 21.6%, normalized mean bias (NMB): −4.2%) meet the goals for published benchmarks (Emery et al., 2017). I have used ORBIT to produce marginal health damage estimates per 1,000 kg pollution for each primary pollutant, month, and time of day, for each subdistrict in South Asia.

ORBIT can natively report health damages, and has many features that distinguish it from other RCMs:

- *Diurnal & seasonal results.* ORBIT resolves primary and secondary $PM_{2.5}$ for eight 3-hour bins each month, not just annual average.

- *Nonlinear chemistry.* ORBIT captures nonlinear regime changes, including saturation effects where emissions increases do not increase $PM_{2.5}$ concentrations.
- *Pollutant cross-couplings.* ORBIT resolves how one chemical species affects another (*e.g.*, $\downarrow SO_2 \rightarrow \uparrow pNO_3$).
- *Little chemical transport model (CTM) dependence.* Only oxidant and background aerosol fields are taken from a GEOS-Chem simulation, reducing the dependence on a specific, expensive simulation.
- *Fast.* A full South Asia simulation takes ~23 min per month (8 CPU cores, 3.6 GB), and months are parallelizable. Once the baseline is run, marginal simulations (for small emissions perturbations) take only ~5 min (5.5 GB).

2 Methods

2.1 ORBIT Solver

2.1.1 Continuity Equation

In a compressible atmosphere with air density ρ and wind velocity $\mathbf{u}$, the mass continuity equation for a trace pollutant with mass mixing ratio q is given by

$$\frac{\partial(\rho q)}{\partial t} + \nabla \cdot (\rho q\, \mathbf{u}) = \rho\,(P - Q) + E - D, \tag{1}$$

where P and Q are the chemical production and loss rates of the mixing ratio, and E and D are the emission and deposition rates per unit volume.

To match the 3-hourly fields of the input meteorology (Section 2.5.2), I discretize time into 3-hour intervals, denoted by k. Within each interval, the wind and process rates are held constant, and the covariance between their faster fluctuations and the mixing ratio is ignored.

ORBIT and its input meteorology represent the wind and the mixing ratio as averages over each grid cell, so each can be decomposed into its cell average and a fluctuation about it:

$$\mathbf{u} = \bar{\mathbf{u}} + \mathbf{u}', \qquad q = \bar{q} + q'. \tag{2}$$

The fluctuations carry pollution through the cell faces by eddies and convective updrafts smaller than a cell, so ORBIT must also model this sub-grid flux, $\mathbf{J} \equiv \rho\,\overline{\mathbf{u}'q'}$ (neglecting fluctuations of density within the cell). Averaging Equation 1 over the cell then gives

$$\frac{\partial(\rho \bar{q})}{\partial t} + \nabla \cdot (\rho \bar{q}\, \bar{\mathbf{u}}) = -\nabla \cdot \mathbf{J} + \rho\,(P - Q) + E - D. \tag{3}$$

From here on I drop the bars, so that q and $\mathbf{u}$ denote the cell averages $\bar{q}$ and $\bar{\mathbf{u}}$.

For each cell in the domain, I integrate Equation 3 over the volume V_j of cell j. If the cell boundaries are fixed in time, the time derivative can be moved outside the integral:

$$\int_{V_j} \frac{\partial(\rho q)}{\partial t}\, \mathrm{d}V = \frac{\mathrm{d}}{\mathrm{d}t} \int_{V_j} \rho q\, \mathrm{d}V = \frac{\mathrm{d}(m_j q_j)}{\mathrm{d}t}, \tag{4}$$

where $m_j = \int_{V_j} \rho \, \mathrm{d}V$ is the air mass of the cell and $q_j = m_j^{-1} \int_{V_j} \rho q \, \mathrm{d}V$ is its mass-weighted mean mixing ratio. On a hybrid pressure grid, the cell boundaries move only if the surface pressure changes, so within each interval (where the meteorology is constant) the boundaries are fixed.

The last expression simplifies to $m_j \, \mathrm{d}q_j/\mathrm{d}t$ if m_j is constant in time, which, again, is true within each 3-hour interval because the cell boundaries are fixed. ORBIT also enforces air-mass continuity, so that the air entering each cell balances the air leaving it (Section 2.2.3). The balance is not exact, because the true surface pressure (and thus the air mass of each cell) can change within each interval, and furthermore because ORBIT forms the air fluxes from the 3-hourly means of the wind and pressure thickness, omitting their covariance. On the 2022 ORBIT grid, I find that the imbalance has a median of 1.5% of the air mass in each cell per day, which is 1–2 orders of magnitude smaller than the typical losses from deposition and chemistry.

Assuming hydrostatic balance, $\rho \, \mathrm{d}z = -\mathrm{d}p/g$, which is also assumed by the input meteorology (and holds at grid spacings of about 10 km or larger), I express the air mass in terms of pressure, because on the hybrid grid, the pressure thickness of each cell follows directly from the surface pressure:

$$m_j = \frac{\Delta p_j A_j}{g}, \tag{5}$$

where Δp_j is the cell's pressure thickness (the difference between the pressures at its bottom and top), A_j is its horizontal area, and g is the gravitational acceleration.

To integrate the next term of Equation 3, I use the divergence theorem to express the volume integral as a surface integral over the faces of the cell:

$$\int_{V_j} \nabla \cdot (\rho q \, \mathbf{u}) \, \mathrm{d}V = \oint_{\partial V_j} \rho q \, \mathbf{u} \cdot \mathbf{n} \, \mathrm{d}S = \sum_l \Phi_{jl}^{\mathrm{adv}}, \tag{6}$$

summed across the cells l sharing a face with cell j, where $\mathbf{n}$ is the outward unit normal and Φ_{jl}^{adv} is the net rate at which pollutant mass leaves cell j for cell l. Faces at the ground, the lateral edges, and the top of the model have no neighboring cell, and I apply boundary conditions there instead (Section 2.2.5).

Similarly, I apply the divergence theorem (with the same boundary treatment) to the first term on the right-hand side of Equation 3:

$$-\int_{V_j} \nabla \cdot \mathbf{J} \, \mathrm{d}V = -\oint_{\partial V_j} \mathbf{J} \cdot \mathbf{n} \, \mathrm{d}S = -\sum_l \Phi_{jl}^{\mathrm{sub}}, \tag{7}$$

where Φ_{jl}^{sub} is the net rate at which pollutant mass leaves cell j for cell l by eddy diffusion or (through the faces between layers) convection. When $\sum_l \Phi_{jl}^{\mathrm{adv}}$ is moved to the right-hand side of the equation, the transport terms combine into the net mass flux into cell j through the face it shares with cell l:

$$F_{jl} = -\left(\Phi_{jl}^{\mathrm{adv}} + \Phi_{jl}^{\mathrm{sub}}\right). \tag{8}$$

To integrate the next term in Equation 3, I assume that the reacting pollutants are well mixed within each cell, so that the chemical rates can be evaluated using the cell-mean mixing ratios $\mathbf{q}_j$ of the pollutants:

$$\int_{V_j} \rho \, (P - Q) \, \mathrm{d}V \approx m_j \left[P(\mathbf{q}_j) - Q(\mathbf{q}_j)\right] \equiv m_j \, (P_j - Q_j). \tag{9}$$

In reality, pollutants are not well mixed within a cell: pollution has strong spatial structure, and segregated pollutants within a grid cell react more slowly (whereas colocated pollutants, such as oxidants inside a plume,

can react faster). Because the chemical rates are in general nonlinear in the reactant mixing ratios, the true volume integral depends on how the pollutants are distributed at each point within the cell, which Eulerian models do not resolve.

Finally, integrating the emission and deposition terms gives the cell totals:

$$E_j = \int_{V_j} E \, \mathrm{d}V, \qquad D_j = \int_{V_j} D \, \mathrm{d}V. \tag{10}$$

Collecting the terms, Equation 3 becomes:

$$m_j \frac{\mathrm{d}q_j}{\mathrm{d}t} = \sum_l F_{jl} + m_j \, (P_j - Q_j) + E_j - D_j. \tag{11}$$

The following assumptions then make Equation 11 linear in the mixing ratios q_j. The sub-grid component of the $\sum_l F_{jl}$ term (see Equation 8) depends on $\mathbf{J}$, which ORBIT models as eddy diffusion and convective mass flux, both of which are linear in the mixing ratio. The advective component is also linear, because in ORBIT the mixing ratio carried across a face is a weighted combination of the mixing ratios in the 2 cells sharing the face, with weights that depend on the meteorology but not on the mixing ratios. The fluxes across each face therefore take the form

$$F_{jl} = \alpha_{jl} \, q_l - \alpha_{lj} \, q_j, \qquad \alpha_{jl}, \alpha_{lj} \geq 0, \tag{12}$$

where α_{jl} is the rate, in kg of air per second, at which air carries pollutant from cell l into cell j (Section 2.2). Many chemical transport models instead use higher-order advection schemes with flux limiters, which reduce numerical diffusion while preventing spurious oscillations, and which make the advection a nonlinear function of the surrounding cells' mixing ratios. ORBIT forward simulations apply this kind of flux correction outside the linear operator (Section 2.2.1).

For the chemical term $m_j(P_j - Q_j)$, I assume that chemical loss is proportional to the mixing ratio, $Q_j = k_j \, q_j$, with a rate coefficient k_j that does not depend on q_j (Section 2.3), and that production P_j depends only on the mixing ratios of precursor species, which ORBIT solves first (Section 2.1.4).

For the deposition term D_j, I assume that deposition is a first-order loss,

$$D_j = m_j \, d_j \, q_j, \tag{13}$$

with a removal rate coefficient d_j that does not depend on q_j (Sections 2.2.6 and 2.2.7). For species that partition between the gas and particle phases, d_j depends on the fraction in each phase, which I hold fixed within each periodic solve. The outer iteration (Section 2.1.5) later updates the fractions at the solved concentrations.

Substituting the face fluxes (Equation 12), the first-order chemical loss, and deposition (Equation 13) into Equation 11, and dividing by m_j, gives for each cell

$$\frac{\mathrm{d}q_j}{\mathrm{d}t} = \sum_l \frac{\alpha_{jl}}{m_j} \, q_l - \Big(\sum_l \frac{\alpha_{lj}}{m_j} + k_j + d_j \Big) q_j + \frac{E_j}{m_j} + P_j, \tag{14}$$

Note that the outflow of pollution through the boundary of the regional domain is included in the second sum (the loss term), but there is no inflow of pollution from the domain boundary in the first sum (Section 2.2.5).

ORBIT carries concentrations at local density, which are needed to relate the pollution to emissions, chemistry, and exposure, so I convert Equation 14 to concentrations accordingly. The concentration in cell

j is $c_j = \rho_j q_j$, where $\rho_j = m_j/V_j$ is the cell's mean air density. Because ρ_j is constant within an interval, multiplying Equation 14 by ρ_j and substituting $q_l = c_l/\rho_l$ gives

$$\frac{\mathrm{d}c_j}{\mathrm{d}t} = \sum_l \frac{\rho_j}{\rho_l}\frac{\alpha_{jl}}{m_j}\, c_l - \Big(\sum_l \frac{\alpha_{lj}}{m_j} + k_j + d_j\Big)c_j + \frac{E_j}{V_j} + \rho_j P_j. \tag{15}$$

Collecting the equations for all N cells gives, for each pollutant and interval k,

$$\frac{\mathrm{d}\mathbf{c}}{\mathrm{d}t} = \mathbf{L}_k\mathbf{c} + \mathbf{e}_k, \tag{16}$$

where $\mathbf{c} \in \mathbb{R}^N$ holds the cell concentrations in $\mu\mathrm{g\,m^{-3}}$, and the operator $\mathbf{L}_k$ is the $N \times N$ matrix of the coefficients in Equation 15,

$$L_{jl} = \frac{\rho_j}{\rho_l}\frac{\alpha_{jl}}{m_j} \quad (l \neq j), \qquad L_{jj} = -\Big(\sum_l \frac{\alpha_{lj}}{m_j} + k_j + d_j\Big). \tag{17}$$

The off-diagonal entry L_{jl} is the rate coefficient at which air from cell l brings pollution into cell j, and $-L_{jj}$ is the total rate coefficient at which cell j loses pollutant by outflow, chemical loss, and deposition. The source vector $\mathbf{e}_k$ has entries $e_j = E_j/V_j + \rho_j P_j$, the emissions per unit volume and the production from precursor species, through which chemistry couples the equations for different species.

Equation 16 expresses the continuity equation for every cell in the domain as a single linear system. Its operator has 4 important properties that follow directly from the physics (Figure 1).

Sign structure. The off-diagonal entries of the operator are non-negative and its diagonal entries are non-positive (Equation 17). Physically, air arriving from a neighboring grid cell can only bring pollution in, and outflow, chemical loss, and deposition can only remove it.

Mass conservation. The columns of the operator sum to zero for transport, and the losses on the diagonal make them negative. Column l describes where the pollution in cell l goes. Its diagonal entry is the rate at which cell l loses pollution, and its off-diagonal entries are the rates at which that pollution arrives in the neighboring cells. Transport only moves pollution between cells, so everything that leaves cell l by transport arrives in a neighbor, and the column sums to zero. Outflow, chemical loss, and deposition appear only on the diagonal, so they make the sum negative. Strictly, the sums are weighted by cell volume, $\sum_j V_j L_{jl} \leq 0$, because the operator acts on concentrations rather than masses. Emissions and chemical production, the only sources of pollutant mass, enter through $\mathbf{e}_k$ rather than $\mathbf{L}_k$, so within each interval the total pollutant mass in the domain changes only through them and through outflow, chemical loss, and deposition. ORBIT carries the concentration across the boundary between intervals, where cell volumes change slightly with surface pressure, so pollutant mass changes there in proportion to the change in each cell's volume.

Sparsity. The operator is sparse, because the physics is local. Transport links only cells that share a face, and chemistry and deposition act within a cell, so each row of $\mathbf{L}_k$ has at most 7 nonzero entries, the cell and its 6 neighbors, out of $N = 69{,}225$.

Preservation of uniform mixing ratios. Transport leaves a uniform mixing ratio unchanged, and thus (given the sign structure) cannot create spurious peaks or holes in a well-mixed field. In matrix terms, the rows of the transport part of $\mathbf{L}_k$, weighted by the density of each cell, sum to zero, $\sum_l L_{jl}\rho_l = 0$ (except in cells at the lateral boundaries, where no pollution enters from outside). The sums vanish because of air-mass continuity, which makes the air entering each cell equal to the air leaving it, to within the residual given above.

The sign structure and mass conservation together make $-\mathbf{L}_k$ an M-matrix (Berman & Plemmons, 1994), so that $\mathbf{I} - \Delta\tau\mathbf{L}_k$ has a non-negative inverse for any time step $\Delta\tau > 0$, which keeps the implicit steps of the

periodic solution non-negative however long they are (Section 2.1.2). Sparsity makes the system tractable, because a dense operator for 1 species and bin would occupy tens of gigabytes, which would be hard to store or read into memory, let alone factor, whereas the sparse operator has at most $7N$ nonzero entries and can be factored directly (Section 2.1.4). The preservation of uniform mixing ratios (given the sign structure) means that every peak or hole in the modeled concentrations arises from emissions, chemistry, or removal.

2.1.2 Periodic Solution

ORBIT directly solves for the periodic steady state of each month. Each month is represented by 1 day, with 8 three-hour operators $\mathbf{L}_0, \ldots, \mathbf{L}_7$, derived from meteorology averaged separately for each time bin across all the days in the month (Section 2.1.3).

Within bin k, the operator $\mathbf{L}_k$ and the sources $\mathbf{e}_k$ are constant, so integrating Equation 16 over the bin gives the concentration at its end exactly,

$$\mathbf{c}_{k+1} = e^{\Delta\tau \mathbf{L}_k}\mathbf{c}_k + \mathbf{b}_k, \tag{18}$$

where $\Delta\tau = 3$ h is the length of the bin, $\mathbf{c}_k$ is the concentration at its start, and $\mathbf{b}_k$ is the concentration that the sources $\mathbf{e}_k$ produce by the end of the bin. The matrix exponential $e^{\Delta\tau\mathbf{L}_k}$ describes how the transport, chemical loss, and deposition transform the concentrations from the start of the bin to the end.

However, the exponential of a sparse matrix is dense, so instead I approximate the time derivative in Equation 16 over the bin, as

$$\frac{\mathbf{c}_{k+1} - \mathbf{c}_k}{\Delta\tau} = \mathbf{L}_k\mathbf{c} + \mathbf{e}_k. \tag{19}$$

Here, we have a choice of whether to evaluate the concentration on the right-hand side, $\mathbf{c}$, as the concentration at the start or end of the bin. The simplest explicit scheme, forward Euler, uses the starting concentration, $\mathbf{c} = \mathbf{c}_k$, so that $\mathbf{c}_{k+1} = (\mathbf{I} + \Delta\tau\mathbf{L}_k)\mathbf{c}_k + \Delta\tau\mathbf{e}_k$. However, the forward Euler approximation can produce negative concentrations, if it is assumed to lose more pollution over a time step than it holds. For example, this issue arises from transport, when the time step chosen is more than the time taken for the wind to traverse across a cell (the Courant limit).

