



TARSA v1.0: a lightweight Eulerian forward model for aerosol transport – description, verification, and observation-based evaluation

Konstantin Kuznetsov¹, Oleg Dubovik^{1,2}, Anton Lopatin¹, Elena Lind¹, and David Fuertes¹

¹GRASP Earth, Lille, France

²Laboratoire d'Optique Atmosphérique (LOA), UMR 8518, Université de Lille, CNRS, Lille, France

Correspondence: Konstantin Kuznetsov (konstantin.kuznetsov@grasp-earth.com)

Abstract. Atmospheric retrieval, inverse-modelling, and data-assimilation workflows require transparent transport models that can be evaluated repeatedly and provide consistent sensitivities. We present TARSA v1.0, a lightweight regional three-dimensional Eulerian model for inert tracers and simplified aerosol applications under prescribed ERA5 meteorology. TARSA advances tracer mass mixing ratios with a conservative finite-volume core comprising monotonic advection, vertical turbulent diffusion, sedimentation, and parameterised dry and wet deposition. The default advection scheme uses a TVD-admissible, cubic-Hermite-regularised Koren limiter to reduce selected gradient-error roughness. Analytic, tangent-linear, discrete-adjoint, and automatic-differentiation gradients agree to single precision in a scalar emission-amplitude check. Numerical verification includes mass-conservation, analytical plume, and manufactured-solution tests. Independent evaluation against european and north-american regional experiments (ETEX-1, CAPTEX-1) shows agreement in plume timing and concentration amplitude, with performance broadly comparable to published transport-model results. A Saharan dust-column case during Fennec, driven by prescribed MERRA-2 emissions without externally prescribed dust initial or lateral boundary concentrations, shows moderate agreement with GRASP/POLDER coarse-dust column estimates. TARSA provides a reproducible forward operator for regional transport and observation-space studies, supporting future air-quality applications using satellite aerosol observations. Its prescribed process set and local sensitivity structure support future retrieval and inversion applications, while end-to-end inverse-modelling performance remains to be demonstrated.

1 Introduction

Atmospheric aerosols and reactive trace gases affect air quality, human health, and the climate system's radiative balance. Their concentrations reflect emissions, chemistry, removal, and three-dimensional transport from boundary-layer to synoptic scales. Constraining emissions increasingly requires atmospheric models to be combined with satellite and ground-based observations. Retrieval, inversion, and uncertainty-quantification workflows consequently need forward operators that are robust, transparent, and inexpensive enough for repeated evaluation.

The growth of the NewSpace sector is motivating nanosatellite constellations equipped with multi-angular polarimeters (MAP) for frequent atmospheric observations. The GRASP–AirPhoton Multi-Angle Polarimeter (GAPMAP) concept exem-



plifies this approach: its instruments measure reflected radiation at multiple viewing angles, wavelengths, and polarization states to characterise atmospheric aerosols (Fuertes et al., 2022). Increasing observation frequency motivates fast transport models that can assimilate successive measurements and propagate their information through a shared three-dimensional atmospheric state. Lightweight Eulerian models with consistent sensitivities are particularly useful for repeated gradient-based assimilation and emission inversion, supporting the development of end-to-end services that connect satellite observations to air-quality analyses and forecasts.

TARSA's medium-term objective is to support monitoring and forecasting of near-surface $PM_{2.5}$ and PM_{10} for air-quality applications, including support for environmental directives and assessment of industrial emissions. The inert-tracer and Saharan dust experiments presented here establish the transport and observation-comparison foundations for this development.

Atmospheric composition models range from plume and puff systems such as ADMS (Carruthers et al., 1994), AERMOD (Cimorelli et al., 2005), CALPUFF (Scire et al., 2000), and SCIPUFF (Sykes et al., 1998) to comprehensive Eulerian chemistry–aerosol systems such as WRF-Chem (Grell et al., 2005), GEOS-Chem (Bey et al., 2001), and CAMS/IFS (Flemming et al., 2015). The former provide mature tools for local source-impact and regulatory applications; the latter are indispensable for forecasting and full-budget studies. Regional observation-oriented applications can benefit from an intermediate model with a gridded three-dimensional state, separable processes, controllable numerical settings, and a clear mapping from controls to observations. Restricting the physical processes to prescribed-coefficient transport and first-order removal makes such a model easier to configure, inspect, and evaluate repeatedly. Table 1 summarises this positioning.

Lagrangian systems such as FLEXPART (Stohl et al., 2005), HYSPLIT (Stein et al., 2015), and STILT (Lin et al., 2003) offer complementary advantages: particle methods preserve narrow plumes and can efficiently simulate sparse releases over large domains. They also have mature sensitivity machinery, including backward-mode source–receptor relationships (Seibert and Frank, 2004) and gridded STILT footprints used in atmospheric inversions. These calculations are organised per receptor and carry particle-sampling noise. TARSA instead uses deterministic finite-volume transport on a shared grid, supporting repeated satellite observation-operator evaluations and comparison of tangent-linear, discrete-adjoint, and automatic-differentiation sensitivities on the same forward trajectory. This choice addresses full-state regional applications while recognising the strengths of backward-mode Lagrangian models for receptor-oriented problems.

TARSA builds on established methods. Tangent-linear and adjoint sensitivity analysis is well developed in atmospheric transport (Sandu et al., 2005); monotone flux-limited schemes are only piecewise differentiable, and limiter activation can affect discrete-adjoint consistency (Liu and Sandu, 2008; Alexe and Sandu, 2007). This has motivated stabilised adjoint formulations (Botchorishvili et al., 2018) and continuously differentiable or regularised limiter families, including regularised Koren constructions (Woodfield et al., 2024; Huang et al., 2025). Differentiable forward solvers are also increasingly used across the geosciences (e.g. Kochkov et al., 2024).

TARSA's contribution is the integration and verification of these elements in a lightweight regional aerosol-transport framework. Open Eulerian frameworks with validated tangent-linear and adjoint transport already exist, including Polair3D within Polyphemus (Mallet et al., 2007; Quélo et al., 2007), STEM (Sandu et al., 2005), and SILAM (Sofiev et al., 2015), where sensitivity capabilities accompany comprehensive chemistry-transport systems. TARSA combines (i) a conservative finite-



60 volume core with separately inspectable physical processes; (ii) a Koren limiter with cubic-Hermite regularisation within the
total-variation-diminishing (TVD) bounds; (iii) analytic, tangent-linear, discrete-adjoint, and automatic-differentiation sen-
sitivity paths checked for an emission-amplitude control; and (iv) station and satellite observation operators in reproducible
prescribed-meteorology configurations. We document the model and verify its numerics, evaluate inert transport against ETEX-
1 (Nodop et al., 1998) with a supporting CAPTEX-1 comparison (Ferber et al., 1986), and assess dust columns during Fennec
(Ryder et al., 2013; Chaboureaud et al., 2016). A CAMS replay (Inness et al., 2019) supplies a supplementary consistency test.
65 TARSA complements comprehensive CTMs for controlled regional studies; the present evidence concerns the forward model
and low-dimensional sensitivity consistency, with operational inverse-modelling performance left to future work.