An alternative is the implicit, backward Euler scheme, which evaluates the concentration at the end of the step, $\mathbf{c} = \mathbf{c}_{k+1}$. In backward Euler, we find the final concentration by solving the linear system

$$(\mathbf{I} - \Delta\tau\,\mathbf{L}_k)\,\mathbf{c}_{k+1} = \mathbf{c}_k + \Delta\tau\,\mathbf{e}_k. \tag{20}$$

For ORBIT, I use 1 backward Euler step per bin, so that $\mathbf{c}_{k+1} = \mathbf{M}_k\mathbf{c}_k + \mathbf{b}_k$, with $\mathbf{M}_k = (\mathbf{I} - \Delta\tau\mathbf{L}_k)^{-1}$ in place of $e^{\Delta\tau\mathbf{L}_k}$ and $\mathbf{b}_k = \Delta\tau\,\mathbf{M}_k\mathbf{e}_k$. Because $-\mathbf{L}_k$ is an M-matrix (Section 2.1.1), $\mathbf{M}_k$ is non-negative for any step length, so the concentrations at the end of each step cannot go negative.

To find the concentrations at the end of the day, we can compose the 8 operators for each bin sequentially,

$$\mathbf{M} = \mathbf{M}_7\mathbf{M}_6\cdots\mathbf{M}_0. \tag{21}$$

Then our periodic steady state assumption can be states as finding a starting concentration that returns to itself by the end of the day, i.e. $\mathbf{c}_8 = \mathbf{c}_0$, so that

$$(\mathbf{I} - \mathbf{M})\mathbf{c}_0 = \mathbf{s}, \qquad \mathbf{s} = \sum_{k=0}^{7} \mathbf{M}_7\cdots\mathbf{M}_{k+1}\,\mathbf{b}_k. \tag{22}$$

Here, $\mathbf{s}$ is the concentration at the end of the day, produced by the transport, loss, and removal of the emissions and chemical production within that day. Rewriting as $\mathbf{c}_0 = \sum_{n=0}^{\infty}\mathbf{M}^n\mathbf{s}$, we can see that solving for $\mathbf{c}_0$ is

	from 1	from 2	from 3	from 4	from 5	from 6	row sum
to 1	−4	1	0	1	0	0	−2
to 2	3	−5	1	0	1	0	0
to 3	0	3	−4	0	0	1	0
to 4	1	0	0	−5	1	0	−3
to 5	0	1	0	3	−6	1	−1
to 6	0	0	1	0	3	−5	−1
column sum	0	0	−2	−1	−1	−3	

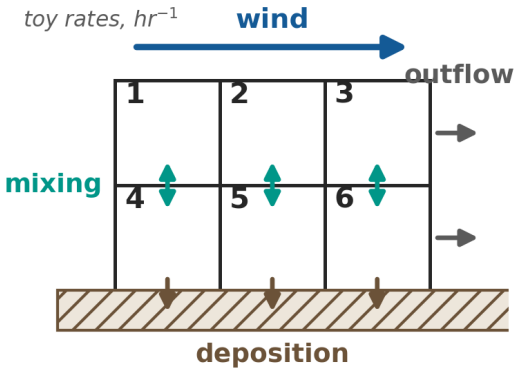

Fig. 1. A toy operator **L** for a domain of 6 equal cells in 2 layers under a uniform wind, with vertical mixing between layers, deposition from the lower layer, and outflow at the downwind boundary (rates in h^{-1}). The diagonal entries are non-positive and the off-diagonal entries are non-negative. Because the cells are equal, the weights of Section 2.1.1 drop out, so each column sums to zero less the deposition and outflow from that cell, and each row sums to zero less the deposition and the inflow missing at the upwind boundary. The sign structure and the column sums make $-\mathbf{L}$ an M-matrix, so that the implicit solves in Section 2.1.2 are stable and non-negative.

equivalent to adding the concentrations from the day to the concentrations that remain from the previous day, and so on *ad infinitum*. Because $\mathbf{M}$ only transports or removes pollution but does not create it, the concentrations do not diverge. We call this matrix $\mathbf{M}$ the monodromy operator (Floquet, 1883).

Solving for the periodic steady state is a departure from chemical transport models, which advance pollutant concentrations through the entire trajectory of the simulated period in small time steps (Byun & Schere, 2006) (Lin & Rood, 1996). They typically use explicit numerical methods that step the system forward in time, with implicit schemes often reserved for vertical diffusion or certain stiff chemical systems (Byun & Schere, 2006). Long, implicit time steps such as those used in ORBIT are often avoided, because they are generally costly (each time step requires solving a large linear system, with operators changing often and with nonlinear chemistry). ORBIT avoids the cost by using linear operators for a repeating system, so the factors can be reused for every application of $\mathbf{M}$ (Section 2.1.4). The closest parallel to ORBIT is InMAP, which averages the meteorology across all time bins and solves for a single static, annual steady state (Tessum, Hill, & Marshall, 2017). By contrast, ORBIT solves for a periodic steady state, effectively asking what the pollutant concentrations would be for an averaged day if it repeated forever. Thus, ORBIT retains diurnal dynamics important for pollutant chemistry, transport, and emissions.

Over the periodic day, the states at the end of the bin balance the emissions and losses exactly: summing over Equation 20 across the 8 bins gives $\sum_k (\mathbf{L}_k \mathbf{c}_{k+1} + \mathbf{e}_k) = 0$. Using Equation 20, I recover the remaining bin states from $\mathbf{c}_0$, and report their mean as the monthly average concentration.

2.1.3 Meteorological Averaging

For each calendar month, ORBIT assembles the operators $\mathbf{L}_k$ across the 8 3-hour bins of each day from the atmospheric processes derived from its input meteorology (Section 2.5.2). The processes are averaged across all the days in the month for each 3-hour bin, never averaging across different bins, so each month keeps its diurnal cycle (Figure 2). There are 96 operators per year (8 bins × 12 months), centred from 01:30–22:30 UTC, (07:00–04:00 Indian Standard Time).

For turbulent mixing, dry deposition velocities, wet scavenging rates, and the vertical air-mass fluxes (Section 2.2.3), I find the monthly mean rates for each time bin. As some of these relationships are nonlinear (for example, in rainfall and surface conditions), I compute the rates for each day and then average, rather than computing the rates from averaged meteorology. Similarly, for wind and vertical air-mass flux, I separately average the positive and negative directions of each component, so that winds blowing in opposite directions do not cancel out (Section 2.2.1). ORBIT neglects the day-to-day covariance between the process rates and mixing ratios, potentially leading to low bias when (for example) weak winds coincide with high concentrations.

2.1.4 Linear Solver

Although Equation 22 is linear in the N cell concentrations of each species, $\mathbf{M}$ is the product of inverses of 8 sparse matrices and thus is dense. Solving for the periodic steady state can be done instead using Krylov subspace methods.

Each matrix $\mathbf{I} - \Delta\tau\mathbf{L}_k$ is sparse (Section 2.1.1) and so I factor them directly into lower and upper triangular factors using UMFPACK (Davis, 2004). To reduce fill-in (and thus keep the factors sparse), I reorder the matrices using METIS nested dissection (Karypis & Kumar, 1998). Because the transport geometry (and thus, sparsity pattern) of all species and bins in each month is identical, I only require 1 symbolic factorization per month, with only the numeric factors needing recomputation for each species and bin within each outer

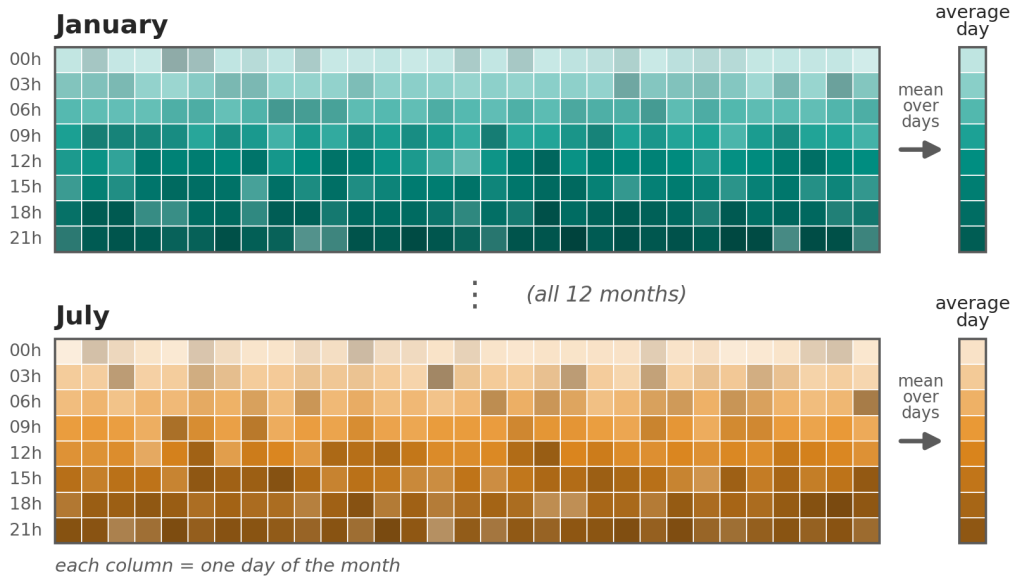

Fig. 2. How ORBIT averages the meteorology. Each row is one 3-hour bin of the day and each column one day of the month. ORBIT averages along the rows, across the days of the month, and never across bins, so each month yields an averaged day of 8 states with its own diurnal cycle. Two months are shown, and the year yields 96 states.

iteration.

Then I solve for the periodic steady state in Equation 22 using matrix-free Generalized Minimal Residual Iteration (GMRES) (Saad & Schultz, 1986), which only requires applying $\mathbf{M}$ to vectors, which can be done by applying 8 sparse back solves sequentially (once per bin) and thus never requiring me to explicitly form $\mathbf{M}$.

I implement GMRES (restart length 30, iteration cap 200, relative tolerance 10^{-4} in outer iteration and 10^{-6} in final solve), and it converges in tens of iterations. Each species is solved sequentially, after its precursors (sulfate after SO_2, nitrate after NO_x, organic bins in descending volatility).

2.1.5 Nonlinear Closure

By construction, $\mathbf{L}$ is linear in concentration. ORBIT fixes the terms that are dependent on concentrations during its linear solve, and then updates these terms, based on the final concentrations after the linear solve, in an outer iteration.

Reduced nitrogen and nitrate are partitioned according to inorganic equilibrium (Section 2.3.5), which depends on sulfate concentrations. Organic partitioning depends on the availability of absorbing organic mass (Section 2.3.6). There is also coupling between the effective deposition rates of transported species and their gas-particle fractions (Sections 2.2.6 and 2.2.7). Furthermore, the anti-diffusive correction of the flux-corrected horizontal transport scheme also depends on the concentrations (Section 2.2.1).

For each periodic solve, I hold the inorganic and organic fractions and the transport correction fixed at their values from the previous outer iteration. Afterwards, I update them at the resulting concentrations in a fixed-point (Picard) iteration. Figure 3 illustrates the periodic solve and the outer iteration for each month.

For the surface nitrate fractions, which converge slowly (especially in warmer months), I use Anderson

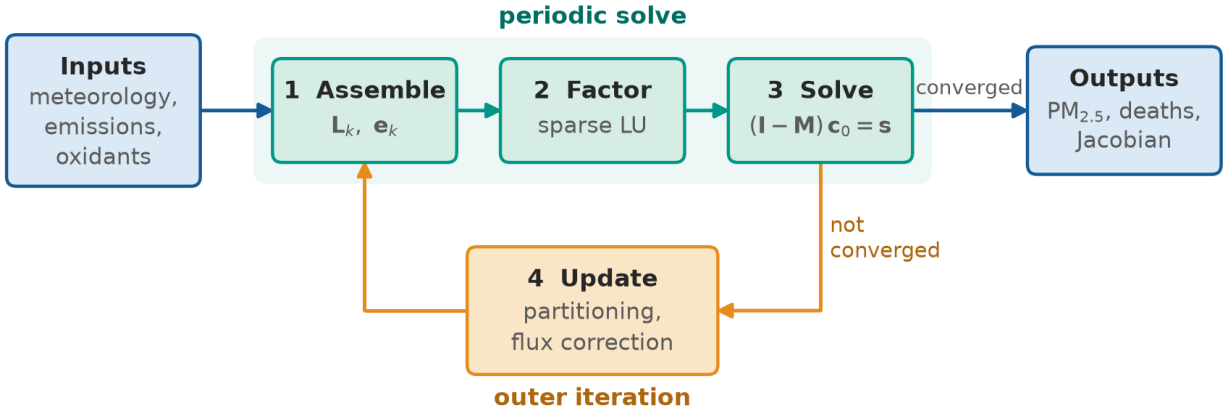

Fig. 3. How ORBIT solves 1 month. The meteorology, averaged into 8 bins of process rates, the emissions, and the oxidants define the operator $\mathbf{L}_k$ and source $\mathbf{e}_k$ of each species and bin (step 1), with the gas-particle fractions and the flux correction held fixed. ORBIT factors the 8 bin matrices of each species (step 2) and solves the periodic system of each species in chemical order, with sulfate after SO_2, nitrate after NO_x, and the organic bins from the most volatile down (step 3). Steps 1 to 3 are linear. Step 4 recomputes the inorganic and organic particle fractions, and in forward and zero-out runs the flux correction, from the solved concentrations, and the loop repeats until the criterion of Section 2.1.5 is met or the iteration cap is reached. The outputs are $PM_{2.5}$, attributable deaths, and the Jacobian for the sensitivity runs.

acceleration with depth 3 (Anderson, 1965; Walker & Ni, 2011) for all 8 bins (scaling the column profiles), after damping the first update by 0.5. I take the ammonium fractions from the equilibrium at each iteration. Near changes in the thermodynamic regime (where the extrapolation may overshoot), I fall back to a Picard iteration. From the fourth iteration, I fix cells that alternate between the 2 thermodynamic branches at the midpoint of the branches. The iteration stops either at 6, or earlier if 90% of the relative change between iterations falls below 2

2.1.6 Simulation Modes

ORBIT has 3 simulation modes (Table 1). Forward mode predicts total and speciated $PM_{2.5}$ and mortality for total emissions. Marginal mode linearizes the model around a forward mode simulation, and predicts changes in total and speciated $PM_{2.5}$ from small changes in emissions, keeping track of pollutant cross-couplings such as nitrate increases following SO_2 reductions. The adjoint of the marginal mode gives the contribution every source has to a single outcome such as $PM_{2.5}$ in one city, or total attributable deaths. The marginal mode and its adjoint are only valid where the gas-particle partitioning stays in the same thermodynamic regime; the zero-out mode estimates how larger changes in emissions (such as removing an entire sector) affect changes in concentrations, by taking the difference between 2 forward simulations.

At the end of the final iteration of a forward simulation, I compute a Jacobian that contains the marginal particle fractions, chemical production, organic aging, and 4 cross-partial terms (computed using nested finite differences) that describe how changes in particulate ammonium and nitrate respond to their total concentrations and the total concentration of sulfate. I find the marginal particle fraction, the change in particulate NH_4^+ or NO_3^- per unit change in total reduced nitrogen or total nitrate, by evaluating the equilibrium at totals 1% above and below the baseline:

$$f^{\mathrm{marg}} = \mathrm{clip}\left(\frac{(fc)_+ - (fc)_-}{c_+ - c_-},\, 0,\, 1\right), \tag{23}$$

where f is the equilibrium fraction and c is the corresponding total. The marginal fraction lies between 0 and 1, but the finite differences across the equilibrium table can fall outside that range near regime boundaries, and can give negative values in cells where ammonium already neutralizes all of the sulfate, such as Delhi in winter, so I clip the marginal fraction to [0, 1]. I compute the cross-partials from the clipped fraction, so that the Jacobian describes the same operator as the forward run.

For marginal mode, I solve for the linearized (tangent-linear) periodic steady state, as in the forward simulation (Sections 2.1.2 and 2.1.4), with fixed baseline marginal fractions, organic partitioning, and oxidants. The linearization keeps 3 effects of gas-particle partitioning. First, a change in one species can alter the particle fraction of another, and thereby the removal of that species, because particles and gases deposit at different rates, so I include coupling terms between the species in each cell that are proportional to the difference between the particle and gas deposition rates. Second, nitrate and ammonium depend on each other, which I resolve by Picard iteration. Third, the $PM_{2.5}$ response includes the increase in organic particle fractions with absorbing mass.

For source attribution, I use the adjoint of the marginal mode linearization. I consider a receptor that is linear in the concentrations, $J = \sum_k \langle \mathbf{S}_k, \mathbf{G}_k \mathbf{c}_k \rangle$, where $\mathbf{S}_k$ weights surface $PM_{2.5}$ in bin k and $\mathbf{G}_k$ assembles surface $PM_{2.5}$ from the species. With $\mathbf{R}_k = \mathbf{G}_k^{\mathsf{T}} \mathbf{S}_k$, the periodic adjoint satisfies

$$\begin{aligned} \boldsymbol{\lambda}_k &= \mathbf{R}_k + \mathbf{M}_k^{\mathsf{T}} \boldsymbol{\lambda}_{k+1}, \\ (\mathbf{I} - \mathbf{M}^{\mathsf{T}})\boldsymbol{\lambda}_0 = \mathbf{s}_{\mathrm{adj}}, \qquad & \frac{\partial J}{\partial \mathbf{e}_k} = \Delta\tau \mathbf{M}_k^{\mathsf{T}} \boldsymbol{\lambda}_{k+1}, \end{aligned} \tag{24}$$

where $\mathbf{s}_{\mathrm{adj}}$ accumulates the receptor terms over the reversed daily cycle. Applying $\mathbf{M}^{\mathsf{T}}$ requires 8 transposed back-solves with the existing LU factors. The adjoint solves the chemistry backwards, so that the SO_2 sensitivity includes the sulfate it forms and each VOC sensitivity includes the organic bins it contributes to. As in marginal mode, the receptor includes the increase in organic particle fractions with organic aerosol mass, evaluated at the baseline concentrations. Thus, one adjoint solve per month and receptor gives the sensitivities for every species, source cell, and bin. For health impacts, $\mathbf{S}_k$ is the deaths gradient (Section 2.4).

Forward and zero-out simulations use flux-corrected transport (FCT; Section 2.2.1), whereas marginal mode simulations and their adjoint default to using the low-order operator without flux correction. This is an M-matrix, so a small addition of an inert primary species can never lower its concentration anywhere, and its sensitivities cannot be negative. An optional linearization of FCT, with the limiter held at its baseline values (the frozen-limiter linearization), is not an M-matrix, so it can produce negative sensitivities, and it takes about 8 times the wall time, because its correction must converge within an outer GMRES iteration (Table 1).

Using emission pulses from a single cell at Delhi in January 2022, I find that the low-order adjoint agrees with the marginal mode to ratios of 0.997 and 0.990 for POA and NO_x pulses. For a primary $PM_{2.5}$ pulse of 1% the black carbon emissions rate in that cell, the marginal response matches the nonlinear zero-out response using the low-order forward model with a slope of 1.000 and a deaths-weighted ratio of 1.000. Furthermore, against the full, forward simulation used in production (including FCT), the same marginal response captures 98% of the change in $PM_{2.5}$ totals across the domain, and 88% of the change in deaths (FCT keeps pollution more concentrated near large sources). When the results are weighted by population or aggregated to subdistricts, the difference is smaller. Using the low-order operator instead of the linearization only changes the total damages over the domain slightly. I further test the linearization with emission changes

Table 1. The 3 simulation modes of ORBIT, with wall time and peak memory per simulated month for the South Asian configuration at default settings, measured for version 1.0.2 on an office workstation with 16 cores and 31 GB of memory (forward: January, July, and October; marginal: January, ammonia perturbation).

Mode	Question answered	Wall time	Memory
Forward	Concentrations and health impacts for a complete emissions field	22 to 23 min	3.6 GB
Marginal	Response to a small change in source location, time, and species	5 min	5.5 GB
Marginal, flux-corrected	Marginal response with the frozen-limiter FCT linearization	42 min	5.7 GB
Zero-out	Nonlinear response to sector removal or a large emission change	2 forward runs	3.6 GB

of 0.1%, 1%, 10%, and 50% across the domain, and record where in the zero-out simulations the gas-particle partitioning changes, to investigate where the changes move the atmosphere into a different thermodynamic regime.

2.2 ORBIT Physics

ORBIT transports 12 chemical species, each with its own linear operator whose entries are the rates of all physical processes acting on the species (such as those illustrated in Figure 4). The transport terms in the operators are assembled to ensure that air mass, $m_j = \Delta p_j A_j / g$, is conserved, and the mixing ratios are converted to concentrations assuming local density, $\rho_j = \Delta p_j / (g\,\Delta z_j)$ (Equation 17).

2.2.1 Horizontal Advection and Diffusion

Horizontal transport includes both advection (the bulk movement of pollution across cells, caused by the wind) and diffusion (the mixing of pollution between neighboring cells, caused by small-scale turbulence and circulations that spread pollution down its concentration gradient). The Péclet number $\mathrm{Pe} = |u|\Delta x / K$, defined as the ratio of the advective to diffusive transport for wind speed u, cell width Δx, and diffusivity K, is a useful measure of which term dominates. Within the boundary layer, the cell diffusivity is the Monin-Obukhov momentum diffusivity $\kappa u_* z (1 - z/h)^2 / \phi_m$, where $\kappa = 0.41$ is the von Kármán constant, u_* is the friction velocity, z is the height above ground, h is the boundary-layer height, and ϕ_m is the stability function for momentum (Businger, Wyngaard, Izumi, & Bradley, 1971; Dyer, 1974). Above the boundary layer, the diffusivity is 0.5 $\mathrm{m^2\ s^{-1}}$ (Wilson, 2004). The face diffusivity K is the harmonic mean of the diffusivities of neighbouring cells, so that exchange of pollution is weak when one of the cells is poorly mixed.

I use the exponential scheme described by Patankar (Patankar, 1980), which assumes that the fluxes across each face are those that would result from the exact solution of steady, one-dimensional advection and diffusion between adjacent cell centers. At each face between a left and a right cell, with $D = K/\Delta x$ and the Péclet number of the bin-mean wind $\bar{u}$, the coefficients across each face are given by

$$
\begin{aligned}
a_L &= D\,A(|\mathrm{Pe}|) + u^+, \qquad a_R = D\,A(|\mathrm{Pe}|) + u^-, \\
A(|\mathrm{Pe}|) &= \frac{|\mathrm{Pe}|}{\exp(|\mathrm{Pe}|) - 1},
\end{aligned}
\tag{25}
$$

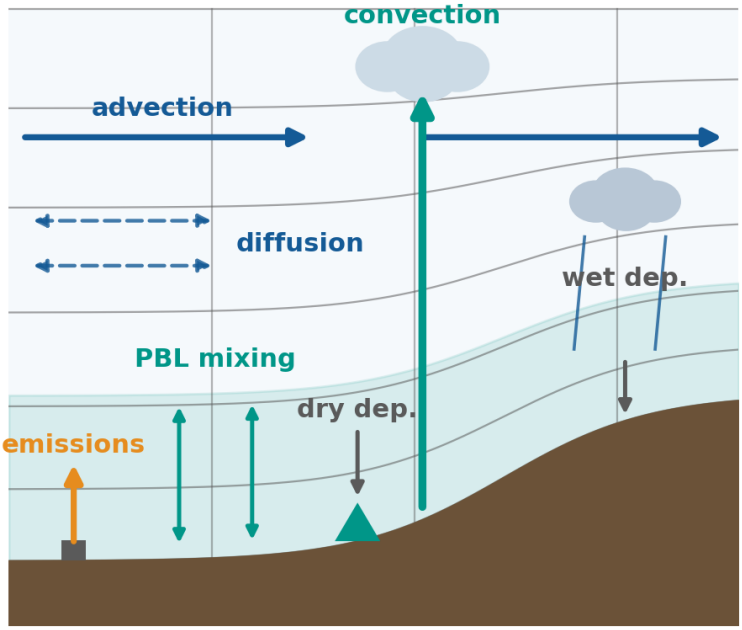

Fig. 4. *Physical processes represented by ORBIT.* An illustration of many of the physical processes represented in the transport operator in ORBIT (see Section 2.2). Chemistry is not shown here but is discussed in Section 2.3. Emissions are shown in orange but are not represented in the operator, and are discussed in detail in Section 2.5.3.

where $u^+ = \langle \max(u, 0) \rangle$ and $u^- = \langle \max(-u, 0) \rangle$ are the bin means of the positive and negative components of the wind in each 3-hour interval (Section 2.1.3). Together, the 2 parts carry the bi-directional exchange produced by wind meander and reversal, equivalent to a diffusivity $\min(u^+, u^-)\Delta x$. In still air (Pe $\to$ 0), the exponential scheme approaches central differencing (where the cells are weighted equally). In strong wind (large |Pe|), it approaches upwind differencing (where the flux depends only on the upwind cell). Both face coefficients remain non-negative at any wind speed, so a cell can never lose more pollution than it has, and the operator retains the sign structure described in Section 2.1.1 that ensures non-negative concentrations.

For each face between neighbouring cells, pollution can move in both directions, with the rate $a_L/\Delta x$ in one direction and $a_R/\Delta x$ in the other direction, with the mixing term being larger in the direction the wind is blowing. I multiply each rate by the ratio of the pressure thickness at the face to that of the cell, $\Delta p_f/\Delta p_j$, to account for the fact that adjacent cells (when the terrain underneath changes height) can hold diferent amounts of air. Thisensures that the air lost by each cell balances with the air gained by its neighbours. For cells facing north to south, the same applies (using the northward wind v and the north-south grid spacing Δy), but I also multiply by the ratio fo the width of the face to the width of the cell, because the longitude lines converge towards the poles, so that the width of the cell in the east-west direction shrinks with latitude as $\Delta x(\phi) = R_E \cos\phi\, \Delta\lambda$ for Earth's radius R_E, latitude ϕ, and longitude spacing $\Delta\lambda$. The resulting rates are the exchange rates α_{jl}/m_j of Equation 14.

As the Péclet number increases, the advection scheme approaches upwind differencing, with the mixing term in the face coefficients decreasing to about a third of its value at Pe = 2 and about 1% at Pe = 6. Upwind differencing is numerically diffusive, with diffusivity $|u|\Delta x/2$, which exceeds the physical diffusivity wherever Pe > 2. Generally, horizontal transport on the South Asian grid is advection-dominated, with median Pe = 6.3 and 70% of horizontal faces with Pe > 2. Overall, the exponential scheme is not much different from upwind differencing, with numerical diffusion supplying 75% of horizontal mixing (Figure 5).

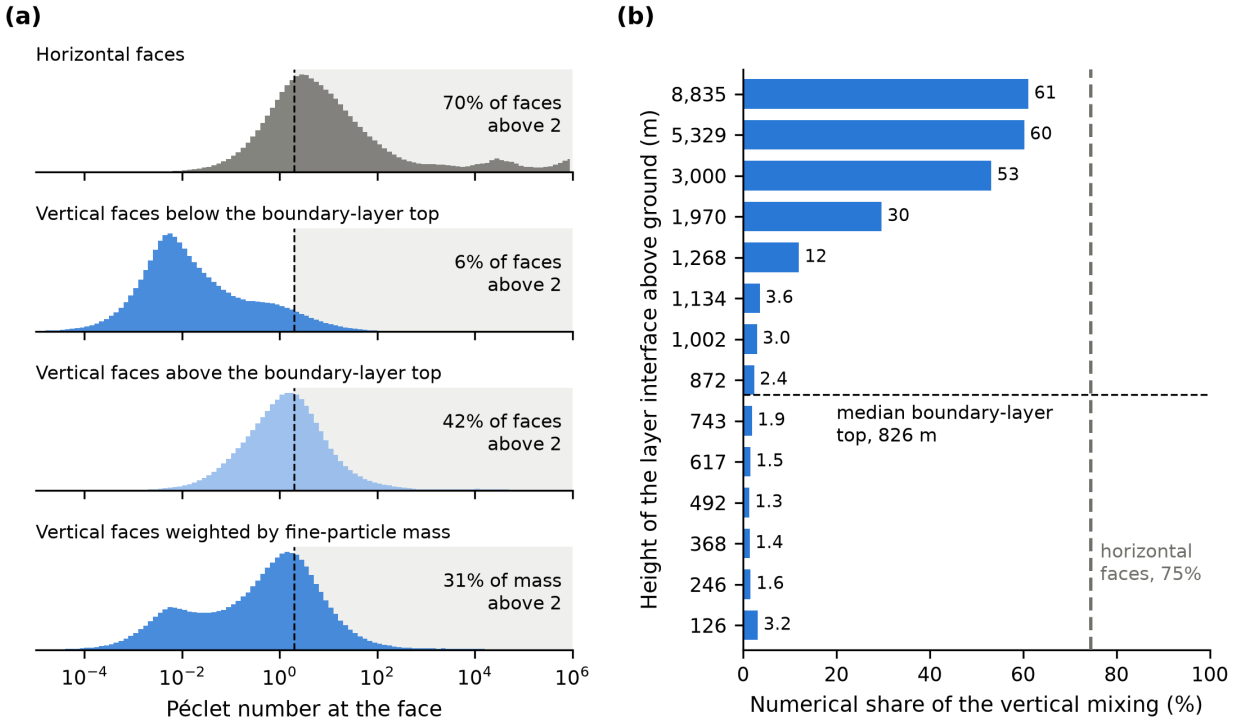

Fig. 5. Péclet numbers across the ORBIT grid in the South Asia domain, 2022 meteorology (about 19 million faces). **(a)** Péclet number distribution for horizontal faces (median: 6.3), vertical faces below the boundary-layer top (median: 0.012), vertical faces above the boundary layer (median: 1.38), vertical faces weighted by $PM_{2.5}$mass (median 0.60). Each face is each scaled to its peak. The shaded region (Péclet number above 2) shows where the numerical diffusion of a first-order upwind scheme exceeds the physical mixing. **(b)** The numerical share of the vertical mixing flux across ORBIT layers, weighted by the $PM_{2.5}$concentration gradient. Below the median boundary-layer top (826 m), only 1.5% of the mixing is numerical, whereas 75% for the horizontal scheme (gray dotted line). The numerical share increases above the boundary layer as the layers increase in size.

The inaccuracy is a general feature of linear advection schemes. By Godunov's theorem (Godunov, 1959), a linear advection scheme is at most first-order accurate unless it creates new maxima or minima.

ORBIT gets around this by using a deferred flux-corrected transport scheme in its forward simulations (keeping the low-order operator for marginal mode and its adjoint, see Section 2.1.6). It retains linear operators in its periodic solve, using the low-order Patankar exponential operator on the left-hand side of Equation 20. On the right-hand side, I add an anti-diffusive correction, the difference between the second-order van Leer fluxes and the low-order fluxes, which I find from the concentrations of the previous outer iteration. I scale each anti-diffusive flux by a factor between 0 and 1 from the multidimensional Zalesak limiter (Zalesak, 1979), so that the correction adds no new maxima or minima. When the factor is 0 and at the domain boundary, ORBIT falls back to first-order transport. The correction gets updated at each outer iteration (Section 2.1.5). With FCT, horizontal transport is second-order accurate where concentrations vary smoothly, and preserves monotonicity at sharp gradients.

2.2.2 Boundary-Layer Mixing

Boundary layer mixing is governed by small-scale eddies that are not resolved by the 3-hourly mean wind fields across the grids, and are instead part of the sub-grid flux **J** of Section 2.1.1, along with horizontal

diffusion (Section 2.2.1) and convection (Section 2.2.4). I model boundary layer mixing as a diffusion process along the vertical concentration gradient, with eddy diffusivity K_{zz}. Within the boundary layer, I assume the mixing follows the K-profile given by the YSU scheme (Hong, Noh, & Dudhia, 2006),

$$K_{zz}(z) = \kappa\, w_s\, z \left(1 - \frac{z}{h}\right)^2 \quad (z < h), \qquad w_s = \left(u_*^3 + 8\kappa\, w_*^3\, z/h\right)^{1/3}, \tag{26}$$

where h is the boundary-layer height from the meteorology data, and the velocity scale w_s combines mechanical turbulence, given by u_*, with convective turbulence, given by the convective velocity $w_* = (g\, h\, \overline{w'\theta'}/\theta_0)^{1/3}$. Here $\overline{w'\theta'}$ is the surface kinematic heat flux and θ_0 is the mean potential temperature of the boundary layer. The convective term is active only when the surface heat flux is upward, which makes the Obukhov length $L_{\mathrm{Ob}} = -\theta_0 u_*^3/(\kappa g\, \overline{w'\theta'})$ negative. The profile peaks within the boundary layer, and is zero at the ground and top of the boundary layer. Unlike the ACM2 scheme (Pleim, 2007), the YSU profile only acts between adjacent layers, preserving the sparsity of the operators and their LU factors (Section 2.1.4). Transport between distant layers is represented by the convective mass flux of Section 2.2.4.

In stable and neutral columns, the boundary layer mixing depends on the local stability of the air, in which case I multiply K_{zz} by the stable-case function $f = 1/(1 + 4.7\,\mathrm{Ri})^2$ of Louis (1979), where $\mathrm{Ri} = (g/\theta)(\partial\theta/\partial z)/|\partial\mathbf{v}/\partial z|^2$ is the gradient Richardson number of potential temperature θ and horizontal wind $\mathbf{v}$ (f is derived in Louis (1979) for a bulk Richardson number), and $f = 1$ where $\mathrm{Ri} \le 0$. Above the boundary layer, K_{zz} is 0.5 $\mathrm{m^2\ s^{-1}}$ (Wilson, 2004) multiplied by f. Everywhere, K_{zz} is at least 0.01 $\mathrm{m^2\ s^{-1}}$ (the value in the top layer), where I assume the air is strongly stable.

For horizontal faces (Section 2.2.1), I take the diffusivity at each interface as the harmonic mean of the neighbouring cell values, so that exchange is weak across an inversion. I scale the exchange coefficient $K_{\mathrm{face}}/[\Delta z_k\, \frac{1}{2}(\Delta z_k + \Delta z_{k+1})]$ by the ratio of the air density at the interface to the air density in the cell, to conserve air mass.

2.2.3 Vertical Advection

In the process of regridding and averaging the MERRA-2 pressure velocity and wind fields, ORBIT must ensure that air mass flux across each level balances the horizontal face fluxes (Section 2.1.1), which is achieved by comparing the air mass flux across vertical levels with the divergence of the horizontal face fluxes. With pressure velocity ω_k (Pa $\mathrm{s^{-1}}$, where positive values are downward) at the bottom interface of layer k, and no flux through the ground,

$$\omega_{k+1} = \omega_k + \frac{1}{A}\sum_f \Delta p_f\, u_f\, \ell_f + (B_k - B_{k+1})\frac{\mathrm{d}p_s}{\mathrm{d}t}, \qquad \omega_0 = 0, \tag{27}$$

where the sum is across each lateral cell face f, u_f is the outward wind normal to the face, Δp_f and ℓ_f are the face pressure thickness and length, and A is the horizontal area of the cell. The final term is the share of the change in surface pressure p_s across the layer between consecutive 3-hourly fields, where B_k is the hybrid coefficient of the lower edge of the layer, so that the air mass of the column changes only with the surface pressure. Because the levels follow the terrain, air moving horizontally up a slope changes pressure without crossing a level, and ω includes only the flux across levels. Before averaging, I compute ω for each 3-hour window (Section 2.1.3).