Table 1. Positioning of TARSA relative to commonly used atmospheric transport and composition model classes.

Model class	Examples	Primary strengths	Limitations for repeated retrieval/inversion-oriented use	TARSA position
Plume, puff, and regulatory dispersion models	ADMS, AERMOD, CALPUFF, SCIPUFF	Efficient local source-impact and regulatory or emergency-response studies	Not primarily designed around a fixed regional 3-D gridded state and satellite observation operators	Gridded regional Eulerian forward operator
Lagrangian particle and trajectory models	FLEXPART, HYSPLIT, STILT	Strong plume-following capability, efficient sparse-release simulations, and exact per-receptor source-receptor sensitivities via backward runs	Stochastic sampling noise and per-receptor organisation make full-state fixed-grid tangent-linear/adjoint pairs and repeated satellite-operator evaluations less direct	Deterministic finite-volume transport on a fixed grid
Comprehensive chemistry-aerosol models	WRF-Chem, GEOS-Chem, CAMS/IFS, SILAM	Full chemistry, aerosol budgets, forecasting, and reanalysis capability	High process complexity and computational cost can be a disadvantage for controlled process-sensitivity and repeated inversion experiments	Lightweight, process-separable complement, not a replacement
Aerosol reanalyses and replay products	CAMS, MERRA-2	Observation-constrained, physically rich reference fields	Agreement can reflect assimilated constraints or imposed boundary and initial fields rather than independent forward-transport skill	Forcing/reference data and consistency checks while keeping the main forward operator inspectable



Sections 2–4 describe the model, numerics, and sensitivity structure; Sects. 5–8 present the experiments and their interpretation. The appendices provide detailed formulations, supporting results, and a quickstart workflow.




Figure 1. TARSA forward-operator workflow. Prescribed meteorology, emissions, initial and boundary conditions, and tracer-specific properties define the regional transport problem. TARSA advances tracer mass mixing ratios with a conservative finite-volume core including advection, vertical turbulent diffusion, sedimentation, and parameterised dry and wet deposition. The resulting fields can be transformed by observation operators into surface concentrations, vertical columns, station samples, and satellite-collocated quantities for verification, evaluation, and future inversion- or data-assimilation-oriented workflows.

2 Model scope and governing equations

2.1 Scope and intended applications

TARSA is a regional three-dimensional Eulerian transport model for generic tracers expressed as mass mixing ratios. In the present version, meteorology is prescribed externally, and the prognostic core contains only physical processes that are linear in the tracer concentration for fixed meteorological forcing. The implemented process set therefore consists of advection, turbulence-driven vertical diffusion, gravitational settling, and parameterised dry and wet deposition. Optional linear source or sink terms can be used for prescribed emissions or simple first-order decay. Nonlinear aerosol microphysics are not solved prognostically. No interactive chemistry module is included in v1.0.

Figure 1 summarises the mapping from prescribed inputs through tracer transport to gridded fields and observation-space quantities.



2.2 Governing transport equation

80 The transport of aerosols and trace gases in TARSA is described by a conservation equation for a generic tracer C , expressed as mass mixing ratio (kg of species per kg of moist air):

$$\frac{\partial(\rho C)}{\partial t} + \frac{\partial(\rho u_x C)}{\partial x} + \frac{\partial(\rho u_y C)}{\partial y} + \frac{\partial[\rho(u_z - w_{\text{sed}})C]}{\partial z} = \frac{\partial}{\partial z} \left(\rho K_z \frac{\partial C}{\partial z} \right) + \rho E + \rho F_{\text{dry}} + \rho F_{\text{wet}}, \quad (1)$$

where (u_x, u_y, u_z) are the resolved wind components, w_{sed} is the effective gravitational settling velocity, K_z is the vertical turbulent diffusivity, ρ is air density, E is an external source term expressed as a mixing-ratio tendency, and F_{dry} and F_{wet} are

85 linear sink tendencies representing dry and wet deposition.

In this notation C is the prognostic mass mixing ratio. The corresponding species mass in a discrete cell is therefore $\rho_p V_p C_p$, and tendency terms such as E have units of mixing ratio per unit time; multiplication by ρ converts them to species-mass tendencies per unit volume.

With prescribed meteorology and source coefficients, Eq. (1) is linear in the tracer. The discrete first-order upwind update is
90 affine, whereas the regularised Koren option, based on a monotonic upstream-centred scheme for conservation laws (MUSCL), uses solution-dependent limiter weights. Section 4 explains the implications for sensitivities and superposition. Neither configuration clips the evolving tracer state for non-negativity during time stepping.

2.3 Meteorological input and vertical coordinate

Meteorological forcing is taken from ERA5 (Hersbach et al., 2020). TARSA uses the pressure-level product because its online
95 CDS-disk access is faster than the ERA5-complete model-level retrieval from the MARS tape archive (Hersbach et al., 2023a). Native horizontal grid points are reprojected onto TARSA's Cartesian coordinates without changing field values (Sect. 3.1); the refined ETEX configuration additionally interpolates the fields (Sect. 5.6). ERA5 pressure velocity ω (Pa s^{-1}) is converted hydrostatically to physical vertical velocity, $u_z = -\omega/(g\rho)$, with $g = 9.81 \text{ m s}^{-2}$.

Air density is diagnosed from pressure, temperature, and humidity as $\rho = p_l/(R_d T_v)$, where $R_d = 287.05 \text{ J kg}^{-1} \text{ K}^{-1}$ is
100 the gas constant for dry air and T_v is virtual temperature, $T_v = T(1 + 0.61q)$, with T the air temperature and q the specific humidity.

Level heights are read directly from ERA5 pressure-level geopotential (z) as $z = \Phi_{\text{ERA5}}/g$. TARSA performs no separate hydrostatic integration and does not subtract surface elevation, so these heights are relative to mean sea level.

Each of the 25 requested pressure levels (Appendix F) is used directly as one layer in the native configurations, without
105 vertical interpolation of wind, temperature, humidity, or geopotential. The ETEX station evaluation refines these fields and their geometry to 49 layers (Sect. 5.6); auxiliary CAMS composition fields are separately interpolated in pressure coordinates onto the ERA5 levels. Hourly meteorological snapshots are held constant across internal model sub-steps until the next snapshot is reached (zero-order hold). Heights and layer thicknesses Δz_k are recomputed from geopotential at each forcing hour and held fixed within that hour. The tracer mixing ratios are not explicitly remapped when geometry changes. Consequently, the
110 fixed-geometry mass-conservation tests of Sect. 6.1.1 do not test this hour-to-hour update.



TARSA v1.0 does not represent terrain explicitly: pressure levels are isobaric surfaces, cells are neither masked nor blocked according to elevation, and nominally sub-surface levels are not checked. Any extrapolated ERA5 values at those levels are used as ordinary air. Over strong orography, including the Alps in the ETEX domain, near-surface transport and deposition therefore depend on the ERA5 pressure-level extrapolation; TARSA neither blocks nor diverts flow around terrain. Reduced
115 near-surface vertical resolution also limits boundary-layer gradients and shallow plume transport. Model-level forcing and WRF meteorological input are planned extensions (Sect. 7).

2.4 Vertical turbulent diffusion parameterisations

Vertical turbulent mixing is represented in flux form, with K_z diagnosed from stability and wind shear using a simplified ECMWF boundary-layer parameterisation (ECMWF, 2024a). Stable, neutral, and unstable regimes are distinguished as de-
120 tailed in Appendix D. The local unstable closure applies independently of layer position relative to the diagnosed boundary-layer height; TARSA has neither the full eddy-diffusivity mass-flux (EDMF) treatment nor a separate convective mixed-layer scheme. Section 7 considers the implications for deep unstable boundary layers.

2.5 Sedimentation and deposition parameterisations

TARSA v1.0 includes the aerosol processes that can be represented as linear operators acting on the transported tracer. These
125 are gravitational settling, dry deposition, and wet deposition. TARSA also provides an optional diagnostic estimate of hygroscopic growth, which is often required when aerosol optical properties are derived from transported mass. The resulting effective particle properties may be used in the sedimentation and deposition parameterisations, but this diagnostic step does not introduce nonlinear coupling in the state update.

Details of the wet- and dry-deposition parameterisations used in the present implementation are given in Appendix E. In the
130 experiments considered here, sedimentation is used for aerosol tracers. Wet deposition is applied to the aerosol tracers with aerosol-type-dependent in-cloud nucleation efficiencies (Appendix E), so that poorly soluble species such as mineral dust are removed mainly by below-cloud scavenging. In the CAMS replay experiment, dry deposition is switched off because lower-boundary concentrations are prescribed directly from CAMS; adding an explicit near-surface sink would otherwise double count removal already implicit in the reanalysis fields.

3 Numerical implementation

3.1 Grid, indexing, and state vector

Equation (1) is solved on a structured finite-volume grid with $N_x \times N_y \times N_z$ cells. Each cell (i, j, k) has dimensions Δx_i , Δy_j , and Δz_k . Horizontal cell sizes are determined from the chosen projection, while vertical cell thicknesses are derived from the ERA5-based height coordinate.



140 The vertical grid follows Sect. 2.3. Horizontal projection is selected from the domain extent: a Universal Transverse Mercator zone for compact domains, a regional equal-area or conformal projection or Web Mercator fallback for larger domains, and a PROJ-free equirectangular projection as a last resort. The resulting Δx_i and Δy_j vary weakly with position except in the equirectangular case.

For efficient implementation, the three-dimensional indices are mapped to a one-dimensional state vector using column-major storage, so that all vertical levels for a fixed horizontal location are contiguous in memory. The mapping is

$$p = (j - 1)N_x N_z + (i - 1)N_z + k, \quad (2)$$

with $p = 1, \dots, N_x N_y N_z$. The state vector at time level n is therefore

$$\mathbf{X}^n = (C_1^n, C_2^n, \dots, C_N^n)^T. \quad (3)$$

3.2 Finite-volume discretisation

150 The governing equation is integrated over each control volume to obtain a conservative semi-discrete finite-volume formulation. Writing the transport equation in conservative form and integrating over cell $V_{i,j,k}$ gives

$$\frac{d}{dt} \int_{V_{i,j,k}} \rho C dV + \int_{\partial V_{i,j,k}} \rho \mathbf{u}_{\text{eff}} C \cdot \mathbf{n} dS = \int_{\partial V_{i,j,k}} \rho K_z \frac{\partial C}{\partial z} \mathbf{e}_z \cdot \mathbf{n} dS + \int_{V_{i,j,k}} \rho (E + F_{\text{dry}} + F_{\text{wet}}) dV, \quad (4)$$

where $\mathbf{u}_{\text{eff}} = (u_x, u_y, u_z - w_{\text{sed}})^T$ and $\mathbf{n}$ is the outward unit normal on the continuous control-volume boundary. In the discrete system, the surface integrals are replaced by flux sums over the six faces of each cell; below, $\mathbf{n}_f$ denotes the unit normal assigned to a particular oriented face f . This yields a conservative balance between the time tendency, advective and sedimentation face fluxes, vertical diffusive fluxes, and local source or sink terms.

3.3 Air-mass consistency of prescribed winds

Prescribed ERA5 winds need not satisfy the discrete continuity equation $\partial \rho / \partial t + \nabla \cdot (\rho \mathbf{u}) = 0$ on the TARSA grid. Inconsistent tracer and air-mass advection can therefore generate spurious tendencies even for uniform mixing ratios (Jöckel et al., 2001).
160 TARSA uses the hydrostatic vertical-velocity conversion and hourly diagnostic density of Sect. 2.3; it does not re-diagnose u_z from horizontal mass-flux divergence, apply a pressure or mass fixer, or carry density prognostically between forcing hours. An optional two-dimensional utility projects horizontal winds to reduce divergence at each level, but does not enforce the joint three-dimensional air-mass continuity budget.

The default mass-consistent (MC) regularised Koren MUSCL operator advances air density ρ and tracer mass density ρC together with the same flux formulation, then recovers $C^* = (\rho C)^* / \rho^*$ after each sub-step. For initially uniform C , tracer-mass fluxes remain proportional to air-mass fluxes, preserving the mixing ratio at the discrete flux level even when the prescribed fluxes do not satisfy continuity. This property does not prevent spurious redistribution of non-uniform tracers by discretely divergent winds. The explicit main-text configurations use this paired update; the monolithic implicit first-order option does not. Section 6.1.2 tests uniform-field preservation under real forcing separately from global tracer-mass conservation.



170 3.4 Advection options implemented in TARSA

TARSA implements two finite-volume advection options. First-order upwind, denoted $\mathcal{A}_{UP}$, is monotone by construction and is used for the monolithic fully implicit CAMS replay configuration, where robustness and a fixed sparse update are more important than plume sharpness. The regularised Koren MUSCL scheme, denoted $\mathcal{A}_{rKM}$, is the default option for the sharp-plume experiments. It is a MUSCL-type reconstruction in the sense of [van Leer \(1979\)](#) with the Koren limiter introduced by
175 [Koren \(1993\)](#): in smooth plume regions it reduces numerical diffusion relative to first-order upwind, while near steep gradients it limits the reconstructed slope toward the monotone upwind value.