I split the cross-level flux into its non-negative downward and upward parts, as $\omega^+ = \max(\omega, 0)$ and $\omega^- = \max(-\omega, 0)$. Layer k loses air downward through its lower interface and upward through its upper

interface, at a rate of $(\omega_k^+ + \omega_{k+1}^-)/\Delta p_k$, and gains air from the layer below at a rate of $\omega_k^-/\Delta p_k$ and from the layer above at $\omega_{k+1}^+/\Delta p_k$. The 4 rates are the vertical-advection exchange rates α_{jl}/m_j of Equation 14.

I compute the vertical Péclet number with K_{zz} in place of K, to compare advection across each interface with the boundary-layer mixing of Section 2.2.2. In the boundary layer, the median face has Pe = 0.012, 94% of faces lie below Pe = 2, and numerical diffusion contributes only 1.5% of mixing. Above the boundary layer, the layer spacing increases from roughly 130 m to over 1 km, which likely explains why the Pe increases (median: 1.38) (Figure 5). Thus, first-order upwind transport is accurate in the boundary layer, where it matters most for exposure, and so I do not apply any flux correction to vertical advection as I did to vertical advection (Section 2.2.1).

2.2.4 Convection

Convective updrafts are also narrower than a grid cell and thus not represented by the pressure velocity of Section 2.2.3 except by their net flux, which does not capture the fast ascent of pollution within the updrafts, and the slower descent around them. I represent convection using the updraft mass flux from MERRA-2, M_c (kg m^{-2} s^{-1}), at each interface. MERRA-2 parameterizes convection with the Relaxed Arakawa–Schubert scheme, whose plumes rise from the top of the boundary layer, mix with the surrounding air, and then detrain at different heights (Molod, Takacs, Suarez, & Bacmeister, 2015). Where M_c decreases with height, air leaves the updraft, so the top of the cloud is given by the profile of M_c, without needing a separate parameter.

At each interface, the updraft carries pollution from layer k to the layer $k+1$ above it. Layer k loses pollution at the rate $M_c g/\Delta p_k$, and layer $k+1$ gains pollution at the rate $M_c g/\Delta p_{k+1}$, where each term is divided by the layer pressure thickness so that the pollutant mass balances. I represent convection as an exchange between adjacent layers, which preserves operator sparsity (Section 2.1.1). Following the cloud model used by the Relaxed Arakawa–Schubert scheme (Molod et al., 2015), the same mass flux M_c is returned downward through each interface as a compensating subsidence, so that convection leaves the air mass of each layer unchanged, consistent with ω carrying the net vertical flux.

2.2.5 Boundary Conditions

As ORBIT is currently a regional model, it must handle the edges of its domain. At the lateral boundaries, ORBIT assumes free outflow and no inflow, so pollution is free to leave the domain horizontally (and does not return), and no pollution enters from outside the domain. Concentrations are therefore likely biased low in grid cells close to the edges of the domain, because ORBIT does not account for pollution entering from outside the domain, such as from the Middle East or Southeast Asia. At the boundary faces, ORBIT uses first-order advection without the flux correction (Section 2.2.1). The vertical flux of pollution by advection is zero at the ground (Equation 27). Pollution can leave through the top of the domain, either by convection (Section 2.2.4) or with the residual of the upward integration in Equation 27 when that residual is upward. When the residual is downward, air enters through the top of the domain, and I assume that this air carries no pollution.

2.2.6 Dry Deposition

Pollution is removed by wet deposition, which involves cloud water or precipitation (Section 2.2.7), and by dry deposition, which is removal at the ground through contact with the surface of the Earth. Dry deposition is generally modeled with a deposition velocity V_d, which describes how fast pollution moves down to the

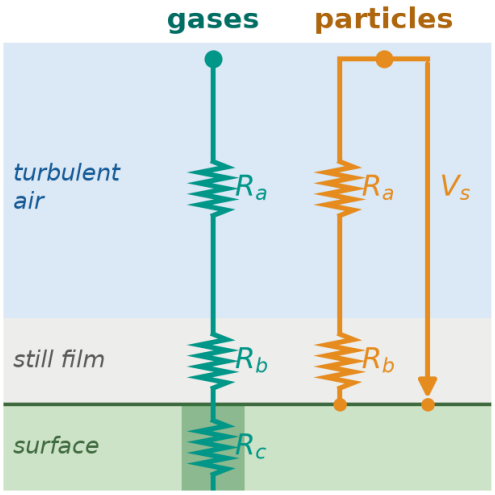

Fig. 6. Resistances to dry deposition. A gas passes through the turbulent air (R_a), the still film of air in contact with the surface (R_b), and uptake by the surface (R_c) in series. A particle passes through R_a and R_b and sticks on contact, or undergoes gravitational settling (V_s).

surface and is taken up by it, and ORBIT removes each species from the lowest layer at the rate $V_d/\Delta z_0$. I compute the deposition velocities of particles, SO_2, NO_x, NH_3, nitric acid, and organic vapors for each 3-hour window of the reanalysis and then average them within each bin (Section 2.1.3), and I report the dry deposition flux of each species as its deposition velocity times its surface concentration. For species that partition between the gas and particle phases, I weight the dry and wet deposition rates of the 2 phases by their partition fractions (Equation 13).

Gases and particles reach a surface in different ways (Figure 6). Dry deposition of gases is modelled by 3 steps in series: turbulent transport through the bulk air, diffusion across the thin film of still air surrounding the surface, and the uptake by the surface itself. The resistances of the 3 steps are additive, giving $V_d = (R_a + R_b + R_c)^{-1}$. Aerodynamic resistance R_a depends on friction velocity and Monin-Obukhov stability. Quasi-laminar resistance is given by $R_b = 5\,\mathrm{Sc}^{2/3}/u_*$ (Seinfeld & Pandis, 2016). The surface resistance R_c follows Wesely Wesely (1989), combining stomatal, mesophyll, cuticular, and ground terms for 5 seasons selected by surface air temperature, and 11 land-use types, mapped from the Olson global land-cover map. I treat NO_x as NO_2, and the surface resistances of SO_2, NH_3, and NO_x differ via their Henry's law constants and reactivity. For organic vapors, I assume their properties match those of formic acid. I use the same resistance model for nitric acid, whose very large Henry's law constant makes its surface resistance practically negligible (Wesely, 1989), so that its deposition velocity is usually controlled by R_a (Zhang, Brook, & Vet, 2003).

Particles do not experience surface resistance, as they are assumed to stick to surfaces on contact. But particles still face aerodynamic and quasi-laminar resistances. Dry deposition of particles is thus modelled with $V_d = (R_a + R_b + R_a R_b V_s)^{-1} + V_s$ (Seinfeld & Pandis, 2016), where V_s is the Stokes-Cunningham settling velocity. Across the quasi-laminar layer, particles reach the surface through Brownian diffusion, impaction, and interception. Deposition velocities can vary by orders of magnitude, depending on particle size. In ORBIT, I make the simplifying assumption that deposition of all aerosols can be represented by a particles with uniform diameter of 0.7 μm and density of 1500 kg m^{-3}.

2.2.7 Wet Deposition

ORBIT models 2 processes that contribute to wet deposition: in-cloud scavenging, where particles become nucleation sites for cloud droplets in which soluble gases also dissolve, and the droplets are then removed as the droplets become precipitation; and below-cloud washout, where the falling precipitation collects particles and absorb gases. For each pathway, particles and gases are treated differently, with the gas and particle rates of each species weighted by their partition fractions (Section 2.2.6).

For in-cloud scavenging, the particles are removed at the rate at which cloud droplets are converted to precipitation (Mari, Jacob, & Bechtold, 2000):

$$\Lambda_{\text{in}} = \min\left(\frac{f_{\text{act}}\, P}{L_{\text{cloud}}},\ \frac{f_{\text{act}}}{\tau_{\text{min}}}\right), \tag{28}$$

where P is the vertical divergence of convective plus large-scale precipitation flux (kg m^{-3} s^{-1}), and L_{cloud} is the cloud liquid water content. I multiply the result by the cloud fraction of each layer. In ORBIT, I assume that the fraction of particle mass in cloud droplets is $f_{\text{act}} = 0.7$, and that the shortest removal time is given by $\tau_{\text{min}} = 300$ s. For below-cloud washout of particles, I follow Simpson et al. (2012), assuming raindrops fall at a terminal velocity of 5 m s^{-1}, which gives a rate proportional to the rainwater content of the layer. In ORBIT, I assume a collection efficiency of 0.1 (EMEP fine-particle value: 0.02).

Soluble gases are scavenged at the rate $WF_{\text{p}}/(h_s \rho_w)$, where F_{p} is the precipitation flux (kg m^{-2} s^{-1}), $h_s = 1000$ m is the characteristic scavenging depth, ρ_w is the density of water, and W is the gas scavenging ratio (Simpson et al., 2012). For layers with cloud water, I use the in-cloud ratio, and for layers below cloud, I use the below-cloud ratio. The ratios are 0.3×10^6 and 0.15×10^6 for SO_2, and 1.4×10^6 and 0.5×10^6 for ammonia and nitric acid. For total reduced nitrogen and total nitrate, I weight the gas rates and the particle rate by the partition fractions. Only the particle fraction of each organic volatility bin is scavenged, as in Simpson et al. (2012), and NO_x is not wet-deposited. I report the wet deposition flux of each species as its scavenging rate times its concentration, integrated over the column.

2.3 ORBIT Chemistry

2.3.1 Transported Species

ORBIT transports 12 chemical species: inert primary $PM_{2.5}$ (which includes black carbon, mineral dust, and sea salt), inert primary organic aerosol (POA), 5 organic volatility bins, and 5 inorganic species (Table 2). Because gas-particle partitioning is nonlinear, ORBIT transports the conserved total of each species that partitions between gas and particle for its periodic solve. These species are reduced nitrogen (NH_3 and NH_4^+), total nitrate (HNO_3 and NO_3^-), and each organic bin (its gas and particle phases). The fractions of each total in the gas and particle phases are determined from the thermodynamic equilibrium (Sections 2.3.5 and 2.3.6), and are used to weight the deposition rates across gas and particle phases (Section 2.2.6). These fractions are then updated in the outer iteration (Section 2.1.5).

ORBIT reports surface $PM_{2.5}$ as the the particulate mass:

$$\text{PM}_{2.5} = c_{\text{PM}} + c_{\text{POA}} + \sum_i F_{p,i}\, c_i + f_{\text{NH}_4}\, c_{\text{NH}}\, \frac{18.04}{14.01} + c_{\text{pSO}_4}\, \frac{96.06}{32.07} + f_{\text{NO}_3}\, c_{\text{NO}_3}\, \frac{62.00}{14.01}, \tag{29}$$

where c_{POA} is primary organic aerosol concentration, c_{PM} is the remaining primary $PM_{2.5}$ concentration, $F_{p,i}$ and c_i are the particle fractions and total concentrations of each organic bin i (Section 2.3.6), f_{NH_4} and

Table 2. The 12 transported chemical species in ORBIT and their chemistry.

Species	Unit	Chemistry
Primary $PM_{2.5}$ (BC, dust, sea salt)	mass	inert
Primary organic aerosol (POA)	mass	inert
Reduced nitrogen (NH_3 + NH_4^+)	N	partitioned by ISORROPIA II
SO_2	S	oxidized to sulfate by OH (gas) and H_2O_2 (aqueous)
NO_x	N	converted to total nitrate by OH + NO_2 (day) and N_2O_5 hydrolysis (night)
Particulate sulfate	S	nonvolatile and enters ISORROPIA II
Total nitrate (HNO_3 + NO_3^-)	N	partitioned by ISORROPIA II
Organic bins, C^* = 1000, 100, 10, 1, 0.1 $\mu g\,m^{-3}$	mass	Pankow partitioning, and aging of the gas phase into the next lower-volatility bin, with a 7.5% mass gain per step

f_{NO_3} are the particle fractions of reduced nitrogen and nitrate (Section 2.3.5). The mass ratios are given to convert from elemental to compound mass. Monthly $PM_{2.5}$ is the mean over the 8 bins. Following the GEOS-Chem $PM_{2.5}$ equation (The International GEOS-Chem User Community, 2026), I apply the growth factors of 1.10 for sulfate, nitrate, and ammonium and 1.05 for POA and SOA, with the factor applied to sea salt during emissions processing.

2.3.2 Oxidants

ORBIT version 1 does not currently implement a module for oxidants. Instead, oxidant concentrations are taken from 3-hourly OH, NO, NO_2, and NO_3 concentrations from the 2016 nested Asian GEOS-Chem (v11-01) simulation on 47 vertical levels described in Thakrar et al. (2022). This means that ORBIT does not model how changes in oxidant concentrations might affect $PM_{2.5}$ as a result of emissions. ORBIT uses the fixed oxidant concentrations to estimate first-order production and loss rates averaged for each time bin in each month, keeping their diurnal variation (which is especially important for nitrate, see Section 2.3.4).

2.3.3 Sulfate

SO_2 is oxidized to sulfate in the gas phase by OH, and in cloud water in the aqueous phase by dissolved oxidants. In the gas phase, the rate constant follows the Troe expression (Burkholder et al., 2020),

$$k_{OH+SO_2} = \frac{k_0(T)[M]}{1 + k_0(T)[M]/k_\infty(T)}\, F_c^{1/[1+(\log_{10}(k_0[M]/k_\infty))^2]}, \tag{30}$$

with $k_0 = 2.9 \times 10^{-31}(T/298)^{-4.1}$ cm^6 molecule^{-2} s^{-1}, $k_\infty = 1.7 \times 10^{-12}(T/298)^{0.2}$ cm^3 molecule^{-1} s^{-1}, and $F_c = 0.6$ (Burkholder et al., 2020). I multiply the rate constant by OH to give the first-order rate.

In cloud water, H_2O_2 oxidizes dissolved SO_2 at the rate $k[H^+][H_2O_2][HSO_3^-]/(1+K[H^+])$, with $k = 7.5\times10^7$ M^{-2} s^{-1} and $K = 13$ M^{-1} at 298 K (Seinfeld & Pandis, 2016). The dissolved concentrations are taken from Henry's law equilibrium with the archived H_2O_2 and the cloud liquid water, with Henry's law constants of 1.23 M atm^{-1} for SO_2 and 1×10^5 M atm^{-1} for H_2O_2, and a first dissociation constant of 1.3×10^{-2} M

for dissolved SO_2 at 298 K (Seinfeld & Pandis, 2016). I assume a cloud-water pH of 3.5. I omit aqueous oxidation by ozone, which at this pH is about 3 orders of magnitude slower than oxidation by H_2O_2 (Seinfeld & Pandis, 2016, Figure 7.19). The SO_2 loss and the sulfate production match, conserving total sulfur mass.

2.3.4 Nitrate

NO_x becomes nitrate via 2 pathways. In the day time, NO_2 reacts with OH, which forms mainly through ozone photolysis, to form nitric acid. In the night time, when the nitrate radical NO_3 is no longer photolyzed, it combines with NO_2 to form N_2O_5, which hydrolyzes on aerosol surfaces to nitric acid.

The daytime rate constant follows Equation 30 with $k_0 = 1.8 \times 10^{-30}(T/298)^{-3}$ cm^6 molecule^{-2} s^{-1}, $k_\infty = 2.8 \times 10^{-11}$ cm^3 molecule^{-1} s^{-1}, and $F_c = 0.6$ (Burkholder et al., 2020). I multiply the rate constant by OH and the NO_2 fraction of NO_x to give the first-order rate.

The GEOS-Chem simulation did not save out N_2O_5, but it can be found assuming equilibrium between NO_2 and NO_3, with equilibrium constant $K_{eq} = 5.8 \times 10^{-27} \exp(10840/T)$ cm^3 molecule^{-1} (Burkholder et al., 2020). Uptake on aerosol is $k_{het} = \frac{1}{4}\gamma\bar{v}A_{fine}$, where $\bar{v}$ is the mean molecular speed of N_2O_5 and A_{fine} is the fine-aerosol surface area. I assume an uptake coefficient of $\gamma = 0.02$ and infer A_{fine} from the GEOS-Chem fine-aerosol mass, with a particle radius of 0.15 μm and a density of 1700 kg m^{-3}. The fine-aerosol surface area remains fixed during the outer iteration. Each hydrolysis reaction converts 2 nitrogen atoms into nitrate, giving $k_{night} = 2k_{het}[N_2O_5]/[NO_x]$. OH is approximately zero at night, and archived NO_3 falls to its lowest level by day, so each pathway is effectively switched off while the other operates. I apply the sum of their rates, $k_{day} + k_{night}$, to get the loss of NO_x and the production of total nitrate.

2.3.5 Inorganic Equilibrium

After the species are transported, I partition the reduced nitrogen and total nitrate between their gas and particle phases by inorganic thermodynamic equilibrium, which depends on the concentration of sulfate, the other cations, temperature, and relative humidity. Ammonium nitrate condenses when the air is cold and humid, and evaporates when it is warm and dry. Sulfate, which is not volatile, takes up ammonia first, so a reduction in sulfate can increase particulate nitrate. Following the partition, the particulate ammonium and nitrate contributes to $PM_{2.5}$ (Equation 29) and the partition fractions go on to weight the removal rates of the 2 totals in the next iteration (Section 2.2.6).