Under the CFL conditions used here, $\mathcal{A}_{rKM}$ suppresses spurious oscillations and helps preserve non-negative concentrations. Cubic-Hermite patches replace the two positive-ratio Koren branch transitions, matching the adjacent values and slopes while leaving the limiter unchanged elsewhere. This local regularisation retains TVD-bounded behaviour and reduces selected
180 derivative discontinuities without the wider coupling of a global spline. Appendix A gives the reconstruction, limiter formula, and diagnostic; Sect. 4 states the limits of the resulting local sensitivities.

3.5 Vertical diffusion discretisation

Vertical diffusive fluxes are discretised with a second-order central difference formulation. The vertical gradients are approximated at cell interfaces, and both K_z and ρ are interpolated to those interfaces. Because the vertical grid is generally non-
185 uniform, TARSA uses distance-weighted linear interpolation at faces, which reduces to the arithmetic mean on uniform spacing. The resulting implicit vertical-diffusion operator couples only neighbouring levels within each column.

3.6 Boundary conditions for the explicit configuration

TARSA supports two boundary-condition types on the lateral and bottom faces of the domain: homogeneous Neumann (zero-gradient) and Dirichlet (prescribed-value). The lateral choice applies symmetrically to all four lateral faces, the bottom face
190 is chosen separately, and the treatment of the model top is not itself selectable: it is fixed by the advection scheme and by the vertical-diffusion solver, as described below (in the simulation builder these correspond to the keywords `lateral_bc` and `bottom_bc` with values `:neumann0` and `:dirichlet`, both defaulting to `:neumann0`; Listing 1 and the model documentation give the interface details). TARSA stores no padded ghost-cell array; each boundary face's exterior state is instead evaluated on the fly inside the flux kernels, substituting either the nearest interior cell value (zero-gradient) or a value
195 drawn from the prescribed boundary field (Dirichlet). The regularised Koren MUSCL reconstruction additionally sets the limited slope to zero at the two end cells of every stencil line (Appendix A2), so the reconstruction degrades to first order at the domain edge rather than extrapolating from a padded halo.

At a Dirichlet lateral face, the prescribed value is drawn from the same four-dimensional boundary-condition array that supplies the simulation's initial condition, evaluated at the current forcing time; this single array therefore serves both as the
200 $t = 0$ initial condition and, wherever a Dirichlet face is active, as the time-varying inflow concentration. Air density at lateral faces always uses a zero-gradient wall regardless of the tracer boundary choice, so only the tracer field itself can be pinned to



a prescribed inflow value. No special outflow treatment is needed: the smooth upwind flux split (Appendix A1) already selects the interior state at an outflow face, so the zero-gradient condition covers both inflow-free and outflow lateral boundaries.

At the bottom boundary, a Dirichlet condition applies a strong Dirichlet row that pins the lowest layer to the prescribed value in the vertical-diffusion step and substitutes that same value for the lowest-layer state wherever the vertical-advection flux stencil reads it. A zero-gradient bottom instead imposes zero diffusive flux at the lowest interface and copies the nearest interior tracer state for the advective exterior state; this state condition does not itself set the face velocity. The high-level simulation builder additionally zeroes the bottom-face vertical velocity by default for this case (unless its explicit opt-out is selected), giving an impermeable advective wall. The direct ETEX-1 validation drivers, however, supply the face velocities directly and retain the ERA5 vertical velocity at the lowest face. Their reported lower boundary is therefore zero-gradient for tracer reconstruction and diffusion but open to prescribed vertical advection at the nominal 1000 hPa truncation surface; it should not be interpreted as a terrain-following impermeable ground wall. Dry deposition and sedimentation are implemented independently of the bottom condition, as a first-order volumetric loss confined to the lowest layer and as a reduction of the vertical wind by the settling velocity, respectively (Appendix E); combining dry deposition with a Dirichlet bottom can double-count near-surface removal. At the model top, the implicit vertical-diffusion solve always applies a strong Dirichlet condition using the prescribed top-layer value (zero by default), independent of the bottom choice, whereas the regularised Koren MUSCL advection operator applies a zero-gradient condition at the top face; with the zero-valued top boundary array used throughout this paper, both conventions let a vanishing tracer field pass through the top essentially unchanged.

Table 2 lists the experimental boundary conditions. In the Gaussian verification, zero-valued lateral Dirichlet boundaries act as an absorbing far field for the compact interior source.

3.7 Time integration and solver configurations

TARSA provides multiple solver configurations, and the experiments in this paper use different ones depending on their numerical purpose.

For the Gaussian verification, ETEX-1 evaluation, minimal gradient check, and Fennec dust-column evaluation, we use the explicit regularised Koren MUSCL advection operator advanced with a third-order strong-stability-preserving Runge–Kutta method (SSP-RK3) and combine it with implicit vertical diffusion using second-order Strang splitting. This is the default sharp-plume configuration of the present paper.

For the CAMS replay experiment, we use a monolithic fully implicit backward-Euler configuration with first-order upwind advection and implicit vertical diffusion assembled into a global sparse linear system per time step. This configuration is more diffusive but robust under strongly constrained, time-varying boundary conditions.



Let $L(\mathbf{X})$ denote the spatial advection operator. The third-order strong-stability-preserving Runge–Kutta (SSP-RK3) scheme advances the solution through three stages:

$$\mathbf{X}^{(1)} = \mathbf{X}^n + \Delta t L(\mathbf{X}^n), \quad (5)$$

$$\mathbf{X}^{(2)} = \frac{3}{4} \mathbf{X}^n + \frac{1}{4} (\mathbf{X}^{(1)} + \Delta t L(\mathbf{X}^{(1)})), \quad (6)$$

$$235 \quad \mathbf{X}^{n+1} = \frac{1}{3} \mathbf{X}^n + \frac{2}{3} (\mathbf{X}^{(2)} + \Delta t L(\mathbf{X}^{(2)})). \quad (7)$$

The vertical diffusion step leads to an independent tridiagonal solve in each (i, j) column,

$$\mathbf{T}_{i,j} \mathbf{X}_{i,j}^{n+1} = \mathbf{r}_{i,j}, \quad (8)$$

which is solved with the Thomas algorithm (Thomas, 1949) in $\mathcal{O}(N_z)$ operations per column.

In the sharp-plume configuration, advection and vertical diffusion are combined using second-order Strang splitting,

$$240 \quad \mathbf{X}^{n+1} = \Phi_{\text{diff}}(\Delta t/2) \Phi_{\text{adv}}(\Delta t) \Phi_{\text{diff}}(\Delta t/2) \mathbf{X}^n + \mathbf{q}, \quad (9)$$

where $\mathbf{q}$ collects source, sink, and boundary contributions over the split step.

3.8 Experiment-specific numerical configurations

Table 2 summarises the experimental configurations. The real-case evaluations use the explicit sharp-plume configuration, except for the supplementary CAMS replay, which uses monolithic implicit first-order upwind. The MMS diagnostic compares
245 both options.

3.9 Software implementation

TARSA is implemented in Julia (version 1.12 or later). The complete dependency environment, including the exact versions of every package used for this paper, is pinned by the `Project.toml` and `Manifest.toml` files in the repository root and in the archived release. Model state, meteorology, emission, and boundary-condition arrays are stored in single precision by
250 default; the closed-domain mass-conservation checks of Sect. 6.1.1 are the only experiments run in double precision, consistent with the machine-precision drifts reported there. NetCDF input and output, ERA5 retrieval from the Copernicus Climate Data Store, and optional regridding of auxiliary composition fields are handled by native Julia libraries, so no external language runtime is required.

Sensitivities use two automatic-differentiation routes. Enzyme.jl (v0.13 series, as pinned in the archived manifest) differenti-
255 ates the directional-split transport operator in forward mode for the inversion-oriented gradient interface (Sect. 4), whereas the autodiff column of the structural gradient check (Sect. 6.1.5) uses ForwardDiff.jl dual numbers to differentiate the SSP-RK3 stepper, which the pinned Enzyme version could not differentiate end to end at the time of writing; an Enzyme-based RK3 gradient-check script is nevertheless distributed with the code so this path can be reassessed as Enzyme evolves. The regularised Koren advection kernels and their tangent-linear and adjoint counterparts are parallelised with shared-memory threads,



Table 2. Numerical configuration used in each verification or evaluation experiment. “Regularised Koren MUSCL” denotes $\mathcal{A}_{\text{TKM}}$, the total-variation-diminishing (TVD)-limited MUSCL reconstruction with the cubic-Hermite-regularised Koren limiter. “First-order upwind” denotes $\mathcal{A}_{\text{UP}}$. “Implicit BE” denotes the monolithic fully implicit backward-Euler configuration. Boundary conditions are listed as lateral / bottom; the top-boundary treatment is fixed by the solvers (Sect. 3.6).

Experiment	Advection	Time integration	Vertical diffusion	Splitting	Δt / CFL	Lateral / bottom BC
Gaussian plume (verification)	Regularised Koren MUSCL	SSP-RK3	Implicit (Thomas)	Strang	CFL $\approx$ 0.5	Dirichlet (zero) / zero-gradient
Minimal gradient check	Regularised Koren MUSCL	SSP-RK3	Implicit (Thomas)	Strang	CFL $\approx$ 0.5	Dirichlet (zero) / zero-gradient
ETEX-1 (independent evaluation) ^a	Regularised Koren MUSCL	SSP-RK3	Implicit (Thomas)	Strang	1200 s, CFL $\approx$ 0.9	zero-gradient / zero-gradient
Fennec dust-column evaluation	Regularised Koren MUSCL	SSP-RK3	Implicit (Thomas)	Strang	3600 s, CFL $\approx$ 0.7	zero-gradient / zero-gradient
CAMS replay (supplementary appendix)	First-order upwind	Implicit BE	Implicit (monolithic)	none	3600 s	Dirichlet (CAMS) ^b / Dirichlet (CAMS)

^a Configuration used for the reported station comparisons (Sect. 5.6), on ERA5 meteorology bilinearly/linearly refined by a factor of two from the native grid. A corresponding native-resolution forward configuration, with a 3600 s forcing step and CFL $\approx$ 0.5, is used only for the configuration comparison; the real-forcing uniform-tracer diagnostic of Sect. 6.1.2 reuses its native grid, forcing, and advection kernel but has its own diagnostic stepping. For both direct ETEx forward drivers, “zero-gradient bottom” denotes the tracer-reconstruction and diffusive conditions; ERA5 vertical velocity is retained at the lowest advective face.

^b Imposed at inflow faces only; zero-gradient at outflow (Appendix G).

260 one thread per line of cells along the sweep direction; the remaining kernels execute serially, and no distributed-memory or GPU code path is implemented. The fully implicit configuration solves its global sparse linear system with a restarted, ILU-preconditioned GMRES iteration (Appendix G); the corresponding adjoint systems are solved directly.

3.10 Computational cost

265 Table 3 reports ETEx-1 wall-clock time and peak memory on an AMD Ryzen 9 5950X workstation (16 cores, 32 GB RAM) with eight threads. The native 120-hour run completes in 66 s, including about 42 s of time stepping; the refined configuration used for station statistics completes in 6.5 min. The output-free advection–diffusion diagnostic on the native grid takes 39 s (Sect. 6.1.2), separating kernel cost from loading, output, and driver overhead.

These runtimes permit repeated regional simulations on a workstation. Gradient-cost scaling is reported separately in Sect. 6.1.5.



Table 3. Measured wall-clock cost on a single AMD Ryzen 9 5950X workstation (16 physical cores, 8 threads). All rows simulate the same 120-hour forcing window; cells×hours/s normalises by simulated forcing time rather than by the number of stored output frames, which differs between the native and refined ETEX-1 configurations.

Configuration	Grid ($N_x \times N_y \times N_z$)	Forcing h	Threads	Wall time	Peak RSS	Cells×h/s
ETEX-1, native grid	$101 \times 81 \times 25$	120	8	65.8 s	2.94 GB	3.7×10^5
ETEX-1, refined ×2 (headline)	$201 \times 161 \times 49$	120	8	398.5 s	12.84 GB	4.8×10^5
<i>Supplementary: advection+diffusion kernels only, no emissions/deposition/output (Sect. 6.1.2)</i>						
Advection+diffusion kernel, native grid	$101 \times 81 \times 25$	120	8	39.2 s	2.36 GB ^a	6.3×10^5

^a Peak RSS for the whole diagnostic run, not isolated to the advection+diffusion case alone.

270 4 Differentiability and forward-operator structure

Retrieval and inversion applications require sensitivities of the discrete map from controls to observations. TARSA combines prescribed-coefficient linear physics with the regularised advection scheme of Sect. 3.4. Appendix B gives the full tangent-linear and discrete-adjoint formulations and their relation to automatic differentiation.

For a fixed grid, time step, meteorological forcing, boundary-condition treatment, and solver configuration, TARSA advances
275 the discrete tracer state through a one-step map

$$\mathbf{X}^{n+1} = \mathbf{F}_n(\mathbf{X}^n, \theta), \quad (10)$$

where $\mathbf{X}^n$ collects the transported tracer mass-mixing-ratio values at time step n , and θ may contain emissions, initial conditions, boundary values, or selected physical parameters.