The partitioning is found using the metastable branch of ISORROPIA II (in which particles are always aqueous) (Fountoukis & Nenes, 2007) as implemented in HETP, a solver based on ISORROPIA II (Miller, Makar, & Lee, 2024). To avoid solving for every cell, bin, and outer iteration during simulation time, I precompute it as a 7-dimensional lookup table with sulfate (25 logarithmic points, 0.001–80 μg m^{-3} as S), reduced nitrogen and nitrate (25 points each, 0.001–100 μg m^{-3} as N), calcium and sodium (6 values each, 0–5 μg m^{-3}), temperature (20 points, 200–320 K), and relative humidity (20 points, 0.01–0.99). I assume calcium to be 0.036 of the GEOS-Chem fine dust mass, with potassium and magnesium in fixed molar ratios to calcium. Sodium is 0.306 of the archived GEOS-Chem fine sea salt, and chloride is 1.17 times sodium by mole, following the composition of sea salt (Seinfeld & Pandis, 2016). Currently, the dust and sea salt concentrations are fixed at the averages from the GEOS-Chem baseline simulation, so dust and sea salt emissions in ORBIT do not change the equilibrium. Of the 225 million table entries, I load and linearly interpolate only the ammonium and nitrate particle fractions (1.8 GB).

2.3.6 Secondary Organic Aerosol

from vegetation, biomass burning, and industry, yields less volatile products that condense into the organic part of the aerosol, both primary organic aerosol and SOA, and in ORBIT not into the inorganic particles. Further oxidation of the products by OH, known as aging, lowers their volatility further.

Secondary organic aerosol (SOA) forms from the oxidation of non-methane volatile organic compounds (VOCs) emissions. ORBIT groups the VOCs into 5 volatility bins, with saturation concentrations C^* = 0.1, 1, 10, 100, and 1000 $\mu g\,m^{-3}$ (Donahue, Robinson, Stanier, & Pandis, 2006). The saturation concentration is the organic aerosol mass at which half of the product is in the particle phase (Equation 31), so a lower C^* means the product is less volatile. I assume that C^* does not vary with temperature and fix each C^* at its value for 298 K.

Instead of transporting parent VOCs, I distribute their first-generation products into the volatility bins at the time of emission. Isoprene and monoterpenes react with OH within hours (Hodzic et al., 2016), however the approximation may be less accurate for slower-reacting precursors (such as aromatics, which react over days), increasing their concentrations closer to their sources than they are actually formed. I map 5 classes of parent VOCs into volatility bins with published mass yields (Table 3). The 5 classes are: anthropogenic VOCs, monoterpenes, isoprene, biomass-burning VOCs, and intermediate-volatility compounds. I split biogenic VOC into 0.14 monoterpene and 0.86 isoprene by mass, following the global totals of Sindelarova et al. (2014). Each class has high- and low-NO_x yields, because the peroxy radicals formed by VOC oxidation react mainly with NO when NO_x is high and with other peroxy radicals when NO_x is low, giving different volatility products (Tsimpidi et al., 2010). Following Hodzic et al. (2016), I interpolate between the 2 sets, with the fraction of peroxy radicals that react with NO, $\beta = k_{NO}[NO]/(k_{NO}[NO] + k_{HO_2}[HO_2])$, calculated from the surface NO and OH in each 3-hour bin of the day. The archive has no HO_2, so I assume $[HO_2]$ is 50 times [OH].

I calculate the particle fraction of each volatility bin via absorptive partitioning (Pankow, 1994),

$$F_{p,i} = \frac{1}{1 + C_i^*/M_{OA}}, \qquad M_{OA} = c_{POA} + \sum_i F_{p,i} c_i. \tag{31}$$

Primary organic aerosol is also part of the absorbing mass. Because M_{OA} depends on the particle fractions that it sets, I solve Equation 31 in the outer iteration (Section 2.1.5) by Picard iteration, calculating the fractions from the current M_{OA} and then recalculating M_{OA} until the relative change falls below 0.1%. The first solve starts from 10 $\mu g\,m^{-3}$, and later ones start from the preceding value, taking 3–5 updates.

The gas phase of each volatility bin is aged by OH at $k = 4 \times 10^{-11}$ cm^3 $molecule^{-1}$ s^{-1}, which moves it into the next lower volatility bin with 1.075 times its mass, a 7.5% gain for added oxygen (Tsimpidi et al., 2010). No further ageing occurs in the lowest volatility bin. I report SOA as $\sum_i F_{p,i} c_i$.

2.4 ORBIT Health Impacts

ORBIT implements 3 concentration-response relationships to estimate excess deaths attributable to $PM_{2.5}$ concentrations, using data on population and baseline mortality. The results in the main text use the Global Exposure Mortality Model (GEMM) relationship for 5 causes (ischemic heart disease, stroke, chronic obstructive pulmonary disease, lower respiratory infection, and lung cancer) for adults aged 25 and over, including the parameters fitted from the Chinese cohort (R. Burnett et al., 2018). For GEMM, the counterfactual concentration is 2.4 $\mu g\,m^{-3}$, which is the lowest observed exposure in its underlying 41 cohorts, under which it is assumed that there is no change in mortality risk (R. Burnett et al., 2018).

Table 3. First-generation mass yields of the organic aerosol volatility bins by parent VOC class and NO_x regime. Anthropogenic, monoterpene, and isoprene yields are the ARO1, TERP, and ISOP sets from Tsimpidi et al. (2010). Biomass-burning yields weight the Tsimpidi et al. (2010) species sets by the mass shares from 2022 GFED5.1 VOC over the domain. Intermediate-volatility yields are from Hodzic et al. (2016), with products at $\log_{10} C^* \le -1$ summed into the 0.1 $\mu g\,m^{-3}$ volatility bin.

VOC class	NO_x	$C^* = 0.1$	1	10	100	1000 $\mu g\,m^{-3}$
Anthropogenic	high	0	0.003	0.165	0.300	0.435
	low	0	0.075	0.225	0.375	0.525
Monoterpene	high	0	0.012	0.122	0.201	0.500
	low	0	0.107	0.092	0.359	0.600
Isoprene	high	0	0.001	0.023	0.015	0
	low	0	0.009	0.030	0.015	0
Biomass burning	high	0	0.005	0.112	0.187	0.346
	low	0	0.063	0.138	0.266	0.423
Intermediate volatility	high	0.276	0.069	0.019	0.010	0.012
	low	0.488	0.046	0.010	0.007	0.008

ORBIT also implements the GBD 2019 risk curves, which follow the integrated exposure-response approach R. T. Burnett et al. (2014) by fitting splines to studies of ambient air pollution, household solid-fuel use, and secondhand smoke. The GBD curves treat the counterfactual concentration as uncertain, uniformly distributed between 2.4 and 5.9 $\mu g\,m^{-3}$ (Global Burden of Disease Collaborative Network, 2021). Furthermore, ORBIT implements the GEMM concentration-response function for noncommunicable diseases and lower respiratory infections (NCD+LRI) combined. However, GBD does not publish baseline mortality rates for the combined group, so I apply this function to the all-cause mortality rates, which include deaths from other causes and so overstate the burden.

Baseline mortality rates are observed rates for whole regions, so they already include the deaths caused by pollution. I attribute deaths to cells in proportion to their excess risk, so that the cells of each region sum to the region's total. In cell i of region L,

$$D_i = \sum_a \sum_j P_{a,i}\, y_{0,a,j,L}\, \frac{\mathrm{RR}_j(C_i) - 1}{\widehat{\mathrm{RR}}_{a,j,L}}, \qquad \widehat{\mathrm{RR}}_{a,j,L} = \frac{\sum_{k \in L} P_{a,k}\, \mathrm{RR}_j(C_k)}{\sum_{k \in L} P_{a,k}}, \tag{32}$$

where $P_{a,i}$ is the population in age group a, $y_{0,a,j,L}$ is the mortality rate for cause j for the region, $\mathrm{RR}_j(C)$ is the relative risk at concentration C, and $\widehat{\mathrm{RR}}_{a,j,L}$ is its population-weighted mean over the region. The cells of each region then sum to $P_{a,L}\, y_{0,a,j,L}(1 - 1/\widehat{\mathrm{RR}}_{a,j,L})$. Population and age structure use WorldPop's 2023 constrained, UN-adjusted 1 km total and age- and sex-specific estimates, release R2025A (Bondarenko et al., 2025b, 2025a), aggregated to 0.1°. Cause- and age-specific rates for adults aged 25 and over are from the Global Burden of Disease Study 2023 for subnational regions (including states in India, and provinces in Pakistan) (Global Burden of Disease Collaborative Network, 2025).

Because the risk curves rise steepest at low concentrations, an additional 1 $\mu g\,m^{-3}$ raises the relative risk more in clean air than in polluted air. The marginal emissions damages thus depend on the total concentration at each location. I evaluate the totals using the annual-mean surface $PM_{2.5}$ of the ORBIT forward simulation, which is consistent with the marginal calculation. Differentiating the regional totals gives the deaths gradient:

$$\frac{\partial D}{\partial C_i} = \sum_a \sum_j P_{a,i}\, y_{0,a,j,L}\, \frac{\mathrm{RR}_j'(C_i)}{\widehat{\mathrm{RR}}_{a,j,L}^{\,2}}, \tag{33}$$

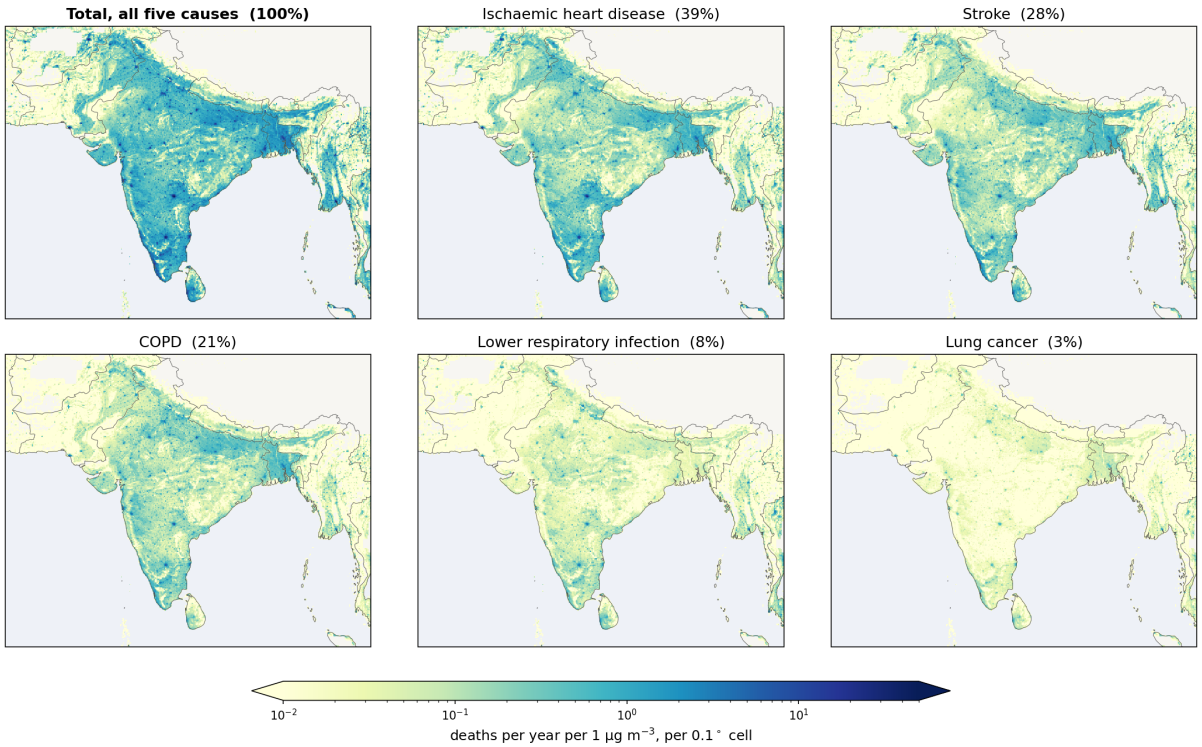

Fig. 7. The number of deaths per year per 1 μg m^{-3} of $PM_{2.5}$ in each 0.1° cell, for the 5 causes of the GEMM function and their sum. A uniform reduction of 1 μg m^{-3} across the domain corresponds to about 29,200 deaths per year, of which two-thirds are ischemic heart disease and stroke.

which I aggregate from the 0.1° grid to the model grid (Figure 7). ORBIT supports Monte Carlo draws of the response parameters and baseline rates (the central values are used in the gradient provided in the open data of the ORBIT release). The marginal deaths from a unit of emissions are given by the gradient applied to the concentration change caused by those emissions. I use the gradient as the receptor weight $\mathbf{S}_k$ of the adjoint, to give the marginal deaths from every source at once (Section 2.1.6). Subdistrict sensitivities use GADM level 2 polygons (GADM, 2022) across the full ORBIT grid, area-weighted by the intersections of the polygons over all cells.

I report the annual deaths associated with 1000 kg of each emitted compound, such as SO_2 and NH_3, converted from the elemental sulfur and nitrogen transported by ORBIT, emitted uniformly through a year or through a specified month. The response functions relate annual deaths to annual average exposure. For a 1 month pulse of emissions, the changes in the annual mean are given by its effect on the concentrations in that month (neglecting any pollution carrying into the next month) multiplied by the fraction of the days of the year in that month. The result for 1000 kg emitted uniformly through the year is the day-weighted mean of the 12 monthly results.

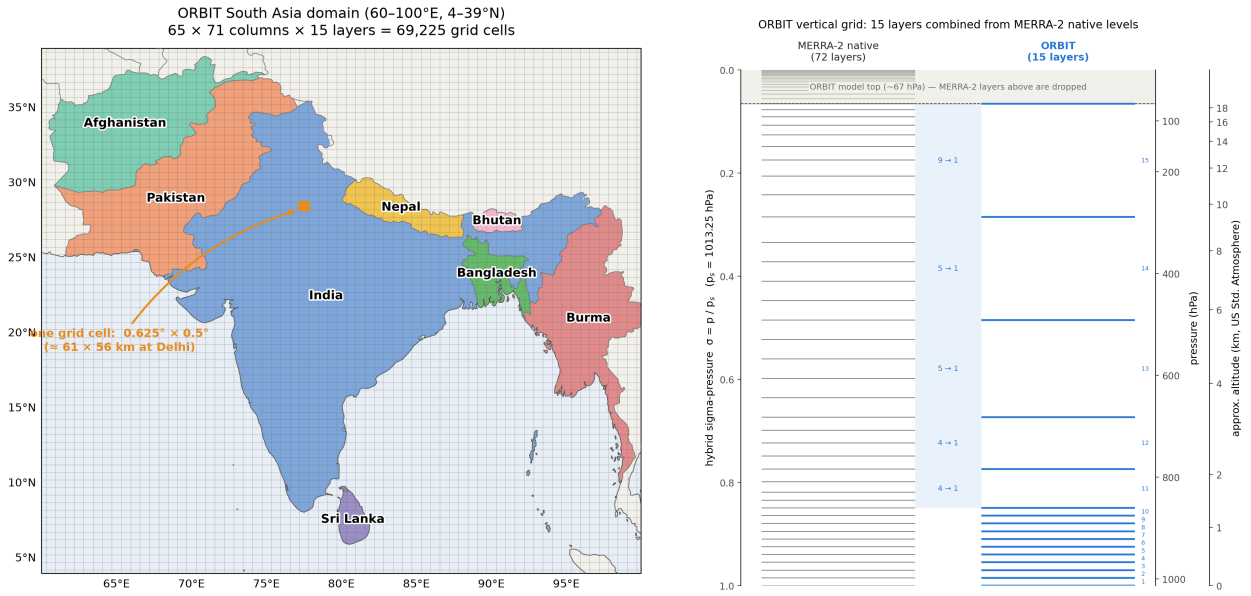

Fig. 8. The South Asian configuration of ORBIT. Left: the domain and the 0.625° × 0.5° horizontal grid, with one cell highlighted over Delhi. Right: the 15 vertical layers, formed from the lowest 38 of the 72 native MERRA-2 levels.

2.5 Implementation for South Asia

2.5.1 Domain and Grid

ORBIT is implemented for South Asia (59.69–100.31°E, 3.75–39.25°N, see Figure 8) for 2022. The domain includes all of Afghanistan, Bangladesh, Bhutan, Nepal, Pakistan, and Sri Lanka, all of India including the Andaman and Nicobar Islands and Lakshadweep, and all of Myanmar except the eastern edge of Shan State. It also contains the Maldivian atolls north of 3.75°N, which amounts to about half its land area. The domain covers most of the Tibetan Plateau, parts of China to its north and east, and (using areas computed from GADM version 4.1 (GADM, 2022)), 81% of Tajikistan, 44% of Thailand, 36% of Turkmenistan, 14% of Iran, 11% of Uzbekistan, and northern Sumatra in Indonesia.

ORBIT uses the native MERRA-2 spacing of 0.5° latitude by 0.625° longitude, which is about 56 by 61 km at Delhi, giving 71 rows and 65 columns (69,225 cells), and 15 vertical layers. In the vertical direction, ORBIT uses the lowest 38 of the 72 native MERRA-2 hybrid sigma-pressure levels. The lowest 10 layers (up to 1.3 km on average) are kept at native resolution. The top of the model domain is 66.6 hPa (about 18 km), just above the tropical tropopause.