The default regularised Koren MUSCL update has solution-dependent limiter weights and a trajectory-dependent Jacobian.
280 Its cubic-Hermite patches smooth selected positive-ratio transitions on nondegenerate stencils; they do not establish global differentiability of the full model. The first-order upwind configuration instead gives an affine sparse update for fixed meteorology, coefficients, and boundary treatment, whose state Jacobian is independent of the tracer trajectory.

When local derivatives are well defined along a reference solution, the corresponding tangent sensitivity has the compact form

$$285 \quad \mathbf{S}^{n+1} = \mathbf{J}_n \mathbf{S}^n + \partial_\theta \mathbf{F}_n, \quad \mathbf{J}_n = \frac{\partial \mathbf{F}_n}{\partial \mathbf{X}^n}, \quad (11)$$

where $\mathbf{S}^n = \partial \mathbf{X}^n / \partial \theta$ and $\mathbf{J}_n$ is evaluated along that trajectory. The scalar experiment in Sect. 6.1.5 verifies that these gradient paths agree to single precision for a deliberately simple emission-amplitude control, within the bounded scope stated in the introduction.

5 Verification and evaluation setup

5.1 Hierarchy of evidence

The experiments reported here serve distinct purposes. The manufactured Gaussian plume test provides *verification* of the core advection–diffusion numerics under controlled conditions. The ETEX-1 tracer release provides the main *independent evaluation* of long-range inert transport against observations, supported by a shorter CAPTEX-1 field comparison (Ferber et al., 1986) outside the European ETEX setting. The Fennec dust-column comparison (Ryder et al., 2013; Chaboureaud et al., 2016) extends the evaluation to aerosol applications, assessing simulated dust columns and the workflow for comparing model output with remote-sensing observations. It uses realistic meteorology, prescribed MERRA-2 dust emissions, a four-day spin-up, and zero-gradient (homogeneous Neumann) lateral boundary conditions for dust, without importing MERRA-2 dust concentrations through the initial state or lateral boundaries. These results do not independently validate the prescribed dust sources. A separate reference configuration, MERRA-2-BC/IC, prescribes MERRA-2 dust concentrations in the initial and lateral boundary conditions; it is archived to allow the influence of these external dust fields to be quantified, but its results are not analysed here. Finally, the replay in Appendix C uses CAMS initial and lateral boundary conditions (Inness et al., 2019) to test consistency under these constraints and demonstrate the replay workflow. It therefore carries less weight as independent evidence of model skill.

5.2 Common forcing and domain settings

All real-case experiments use the ERA5 forcing conventions of Sect. 2.3, with the case-specific grids described below. Synthetic verification uses prescribed fields. Numerical and boundary settings are summarised in Table 2; the model-top treatments follow Sect. 3.6.

5.3 Tracer configurations

Table 4 lists the main-text tracers and active processes: inert tracers for Gaussian and ETEX verification/evaluation, and five mineral-dust bins for Fennec. The supplementary CAMS replay uses CO, dust bins 1–2, OM, BC, and sea-salt bin 1 (Appendix C).

Table 4. Tracer and process configuration used in the main-text experiments discussed in this paper. The Gaussian grid is the reference resolution of the grid-sensitivity study in Sect. 6.1.3; the ETEX-1 grid is the refined station-evaluation grid (Sect. 5.6).

Species / tracer	Type	Adv.	Turb. diff.	Sedim.	Wet dep.	Dry dep.	$N_x \times N_y \times N_z$
Gaussian test tracer	passive	+	+	no	no	no	$121 \times 11 \times 61$
ETEX-1 PMCH	inert gas	+	+	no	no	no	$201 \times 161 \times 49$
Fennec dust bins 1–5	aerosol	+	+	+	+	+	$113 \times 73 \times 25$



5.4 Manufactured Gaussian plume verification

The first verification case compares TARSA to a steady analytical plume solution in a uniform flow. The aim is to verify the combined action of advection and vertical turbulent diffusion in a controlled setting. We use a Cartesian domain

$$[-1500, 1500] \times [-400, 400] \times [0, 100] \text{ m},$$

discretised on a uniform grid. The reference configuration uses constant wind $u_x = 5 \text{ m s}^{-1}$, $u_y = u_z = 0$, constant vertical eddy diffusivity $K_z = 0.05 \text{ m}^2 \text{ s}^{-1}$, and no horizontal diffusion, matching TARSA's vertical-diffusion-only turbulence closure; the lateral boundaries are open (zero-valued Dirichlet, acting as an absorbing far field) and the lower boundary is reflecting. A steady source of strength $Q = 100$ (mixing-ratio units $\times \text{m}^3 \text{ s}^{-1}$) is injected into the single cell at the horizontal quarter-point of the domain ($x_0 \approx -750 \text{ m}$), centred in y , at mid-height $z_0 = 50 \text{ m}$. TARSA is integrated from a zero initial condition for 900 s (roughly two domain-transit times, sufficient to reach a pseudo-steady state downwind), and the result is compared against the analytical reference using scatter plots and scalar error metrics.

Because the configuration has neither lateral wind nor lateral diffusion, the appropriate analytical reference is the steady two-dimensional (x - z) plume of an infinite crosswind line source with ground reflection,

$$C(x, z) = \frac{Q_l}{\sqrt{2\pi} u_x \sigma(x)} \left\{ \exp \left[-\frac{(z - z_0)^2}{2\sigma^2(x)} \right] + \exp \left[-\frac{(z + z_0)^2}{2\sigma^2(x)} \right] \right\}, \quad \sigma^2(x) = \frac{2K_z (x - x_0)}{u_x} \quad (x > x_0), \quad (12)$$

where $Q_l = Q/\Delta y$ is the line-source strength corresponding to the single-cell volumetric source of lateral cell width Δy . The analytical field is cell-averaged over each (x, z) cell, and the comparison uses all cells downwind of the source on the source-aligned y slice.

As a separate time-dependent manufactured-solution diagnostic, we also use the method of manufactured solutions (MMS): a periodic image-sum Gaussian translated by the mean flow in a prescribed three-dimensional velocity field with a smooth horizontal vortex component and constant vertical velocity. The corresponding analytical source term is evaluated at cell centres and injected through the same volumetric source interface used by TARSA. This diagnostic is run for both a sharp and a smooth Gaussian and for two solver configurations: the default sharp-plume regularised Koren MUSCL, SSP-RK3, Strang-split configuration and the monolithic fully implicit first-order upwind/backward-Euler replay configuration. The MMS driver applies periodic transport boundaries in all three directions and compares the cell-centre numerical field with the periodic exact solution after translating the Gaussian centre modulo the domain length.

5.5 Scalar structural gradient check

We compare analytic, finite-difference, tangent-linear, adjoint, and automatic-differentiation gradients on a $41 \times 21 \times 7$ grid for $N_t = 80$ steps with $\Delta t = 3600 \text{ s}$ in single precision. Pure amplitude scaling of one fixed source pattern preserves trajectory proportionality in this zero-background configuration, although the solver is nonlinear for general perturbations.

A single scalar parameter α scales the volumetric emission source, $E(\alpha) = \alpha E_0$, while meteorology and all transport coefficients remain fixed. In the manufactured zero-background setup, the trajectory is positively homogeneous in the source



amplitude,

$$\mathbf{X}^n(\alpha) = \alpha \mathbf{X}^n(1) \quad (13)$$

345 for $\alpha \geq 0$.

For sampled concentrations, we use the scalar least-squares objective

$$J(\alpha) = \frac{1}{2} (\mathbf{H}\mathbf{X}(\alpha) - \mathbf{y})^T \mathbf{R}^{-1} (\mathbf{H}\mathbf{X}(\alpha) - \mathbf{y}), \quad (14)$$

with $\mathbf{R} = \mathbf{I}$ in this minimal experiment. The positive-homogeneity property gives the analytic scalar gradient

$$\frac{dJ}{d\alpha} = g_{\text{an}} = (\mathbf{H}\mathbf{X}(\alpha_0) - \mathbf{y})^T \mathbf{R}^{-1} (\mathbf{H}\mathbf{X}(1) - \mathbf{H}\mathbf{X}(0)). \quad (15)$$

350 The tangent-linear, discrete-adjoint, and automatic-differentiation formulations are those given in Appendix B. All gradient evaluations are made at $\alpha_0 = 0.9$. As an independent perturbation-based reference, we compute the centred finite-difference approximation

$$g_{\text{FD}}(\varepsilon) = \frac{J(\alpha_0 + \varepsilon) - J(\alpha_0 - \varepsilon)}{2\varepsilon}. \quad (16)$$

5.6 ETEX-1 independent evaluation

355 In the first European Tracer Experiment (ETEX-1), 340 kg of perfluoromethylcyclohexane (PMCH) was released from Montfermeil, France (48.0583° N, 2.0083° W), between 23 October 1994 16:00 UTC and 24 October 1994 03:50 UTC. TARSA injects the release into the lowest model layer over the half-open 42,600 s interval, distributing it horizontally with mass-conserving bilinear weights. Partial substeps at the release endpoint are overlap-weighted. The rate is $340/42,600 \approx 7.9812 \times 10^{-3} \text{ kg s}^{-1}$; the archived driver records the exact interpolation weights. No tuning or empirical mass rescaling is applied.

360 The ERA5 forcing domain is 40–60° N, 5° W–20° E, at 0.25° resolution (101×81 horizontal points; approximately 28 km meridionally and 14–21 km zonally). Its 25 pressure levels span 1000 hPa (approximately –0.2 to +0.1 km over this forcing window, including extrapolated sub-surface values) to 50 hPa (20.2–20.6 km). The native configuration uses a 3600 s forcing step and CFL ≈ 0.5 and is retained for a supplementary configuration comparison. The reported station evaluation bilinearly refines the hourly meteorology to 201×161 horizontal points (0.125°; approximately 14 km meridionally and 7–11 km zonally), linearly refines the vertical fields and geometry to 49 levels, and uses a 1200 s model/output step with CFL ≈ 0.9 . Both configurations use mass-consistent SSP-RK3 Koren MUSCL advection and Strang-split implicit diffusion, differing only in refinement, time step, and CFL target. Their zero-gradient lower tracer boundary retains ERA5 vertical velocity at the nominal 1000 hPa face and is therefore advectively open (Sect. 3.6). The uniform-tracer diagnostic using the native grid is described separately in Sect. 6.1.2.

370 For the comparison with observations, modelled PMCH mixing ratios in the lowest model layer (nominally 1000 hPa) are converted to concentrations and bilinearly sampled in the horizontal at receptor locations. The instantaneous 1200 s model outputs are integrated by the trapezoidal rule over the official 30 three-hour windows. The bundled ETEX v1.1 metadata defines



interval 1 as 23 October 15:00–18:00 UTC, one hour before the release begins; arrival and peak times use the representative midpoint of each interval and are referenced to the 16:00 UTC release time. We retain only stations with at least 11 quantified interval means carrying official flags 01, 11, 21, or 31; exclude F21, D10, and UK03; remove missing or unquantifiable values (−0.99 and −0.88); and exclude receptors outside the 40–60° N, 5° W–20° E simulation domain instead of clamping them to boundary cells. The last geographic criterion removes 12 otherwise eligible stations. We evaluate station-based diagnostics including mean concentration, maximum concentration, peak time, and three threshold-based metrics defined with a common detection threshold of 0.05 ng m^{-3} : dosage (time-integrated concentration at or above threshold), arrival time (midpoint of the first three-hour interval at or above threshold), and duration (total time at or above threshold). Peak time is likewise defined only for threshold-qualified plume hits. Exact definitions of these diagnostics are given in Appendix H. These diagnostics form the main basis for the empirical evaluation of TARSA in this paper.

5.7 Fennec dust-column evaluation

The Fennec case considers Saharan dust during June 2011 (Ryder et al., 2013; Chaboureaud et al., 2016). ERA5 forcing is requested over 20° W–12° E, 12°–34° N; the interior transport grid spans 18° W–10° E, 14°–32° N, with 113×73 points at 0.25° and 25 pressure levels. The simulation runs during 12–25 June, with spin-up on 12–15 June and evaluation on 16–24 June. GRASP/POLDER pixels are collocated on the overlap of their product domain (18° W–14° E, 10°–32° N) with the transport grid. The solver uses the default explicit configuration, a 3600 s step, and $\text{CFL} \approx 0.7$ (Table 2).

The primary run starts from zero dust concentrations and uses homogeneous Neumann lateral and lower tracer boundaries, with sedimentation and dry and wet deposition enabled. MERRA-2 supplies prescribed dust emissions (Gelaro et al., 2017; Randles et al., 2017), but no three-dimensional dust concentrations are imposed at the initial time or lateral boundaries. The post-spin-up comparison therefore evaluates aerosol-column transport under prescribed emissions, without independently validating the dust sources.

An additional MERRA-2-BC/IC configuration with MERRA-2 initial and time-varying lateral dust fields is archived for quantifying the influence of external background and inflow; its results are not presented here.