2.5.2 Meteorology

Like many atmospheric models, ORBIT is driven by offline (i.e., pre-generated) meteorology, which describes the winds, mixing, convection, clouds, precipitation, temperature, humidity, and radiation on which its physics and chemistry depend. ORBIT uses the MERRA-2 reanalysis (Gelaro et al., 2017), in the version prepared for GEOS-Chem, at 0.5° latitude by 0.625° longitude, the highest native global resolution at which MERRA-2 is available. The fields required for transport and mixing are the horizontal winds, surface pressure, boundary-layer height, sensible heat flux, roughness length, and friction velocity, together with the profiles of temperature and wind that set the stability of the air (Section 2.2.2). Roughness length and friction

velocity are also used for deposition. The fields required for convection and removal are the convective updraft mass flux, cloud fraction and cloud liquid water, and the convective and large-scale precipitation fluxes. The fields required for chemistry and deposition are temperature, specific humidity, and downward shortwave radiation at the surface. MERRA-2 provides the single-level fields as hourly averages, most other fields as 3-hourly averages, and temperature, surface pressure, and specific humidity as instantaneous values every 3 hours, at the start and end of each 3-hour window. I therefore average the hourly fields over each 3-hour window, and I take the mean of the instantaneous values at the start and end of each window, so that every field describes the same 3-hour window, before applying the methods of Section 2.1.3.

2.5.3 Emissions

To evaluated ORBIT and provide a baseline forward simulation, I compile total emissions of $PM_{2.5}$ and precursors for South Asia, using openly available global emission inventories (Table 4), to which I assign injection heights and diurnal and seasonal profiles.

Anthropogenic NO_x, SO_2, NH_3, VOCs, black carbon, and organic carbon are from CEDS v_2025_04_18 (R. M. Hoesly et al., 2018; R. Hoesly et al., 2025), as provided for GEOS-Chem as monthly 0.5° fields for 8 different sectors. Organic carbon is represented as POA, with the ratio of organic mass to organic carbon of 1.8, which has been conventionally used for IMPROVE measurements (which may overstate the mass during winter Malm, Schichtel, and Pitchford (2011)). Because they are not provided by CEDS, I take non-carbonaceous primary $PM_{2.5}$(such as fly ash and process dust) from EDGAR v8.1 (Crippa et al., 2024) as $PM_{2.5}$ minus black carbon and 1.8 times organic carbon in each sector, assigned to the CEDS stack heights described below. Fugitive road and construction dust are not included in EDGAR so these sources are missing from the ORBIT baseline simulation.

Biomass burning emissions are given by GFED5.1 (Chen et al., 2023; van der Werf et al., 2025), biogenic VOCs are given by CAMS-GLOB-BIO (Sindelarova et al., 2022), soil NO_x is given by the CAMS-GLOB-SOIL climatology for year 2000 (Simpson & Darras, 2021), oceanic dimethyl sulfide is included as $SO_2$from CAMS-GLOB-OCE (Granier et al., 2019), volcanic SO_2 is given by Carn, Fioletov, McLinden, Li, and Krotkov (2017), and natural NH_3 is given by GEIA (Bouwman et al., 1997). VOCs are mapped to volatility bins, as described in Section 2.3.6, with intermediate-volatility emissions equal to 1.5 times POA (Robinson et al., 2007).

Dust and sea-salt emissions are given by MERRA-2 as fluxes in different size bins. Following GEOS-Chem and MERRA-2 (Yang, Tan, Rajapakshe, Chin, & Yu, 2022), for dust I assume $PM_{2.5}$ includes all of the first size bin (geometric diameter 0.2–2.0 μm) and 0.38 of the second (2.0–3.6 μm). For fine sea salt, I assume $PM_{2.5}$ includes the 2 smallest size bins and 0.48 of the third, and multiply by 1.86 (following GEOS-Chem) (The International GEOS-Chem User Community, 2026).

Emissions from agriculture, transportation, the residential and commercial sector, solvents, shipping, anthropogenic VOCs, and ammonia are all emitted at the surface. Primary $PM_{2.5}$, SO_2, and NO_x from waste, industry, and energy are elevated.

ORBIT assumes 3 different heights (30, 50, and 220 m), with stack diameters (1.5, 3, and 6 m), exit speeds (10, 12, and 20 m s^{-1}), and temperature (380, 380, and 410 K) informed by stack-height regulations for India for the energy sector. The buoyancy flux is given by $F = gV_s d^2(T_s - T_a)/(4T_s)$ for exit speed V_s, stack diameter d, stack exit temperature T_s and ambient temperature T_a (U.S. Environmental Protection Agency, 1995). Under neutral and unstable conditions, I use the ASME formula for plume rise, $\Delta h = 7.4\,(F h_s^2)^{1/3}/\bar{u}$, where h_s is the stack height and $\bar{u}$ is the wind speed at stack height (American Society of Mechanical Engineers, 1973), as given by Seinfeld and Pandis (2016, Table 18.4). Under stable conditions, I use the Briggs formula

Table 4. Emission inventories used in the 2022 ORBIT baseline simulation for South Asia, and their totals. Precursor masses are given as NO_2, SO_2, and NH_3.

Source	Inventory	Species and domain totals (Tg yr^{-1})
Anthropogenic	CEDS v_2025_04_18	POA 6.6, black carbon 1.2, NO_x 14.0, SO_2 12.2, NH_3 15.1, NMVOC 5.5
Anthropogenic	EDGAR v8.1	non-carbonaceous $PM_{2.5}$ 1.8
Biomass burning	GFED5.1	$PM_{2.5}$ 2.2, NO_x 1.7, SO_2 0.3, NH_3 0.2, NMVOC 0.3
Windblown dust	MERRA-2	fine dust 10.0
Sea salt	MERRA-2	fine sea salt 26.9
Biogenic VOC	CAMS-GLOB-BIO	26.2
Soil NO_x	CAMS-GLOB-SOIL	3.8
Oceanic DMS	CAMS-GLOB-OCE	0.8 (as SO_2)
Volcanic SO_2	Carn et al. (2017)	0.06
Natural NH_3	GEIA	1.2

$\Delta h = 2.6\,[F/(\bar{u}s)]^{1/3}$, with stability parameter $s = (g/\theta)\,\partial\theta/\partial z$ (U.S. Environmental Protection Agency, 1995). The plume is injected into the layer containing $h_s + \Delta h$, and the wind, temperature, and stability of each 3-hour bin determine the injection layer in that bin.

Daily emissions are distributed across the diurnal profile, with multipliers of each hour resampled to the 3-hour bins in UTC with an offset of 5.5-hours across the domain to account for local time. Anthropogenic emissions follow the sector profiles in EDGAR Crippa et al. (2020). Surface primary $PM_{2.5}$, POA, NO_x, and VOC follow the profile for road transport, SO_2 and the industry and energy elevated emissions follow the profile for industry, and surface primary $PM_{2.5}$ and POA follow the profile for residential combustion in December to February, when residential heating dominates the Indo-Gangetic Plain emissions. Agricultural NH_3, which accounts for nearly all of anthropogenic NH_3, follows surface air temperature in each cell and 3-hour bin using a Henry's-law factor Zhu et al. (2015), normalized to a daily mean of 1 (not including the aerodynamic-resistance factor). Natural NH_3 follows the profile for agricultural soils in EDGAR. The biomass burning profile matches the observed peak in the afternoon (Mu et al., 2011). Soil NO_x keeps its 3-hourly factors given by the climatology (Simpson & Darras, 2021). Dust, sea salt, volcanic, and oceanic emissions are assumed to be constant through the day.

2.5.4 Evaluation

The gold standard for evaluating air quality models is to compare modeled total pollutant concentrations with ground observations. I use the hourly 2022 measurements of India's CPCB Continuous Ambient Air Quality Monitoring Stations (Central Pollution Control Board (CPCB), 2022) to obtain monthly and annual average $PM_{2.5}$ concentrations. I keep only station-months with at least 10 days of measurements, and only stations with at least 9 such months in the year, which retains 290 stations in 129 distinct ORBIT grid cells. Because ORBIT does not attempt to model variation within a grid cell, I average the station measurements within each cell to evaluate ORBIT, and I weight the evaluation statistics by the number of stations. I did not have access to measurements for countries in the domain other than India, or of $PM_{2.5}$ species, that met these criteria.

To better evaluate areas that are not well covered by monitoring networks, I also evaluate ORBIT against

the India-wide satellite-derived surface $PM_{2.5}$ of Kawano et al. (2025), which gives daily concentrations at 10 km horizontal resolution. As a derived product, the 10 km surface likely has greater uncertainties than ground measurements. I average the satellite surface to the model grid, excluding pixel-months with fewer than 5 valid days, model cell-months with fewer than 5 valid pixels, and annual means with fewer than 9 valid months.

Standard benchmarks for evaluating air quality models are provided by Emery et al. (2017). To compare with these benchmarks, I report the correlation R between modeled and observed concentrations, the slope and intercept of an ordinary least-squares regression, the normalized mean bias NMB = $\sum(M - O)/\sum O$, and the normalized mean error NME = $\sum|M - O|/\sum O$.

3 Results

3.1 Model performance

The gold standard of assessing air quality models is against ground monitoring networks. Using standard emissions inventories, ORBIT predicts 2022 annual $PM_{2.5}$ concentrations (at 35% relative humidity) with R: 0.76, NME: 21.6%, NMB: −4.2% (see Fig. 9b). Against monthly pooled measurements (n = 1,667), the same ORBIT simulation has R: 0.77, NME: 34.0%, NMB: −4.7%. Performance variation by month may result from incorrect emissions inventories.

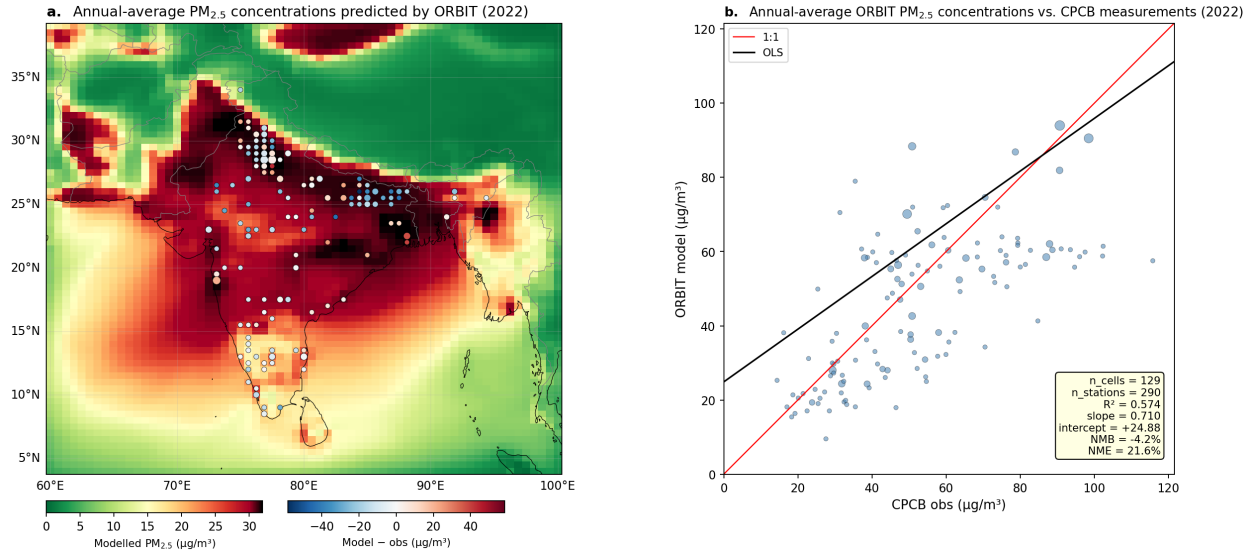

Fig. 9. **(a)** Annual average surface $PM_{2.5}$ concentrations simulated by ORBIT for 2022, with 290 monitor locations (CPCB CAAQM) overlaid and colored by their bias (blue: under-prediction; red: over-prediction). **(b)** Comparison of modeled vs. measured annual $PM_{2.5}$ concentrations.

3.2 Baseline simulation results

Total $PM_{2.5}$ concentrations in the baseline ORBIT simulation for 2022 are shown in Fig. 9a. Annual deaths associated with exposure to total $PM_{2.5}$ are 1,891,000 (95% CI: 1,790,000 – 1,992,000) using the GEMM

5 cause-of-death concentration-response function, which I report as our baseline*. ORBIT results are in a similar range to previous work (McDuffie et al., 2021) for year 2019 (829,000-2,203,000) using the same concentration-response functions. ORBIT shows good model performance, meeting the benchmark goals for correlation, error, and bias (see Section 3.1).

ORBIT predicts monthly average $PM_{2.5}$ concentrations (see Fig. 10), showing an increase in pollution in winter (Dec–Feb) over the Indo-Gangetic Plain, and a drop in pollution across the domain during the monsoon season (Jun–Sep). The seasonality is driven by both emissions inputs and ORBIT's atmospheric chemistry and physics. For example, large pollutant concentrations in Bangladesh and Myanmar in March are driven by biomass burning emissions.

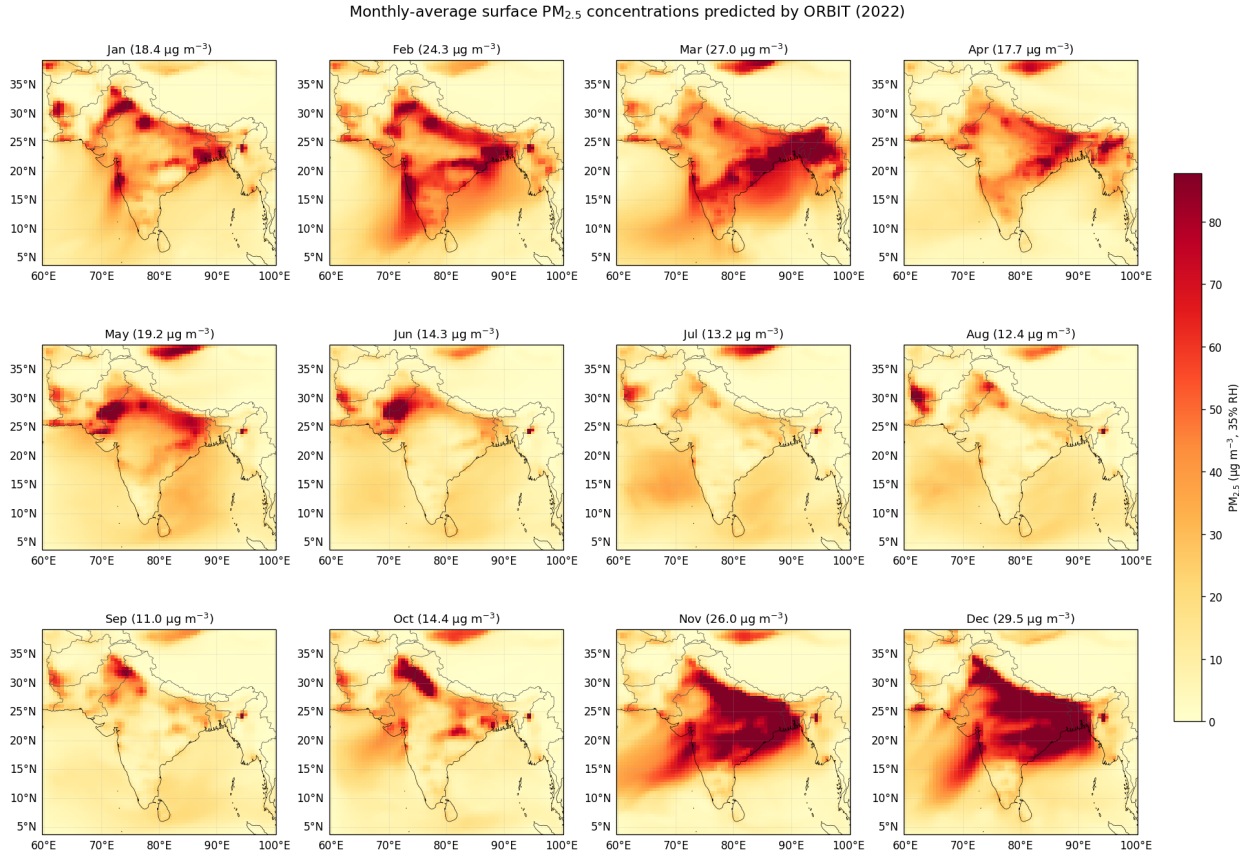

Fig. 10. Monthly average surface $PM_{2.5}$ concentrations in 2022, as simulated by ORBIT.

Speciated annual average $PM_{2.5}$ concentrations are shown in Fig. 11, showing large primary $PM_{2.5}$ concentrations in the western Indo-Gangetic Plain (dust storms from the Thar Desert), large particulate sulfate in Chhattisgarh and Odisha (coal-fired power generation), and large SOA concentrations in Punjab and Northeastern India (biomass burning).

ORBIT also predicts diurnal concentrations for each month, as illustrated for January in Fig. 12. On average across the domain, ORBIT-predicted pollution is highest from night through morning (21:00 – 09:00 IST), and is at its lowest in the afternoon (12:00 – 18:00). The diurnal cycle reflects both ORBIT's chemistry and physics (such as boundary-layer collapse) and the profile of sectoral emissions, as specified by the baseline

*Results are 1,461,000 (1,410,000–1,512,000) deaths using the GBD 2019 risk curves, and 2,961,000 (2,773,000–3,159,000) deaths using the GEMM all-cause function

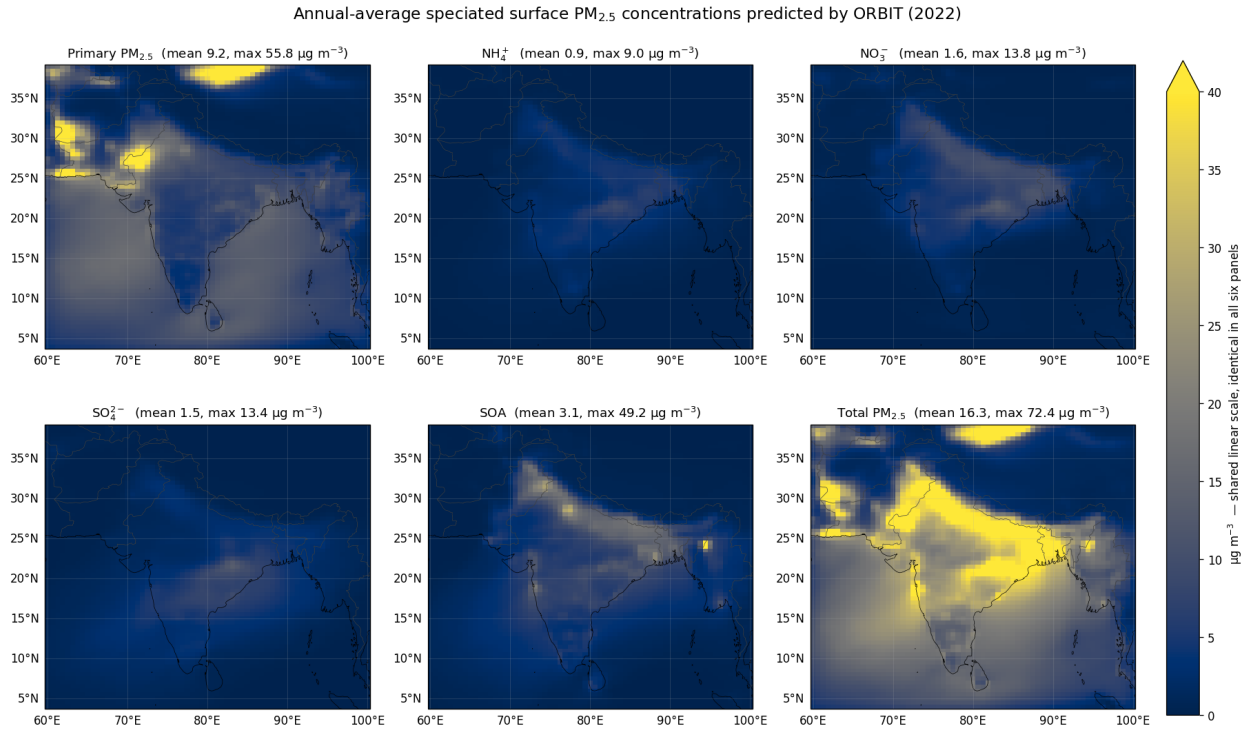

Fig. 11. Annual-average surface concentrations of $PM_{2.5}$ by species for year 2022, as predicted by ORBIT.

simulation. This level of sub-daily resolution is generally not available in reduced-complexity models, and is possible because of ORBIT's periodic-orbit formulation.

3.3 Marginal health damages

ORBIT's marginal mode is able to rapidly estimate health damages from emissions of each primary pollutant, from each grid cell, for each time bin in every month. Here, I illustrate how the deaths associated with annual emissions of 1,000 kg of primary $PM_{2.5}$ (Fig. 13) and SO_2 (Fig. 14) vary by where the pollution is emitted. Deaths are estimated using the GEMM 5 cause-of-death concentration-response function. The range of deaths per 1,000 kg varies considerably. For example, emitting primary $PM_{2.5}$ in Narsingdi, Bangladesh, one of the subdistricts where emissions do the most damage (with Yangon, Myanmar), is 630× more damaging than emitting the same mass of primary $PM_{2.5}$ in the Nicobar Islands (remote Indian islands in the Bay of Bengal), and 760× more damaging than emitting it in Lunana, Bhutan (high in the Himalayas).

Marginal damages per 1,000 kg of SO_2 emissions (Fig. 14) are generally lower than for Primary $PM_{2.5}$, and exhibit a similar spatial pattern (as they are both governed by downwind population), but show differences (owing to, for example, differences in SO_2 oxidation rates, available NH_3, and differences in relative deposition rates of SO_2 versus primary $PM_{2.5}$). There are a small number of places and months (*e.g.*, Kolar in July) where ORBIT predicts that increases in SO_2 above the baseline can *reduce* pollution-related deaths, because in ammonia-poor conditions, added SO_2 can reduce particulate nitrate formation that more than offsets the particulate sulfate, by mass. ORBIT is able to capture these cross-coupling effects that linear models cannot.

ORBIT can capture differences in marginal damages across time as well as space. Fig. 15 shows how deaths associated with 1,000 kg of primary $PM_{2.5}$ emissions change depending on the month of emission, across the

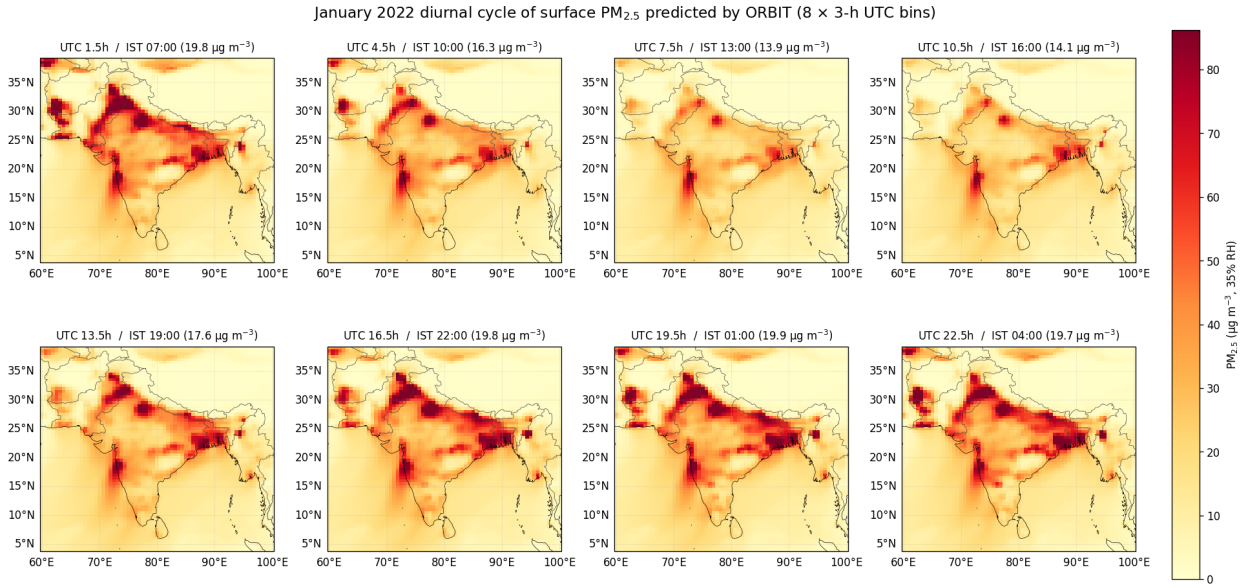

Fig. 12. Diurnal cycle of surface $PM_{2.5}$ concentrations, as predicted by ORBIT, for January 2022.

10 most populous South Asian cities. Emissions are weighted uniformly over each metropolitan area, and assumed to have a uniform rate within each month. Many cities have their highest damages from pollution emitted in the winter, and their lowest from pollution emitted during the monsoon season (June–September). However, the months with the highest and lowest damages vary considerably across cities. For Dhaka, pollution is over 2× more damaging for health when emitted in December than in April. By providing detailed marginal damages across time of day, place, season, and primary pollutant, ORBIT can rapidly inform when and where to target emissions to efficiently improve public health.

3.4 Computational requirements

- Forward simulation: ~23 min walltime and ~3.6 GB per simulated month on 8 CPU cores. Each month can be run in parallel on a cluster, or a full year can run within 5 hours on a laptop.
- Marginal simulation: ~5 min and ~5.5 GB per perturbation per month, enabling hundreds of source × pollutant scenarios.

3.5 Code & data availability

ORBIT is open-source. Code is available at `https://github.com/SumilThakr/orbit` and all data required to run the model is described in the GitHub model documentation.

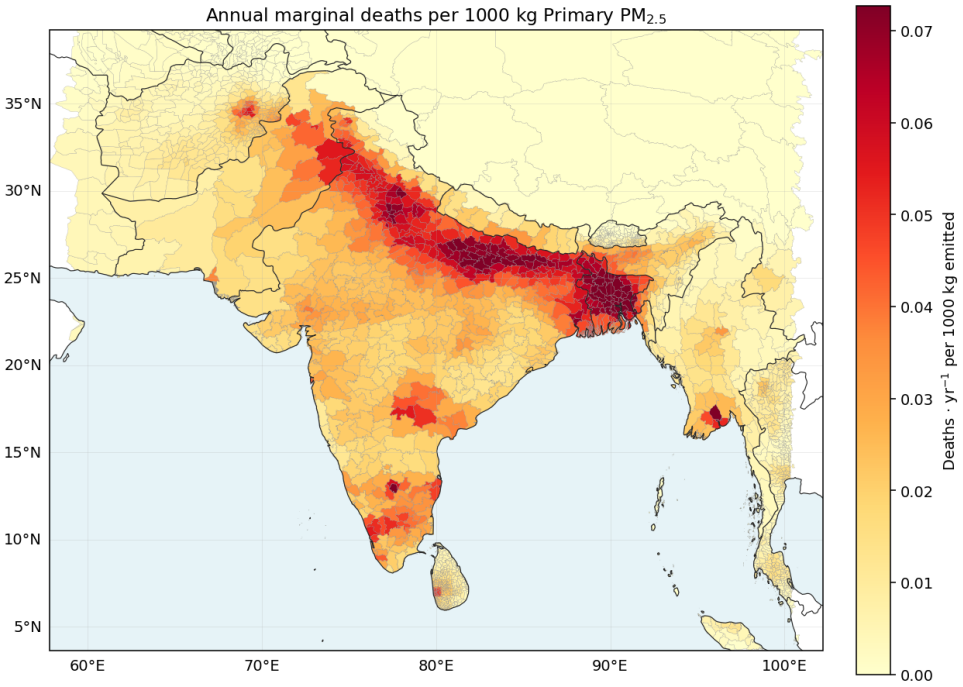

Fig. 13. Annual marginal deaths from continuous emissions of 1,000 kg yr^{-1} primary $PM_{2.5}$, released from every subdistrict in the domain. Deaths can occur outside the subdistrict where emissions occurred.

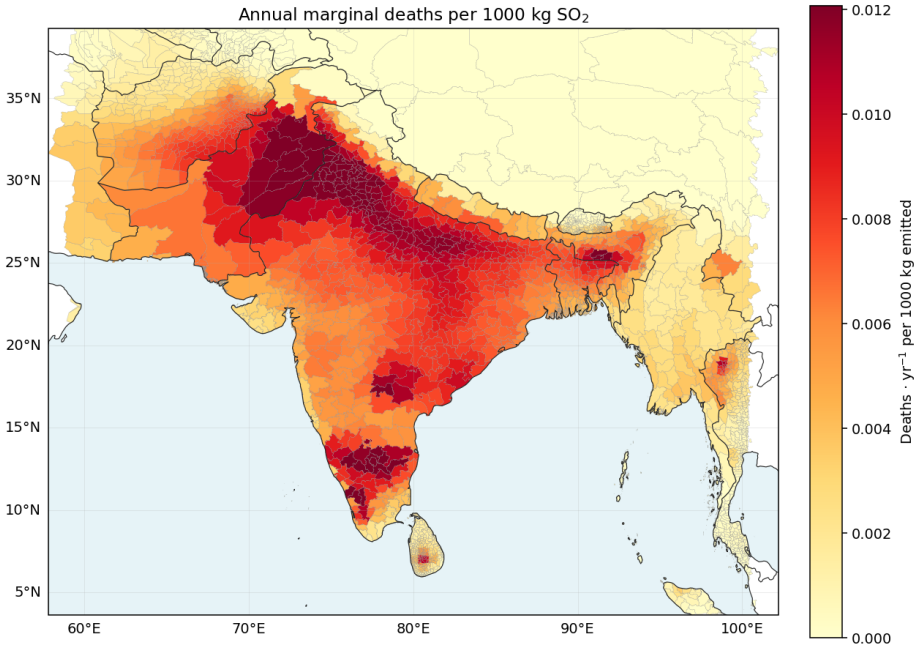

Fig. 14. Annual marginal deaths from continuous emissions of 1,000 kg yr^{-1} SO_2, released from every subdistrict in the domain. Deaths can occur outside the subdistrict where emissions occurred.

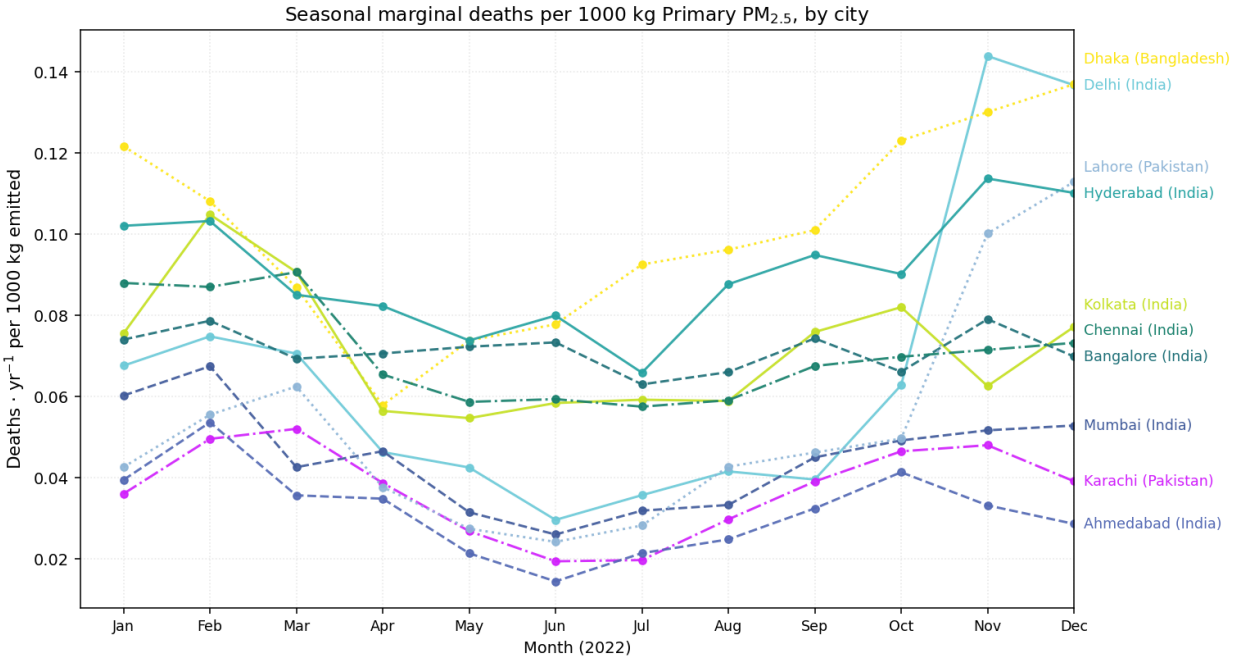

Fig. 15. Seasonal dependence of the marginal health impacts from 1,000 kg yr^{-1} of primary $PM_{2.5}$ emissions, across the 10 most populous South Asian metropolitan areas, emitted at each calendar month of 2022. Hue varies with longitude (from west to east: purple, blue, green, yellow).

References

American Society of Mechanical Engineers. (1973). *Recommended guide for the prediction of the dispersion of airborne effluents* (2nd ed.). New York: American Society of Mechanical Engineers.

Anderson, D. G. (1965). Iterative procedures for nonlinear integral equations. *Journal of the ACM*, *12*(4), 547–560. doi: 10.1145/321296.321305

Berman, A., & Plemmons, R. J. (1994). *Nonnegative matrices in the mathematical sciences*. Philadelphia: Society for Industrial and Applied Mathematics. doi: 10.1137/1.9781611971262

Bondarenko, M., Priyatikanto, R., Tejedor-Garavito, N., Zhang, W., McKeen, T., Cunningham, A., ... Sorichetta, A. (2025a). *The spatial distribution of population broken down by gender and age groupings in 2015–2030 at a resolution of 30 arc seconds (approximately 1 km at the equator), R2025A version v1*. WorldPop, School of Geography and Environmental Science, University of Southampton. doi: 10.5258/SOTON/WP00846

Bondarenko, M., Priyatikanto, R., Tejedor-Garavito, N., Zhang, W., McKeen, T., Cunningham, A., ... Sorichetta, A. (2025b). *The spatial distribution of population in 2015–2030 at a resolution of 30 arc seconds (approximately 1 km at the equator), R2025A version v1*. WorldPop, School of Geography and Environmental Science, University of Southampton. doi: 10.5258/SOTON/WP00845

Bouwman, A. F., Lee, D. S., Asman, W. A. H., Dentener, F. J., Van Der Hoek, K. W., & Olivier, J. G. J. (1997). A global high-resolution emission inventory for ammonia. *Global Biogeochemical Cycles*, *11*(4), 561–587. doi: 10.1029/97GB02266

Burkholder, J. B., Sander, S. P., Abbatt, J., Barker, J. R., Cappa, C., Crounse, J. D., ... Wine, P. H. (2020). *Chemical kinetics and photochemical data for use in atmospheric studies, evaluation no. 19* (Tech. Rep. No. JPL Publication 19-5). Pasadena, CA: Jet Propulsion Laboratory. Retrieved from `http://jpldataeval.jpl.nasa.gov`

Burnett, R., Chen, H., Szyszkowicz, M., Fann, N., Hubbell, B., Pope, C. A., ... Spadaro, J. V. (2018). Global estimates of mortality associated with long-term exposure to outdoor fine particulate matter. *Proceedings of the National Academy of Sciences*, *115*(38), 9592–9597. doi: 10.1073/pnas.1803222115

Burnett, R. T., Pope, C. A., Ezzati, M., Olives, C., Lim, S. S., Mehta, S., ... Cohen, A. (2014). An integrated risk function for estimating the global burden of disease attributable to ambient fine particulate matter exposure. *Environmental Health Perspectives*, *122*(4), 397–403. doi: 10.1289/ehp.1307049

Businger, J. A., Wyngaard, J. C., Izumi, Y., & Bradley, E. F. (1971). Flux-profile relationships in the atmospheric surface layer. *Journal of the Atmospheric Sciences*, *28*(2), 181–189. doi: 10.1175/1520-0469(1971)028<0181:FPRITA>2.0.CO;2

Byun, D., & Schere, K. L. (2006). Review of the governing equations, computational algorithms, and other components of the Models-3 Community Multiscale Air Quality (CMAQ) modeling system. *Applied Mechanics Reviews*, *59*(2), 51–77. doi: 10.1115/1.2128636

Carn, S. A., Fioletov, V. E., McLinden, C. A., Li, C., & Krotkov, N. A. (2017). A decade of global volcanic SO_2 emissions measured from space. *Scientific Reports*, *7*, 44095. doi: 10.1038/srep44095

Central Pollution Control Board (CPCB). (2022). *Continuous ambient air quality monitoring stations (CAAQMS) data, 2022.* `https://airquality.cpcb.gov.in/ccr/`. Central Pollution Control Board, Ministry of Environment, Forest and Climate Change, Government of India.

Chen, Y., Hall, J., van Wees, D., Andela, N., Hantson, S., Giglio, L., ... Randerson, J. T. (2023). Multi-decadal trends and variability in burned area from the fifth version of the Global Fire Emissions Database (GFED5). *Earth System Science Data*, *15*(11), 5227–5259. doi: 10.5194/essd-15-5227-2023

Crippa, M., Guizzardi, D., Pagani, F., Banja, M., Muntean, M., Schaaf, E., ... Vignati, E. (2024). *EDGAR v8.1 global air pollutant emissions.* `http://data.europa.eu/89h/a3af16e4-21ac-420a-b98c`

-b78a9b7723be. European Commission, Joint Research Centre (JRC).

Crippa, M., Solazzo, E., Huang, G., Guizzardi, D., Koffi, E., Muntean, M., ... Janssens-Maenhout, G. (2020). High resolution temporal profiles in the Emissions Database for Global Atmospheric Research. *Scientific Data*, *7*, 121. doi: 10.1038/s41597-020-0462-2

Davis, T. A. (2004). Algorithm 832: UMFPACK V4.3—an unsymmetric-pattern multifrontal method. *ACM Transactions on Mathematical Software*, *30*(2), 196–199. doi: 10.1145/992200.992206

Donahue, N. M., Robinson, A. L., Stanier, C. O., & Pandis, S. N. (2006). Coupled partitioning, dilution, and chemical aging of semivolatile organics. *Environmental Science & Technology*, *40*(8), 2635–2643. doi: 10.1021/es052297c

Dyer, A. J. (1974). A review of flux-profile relationships. *Boundary-Layer Meteorology*, *7*(3), 363–372. doi: 10.1007/BF00240838

Emery, C., Liu, Z., Russell, A. G., Odman, M. T., Yarwood, G., & Kumar, N. (2017). Recommendations on statistics and benchmarks to assess photochemical model performance. *Journal of the Air & Waste Management Association*, *67*(5), 582–598. doi: 10.1080/10962247.2016.1265027

Floquet, G. (1883). Sur les équations différentielles linéaires à coefficients périodiques. *Annales scientifiques de l'École Normale Supérieure*, *12*, 47–88. doi: 10.24033/asens.220

Fountoukis, C., & Nenes, A. (2007). ISORROPIA II: a computationally efficient thermodynamic equilibrium model for K^+–Ca^{2+}–Mg^{2+}–NH_4^+–Na^+–SO_4^{2-}–NO_3^-–Cl^-–H_2O aerosols. *Atmospheric Chemistry and Physics*, *7*(17), 4639–4659. doi: 10.5194/acp-7-4639-2007

GADM. (2022). *GADM database of global administrative areas, version 4.1.* https://gadm.org.

Gelaro, R., McCarty, W., Suárez, M. J., Todling, R., Molod, A., Takacs, L., ... Zhao, B. (2017). The Modern-Era Retrospective Analysis for Research and Applications, Version 2 (MERRA-2). *Journal of Climate*, *30*(14), 5419–5454. doi: 10.1175/JCLI-D-16-0758.1

Gilmore, E. A., Heo, J., Muller, N. Z., Tessum, C. W., Hill, J. D., Marshall, J. D., & Adams, P. J. (2019). An inter-comparison of the social costs of air quality from reduced-complexity models. *Environmental Research Letters*, *14*(7), 074016. doi: 10.1088/1748-9326/ab1ab5

Global Burden of Disease Collaborative Network. (2021). *Global burden of disease study 2019 (GBD 2019) particulate matter risk curves.* Seattle, WA: Institute for Health Metrics and Evaluation (IHME).

Global Burden of Disease Collaborative Network. (2025). *Global burden of disease study 2023 (GBD 2023) results.* https://vizhub.healthdata.org/gbd-results/. Seattle, WA: Institute for Health Metrics and Evaluation (IHME).

Godunov, S. K. (1959). A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics. *Matematicheskii Sbornik*, *47*(3), 271–306.

Granier, C., Darras, S., Denier van der Gon, H., Doubalova, J., Elguindi, N., Galle, B., ... Sindelarova, K. (2019). *The Copernicus Atmosphere Monitoring Service global and regional emissions (April 2019 version)* (Tech. Rep.). Copernicus Atmosphere Monitoring Service (CAMS) report. doi: 10.24380/d0bn-kx16

Heo, J., Adams, P. J., & Gao, H. O. (2016). Reduced-form modeling of public health impacts of inorganic $PM_{2.5}$ and precursor emissions. *Atmospheric Environment*, *137*, 80–89. doi: 10.1016/j.atmosenv.2016.04.026

Hodzic, A., Kasibhatla, P. S., Jo, D. S., Cappa, C. D., Jimenez, J. L., Madronich, S., & Park, R. J. (2016). Rethinking the global secondary organic aerosol (SOA) budget: stronger production, faster removal, shorter lifetime. *Atmospheric Chemistry and Physics*, *16*(12), 7917–7941. doi: 10.5194/acp-16-7917-2016

Hoesly, R., Smith, S., Ahsan, H., Prime, N., O'Rourke, P., Crippa, M., ... Wang, S. (2025). *CEDS v_2025_04_18 gridded emissions data, 0.5 degree.* Zenodo. doi: 10.5281/zenodo.15127477

Hoesly, R. M., Smith, S. J., Feng, L., Klimont, Z., Janssens-Maenhout, G., Pitkanen, T., ... Zhang, Q. (2018). Historical (1750–2014) anthropogenic emissions of reactive gases and aerosols from the

Community Emissions Data System (CEDS). *Geoscientific Model Development*, *11*(1), 369–408. doi: 10.5194/gmd-11-369-2018

Hong, S.-Y., Noh, Y., & Dudhia, J. (2006). A new vertical diffusion package with an explicit treatment of entrainment processes. *Monthly Weather Review*, *134*(9), 2318–2341. doi: 10.1175/MWR3199.1

Karypis, G., & Kumar, V. (1998). A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on Scientific Computing*, *20*(1), 359–392. doi: 10.1137/S1064827595287997

Kawano, A., Kelp, M., Qiu, M., Singh, K., Chaturvedi, E., Dahiya, S., ... Burke, M. (2025). Improved daily $PM_{2.5}$ estimates in India reveal inequalities in recent enhancement of air quality. *Science Advances*, *11*(4), eadq1071. doi: 10.1126/sciadv.adq1071

Lin, S.-J., & Rood, R. B. (1996). Multidimensional flux-form semi-Lagrangian transport schemes. *Monthly Weather Review*, *124*(9), 2046–2070. doi: 10.1175/1520-0493(1996)124<2046:MFFSLT>2.0.CO;2

Louis, J.-F. (1979). A parametric model of vertical eddy fluxes in the atmosphere. *Boundary-Layer Meteorology*, *17*(2), 187–202. doi: 10.1007/BF00117978

Malm, W. C., Schichtel, B. A., & Pitchford, M. L. (2011). Uncertainties in $PM_{2.5}$ gravimetric and speciation measurements and what we can learn from them. *Journal of the Air & Waste Management Association*, *61*(11), 1131–1149. doi: 10.1080/10473289.2011.603998

Mari, C., Jacob, D. J., & Bechtold, P. (2000). Transport and scavenging of soluble gases in a deep convective cloud. *Journal of Geophysical Research: Atmospheres*, *105*(D17), 22255–22267. doi: 10.1029/2000JD900211

McDuffie, E. E., Martin, R. V., Spadaro, J. V., Burnett, R., Smith, S. J., O'Rourke, P., ... Brauer, M. (2021). Source sector and fuel contributions to ambient $PM_{2.5}$ and attributable mortality across multiple spatial scales. *Nature Communications*, *12*, 3594. doi: 10.1038/s41467-021-23853-y

Miller, S. J., Makar, P. A., & Lee, C. J. (2024). HETerogeneous vectorized or Parallel (HETPv1.0): an updated inorganic heterogeneous chemistry solver for the metastable-state NH_4^+–Na^+–Ca^{2+}–K^+–Mg^{2+}–SO_4^{2-}–NO_3^-–Cl^-–H_2O system based on ISORROPIA II. *Geoscientific Model Development*, *17*(6), 2197–2219. doi: 10.5194/gmd-17-2197-2024

Molod, A., Takacs, L., Suarez, M., & Bacmeister, J. (2015). Development of the GEOS-5 atmospheric general circulation model: Evolution from MERRA to MERRA2. *Geoscientific Model Development*, *8*, 1339–1356. doi: 10.5194/gmd-8-1339-2015

Mu, M., Randerson, J. T., van der Werf, G. R., Giglio, L., Kasibhatla, P., Morton, D., ... Wennberg, P. O. (2011). Daily and 3-hourly variability in global fire emissions and consequences for atmospheric model predictions of carbon monoxide. *Journal of Geophysical Research: Atmospheres*, *116*, D24303. doi: 10.1029/2011JD016245

Muller, N. Z. (2014). Boosting GDP growth by accounting for the environment. *Science*, *345*(6199), 873–874. doi: 10.1126/science.1253506

Pankow, J. F. (1994). An absorption model of the gas/aerosol partitioning involved in the formation of secondary organic aerosol. *Atmospheric Environment*, *28*(2), 189–193. doi: 10.1016/1352-2310(94)90094-9

Patankar, S. V. (1980). *Numerical heat transfer and fluid flow*. Washington, DC: Hemisphere Publishing Corporation.

Pleim, J. E. (2007). A combined local and nonlocal closure model for the atmospheric boundary layer. Part I: Model description and testing. *Journal of Applied Meteorology and Climatology*, *46*(9), 1383–1395. doi: 10.1175/JAM2539.1

Robinson, A. L., Donahue, N. M., Shrivastava, M. K., Weitkamp, E. A., Sage, A. M., Grieshop, A. P., ... Pandis, S. N. (2007). Rethinking organic aerosols: Semivolatile emissions and photochemical aging. *Science*, *315*(5816), 1259–1262. doi: 10.1126/science.1133061

Saad, Y., & Schultz, M. H. (1986). GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, *7*(3), 856–869.

doi: 10.1137/0907058

Seinfeld, J. H., & Pandis, S. N. (2016). *Atmospheric chemistry and physics: From air pollution to climate change* (3rd ed.). Hoboken, NJ: John Wiley & Sons.

Simpson, D., Benedictow, A., Berge, H., Bergström, R., Emberson, L. D., Fagerli, H., ... Wind, P. (2012). The EMEP MSC-W chemical transport model – technical description. *Atmospheric Chemistry and Physics*, *12*(16), 7825–7865. doi: 10.5194/acp-12-7825-2012

Simpson, D., & Darras, S. (2021). *Global soil NO emissions for Atmospheric Chemical Transport Modelling: CAMS-GLOB-SOIL v2.2.* Earth System Science Data Discussions [preprint]. doi: 10.5194/essd-2021-221

Sindelarova, K., Granier, C., Bouarar, I., Guenther, A., Tilmes, S., Stavrakou, T., ... Knorr, W. (2014). Global data set of biogenic VOC emissions calculated by the MEGAN model over the last 30 years. *Atmospheric Chemistry and Physics*, *14*(17), 9317–9341. doi: 10.5194/acp-14-9317-2014

Sindelarova, K., Markova, J., Simpson, D., Huszar, P., Karlicky, J., Darras, S., & Granier, C. (2022). High-resolution biogenic global emission inventory for the time period 2000–2019 for air quality modelling. *Earth System Science Data*, *14*(1), 251–270. doi: 10.5194/essd-14-251-2022

Tessum, C. W., Hill, J. D., & Marshall, J. D. (2017). InMAP: A model for air pollution interventions. *PLOS ONE*, *12*(4), e0176131. doi: 10.1371/journal.pone.0176131

Thakrar, S. K., Tessum, C. W., Apte, J. S., Balasubramanian, S., Millet, D. B., Pandis, S. N., ... Hill, J. D. (2022). Global, high-resolution, reduced-complexity air quality modeling for $PM_{2.5}$ using InMAP (Intervention Model for Air Pollution). *PLOS ONE*, *17*(5), e0268714. doi: 10.1371/journal.pone.0268714

The International GEOS-Chem User Community. (2026). *GEOS-Chem 14.7.0.* Zenodo. doi: 10.5281/zenodo.18510368

Tsimpidi, A. P., Karydis, V. A., Zavala, M., Lei, W., Molina, L., Ulbrich, I. M., ... Pandis, S. N. (2010). Evaluation of the volatility basis-set approach for the simulation of organic aerosol formation in the Mexico City metropolitan area. *Atmospheric Chemistry and Physics*, *10*(2), 525–546. doi: 10.5194/acp-10-525-2010

U.S. Environmental Protection Agency. (1995). *User's guide for the Industrial Source Complex (ISC3) dispersion models, volume II: Description of model algorithms* (Tech. Rep. No. EPA-454/B-95-003b). Research Triangle Park, NC: Office of Air Quality Planning and Standards.

van der Werf, G. R., Randerson, J. T., van Wees, D., Chen, Y., Giglio, L., Hall, J., ... Morton, D. C. (2025). Landscape fire emissions from the 5th version of the Global Fire Emissions Database (GFED5). *Scientific Data*, *12*, 1870. doi: 10.1038/s41597-025-06127-w

Walker, H. F., & Ni, P. (2011). Anderson acceleration for fixed-point iterations. *SIAM Journal on Numerical Analysis*, *49*(4), 1715–1735. doi: 10.1137/10078356X

Wesely, M. L. (1989). Parameterization of surface resistances to gaseous dry deposition in regional-scale numerical models. *Atmospheric Environment*, *23*(6), 1293–1304. doi: 10.1016/0004-6981(89)90153-4

Wilson, R. (2004). Turbulent diffusivity in the free atmosphere inferred from MST radar measurements: a review. *Annales Geophysicae*, *22*(11), 3869–3887. doi: 10.5194/angeo-22-3869-2004

Yang, A., Tan, Q., Rajapakshe, C., Chin, M., & Yu, H. (2022). Global premature mortality by dust and pollution $PM_{2.5}$ estimated from aerosol reanalysis of the Modern-Era Retrospective Analysis for Research and Applications, version 2. *Frontiers in Environmental Science*, *10*, 975755. doi: 10.3389/fenvs.2022.975755

Zalesak, S. T. (1979). Fully multidimensional flux-corrected transport algorithms for fluids. *Journal of Computational Physics*, *31*(3), 335–362. doi: 10.1016/0021-9991(79)90051-2

Zhang, L., Brook, J. R., & Vet, R. (2003). A revised parameterization for gaseous dry deposition in air-quality models. *Atmospheric Chemistry and Physics*, *3*, 2067–2082. doi: 10.5194/acp-3-2067-2003

Zhu, L., Henze, D., Bash, J., Jeong, G.-R., Cady-Pereira, K., Shephard, M., ... Capps, S. (2015). Global evaluation of ammonia bidirectional exchange and livestock diurnal variation schemes. *Atmospheric Chemistry and Physics*, *15*(22), 12823–12843. doi: 10.5194/acp-15-12823-2015