Dust is represented by five size bins defined by dry particle radius ranges (in μm) and material densities (in kg m^{-3}), as summarised in Table 5. For the evaluation we compare modelled column-integrated dust volume against the GRASP/POLDER Level-3 0.1° column volume concentration of insoluble coarse particles, mapped to the ERA5/TARSA grid and nearest model output hour. The TARSA dust-column quantity used for the Fennec scatter diagnostic, collocated-statistics and process-sensitivity diagnostic table, and daily map comparisons is the summed model bins DU001–DU004. This bin selection follows the size-range overlap with the retrieved coarse mode (Fig. 2): bins DU002–DU004 (1–6 μm radius) lie squarely within the coarse-mode radius range to which the POLDER retrieval is sensitive; DU001 (0.1–1 μm) is included because its upper part overlaps the lower coarse-mode range, although its submicron portion falls partly outside the retrieved coarse mode, making the proxy imperfect at the fine end; and DU005 (6–10 μm) is excluded because such super-coarse particles contribute little to the retrieved coarse-mode column and are strongly depleted by sedimentation away from sources. We refer to this comparison quantity as a coarse insoluble aerosol column volume proxy rather than as a direct dust truth, because the GRASP/POLDER



product is itself a retrieved aerosol-component field and the TARSA source term is prescribed from MERRA-2. Daily map comparisons are formed as POLDER-observation-weighted means so that TARSA and the MERRA-2 reference are sampled consistently with the satellite coverage.

410 The observational product uses GRASP/Component (Dubovik et al., 2011, 2014; Li et al., 2019), which retrieves aerosol-component volume fractions together with size and shape parameters (Dubovik et al., 2021); its advanced optical properties have been evaluated by Zhang et al. (2021). Coarse-mode insoluble dust comprises non-absorbing silicate minerals (CNAI) and absorbing iron oxides (CAI) (Ou et al., 2021; Li et al., 2022). CAI is capped at 3% of coarse-mode volume in the retrieval (Li et al., 2019). The comparison focuses on the dominant CNAI component as a retrieved dust-loading proxy, consistent with
415 previous GRASP/POLDER studies (Ou et al., 2021; Li et al., 2022; Zhang et al., 2021).

Table 5. Definition of the five Fennec dust aerosol bins. Radii are given in μm and densities in kg m^{-3} .

Bin	1	2	3	4	5
Representative radius	0.73	1.4	2.4	4.5	8.0
Lower radius bound	0.1	1.0	1.8	3.0	6.0
Upper radius bound	1.0	1.8	3.0	6.0	10.0
Density	2500	2650	2650	2650	2650




Figure 2. Comparison of the GRASP/POLDER coarse-mode size distribution with the MERRA-2 dust-bin size ranges used for the prescribed emissions.



6 Results

For any model–reference pair $\{M_i, R_i\}_{i=1}^N$, we use standard diagnostics including the root-mean-square error (RMSE), Pearson correlation coefficient r , and coefficient of determination $R^2 = r^2$. Appendix H gives the exact definitions of these and the other summary statistics (FB, NMSE, FM, FMT, dosage, arrival time, duration, peak time, and the CAPTEX-1 composite rank) used throughout this section.

6.1 Numerical verification

6.1.1 Mass conservation

In a closed domain with $E = F_{\text{dry}} = F_{\text{wet}} = 0$ and purely advective–diffusive transport, the discrete total mass

$$M^n = \sum_p \rho_p V_p C_p^n \quad (17)$$

should be preserved up to solver tolerance. We monitor the relative mass error

$$\varepsilon_{\text{mass}}^n = \frac{|M^n - M^0|}{M^0}, \quad (18)$$

using a closed-box configuration with impermeable boundaries and a divergence-free recirculating flow. The density-weighted definition is retained so that the diagnostic matches the species-mass budget used in the general formulation.

We performed two dedicated checks on a $96 \times 96 \times 16$ grid over a $4\text{ km} \times 4\text{ km} \times 0.8\text{ km}$ domain with circulation scale $U_0 = 6\text{ m s}^{-1}$ (300 time steps, CFL=0.45). In the advection-only case, TARSA conserves mass to machine precision: the maximum and end-of-run drift are both $\max |\Delta M / M_0| = 1.99 \times 10^{-16}$. The imposed face-velocity field is divergence-free to numerical round-off ($\|\nabla \cdot \mathbf{u}\|_\infty = 2.52 \times 10^{-17}$), confirming that the observed conservation is consistent with the discrete flux form.

In a second budget-closure test, we add a constant volumetric source in one cell and a uniform first-order decay rate λ , for which the discrete tendency should satisfy

$$M^{n+1} - M^n = \Delta t \left(\sum_p E_p \rho_p V_p - \sum_p \lambda C_p^n \rho_p V_p \right) + \mathcal{O}(\text{solver/roundoff}). \quad (19)$$

The maximum absolute budget residual is 7.33×10^{-8} in mass units, while the maximum relative residual is 3.81×10^{-16} . These checks confirm that the explicit advection step, together with sources and first-order sinks, preserves the global mass budget to numerical precision. Two scope limitations should be noted. First, neither check exercises the implicit vertical-diffusion step non-trivially: the closed-box cases run without diffusion, and in the uniform-tracer real-forcing drift diagnostic of Sect. 6.1.2 diffusion of a uniform field is a no-op. Interior mass conservation of the diffusion solve follows from its flux form, but its strong Dirichlet top row (Sect. 3.6) permits boundary exchange by construction. Second, the closed-box cases use a fixed synthetic geometry and therefore do not test the hour-to-hour vertical-geometry update of the real-forcing configurations (Sect. 2.3).



6.1.2 Uniform-tracer preservation under ERA5 forcing

445 The uniform-field property of Sect. 3.3 is tested over the cached ETEX-1 ERA5 window (23–27 October 1994), with $C \equiv 1$, no emissions or decay, the native $101 \times 81 \times 25$ grid, the mass-consistent SSP-RK3 kernel, $\Delta t = 3600$ s, and CFL = 0.5. The diagnostic uses zero-gradient lateral and bottom tracer conditions. Its boundary array is set to 1, including the strong Dirichlet top row of the diffusion solve, so all boundaries are consistent with the uniform field. It shares the native ETEX grid, forcing, and advection kernel but tests advection alone and advection with full-step diffusion, separately from the forward driver's

450 Strang composition. The diagnostic script is distributed with the model source.

At the first forcing hour, $|\nabla \cdot \mathbf{u}|_\infty = 2.464 \times 10^{-4} \text{ s}^{-1}$. Despite this divergence, advection preserves the uniform field exactly in single precision over 2886 sub-steps: $\max|C - 1| = \text{mean}|C - 1| = 0$. Including implicit vertical diffusion introduces small floating-point drift from the tridiagonal solve: the maximum deviation is 4.86×10^{-5} at hour 50 and the run-averaged mean absolute deviation is 1.57×10^{-6} . This tests mixing-ratio preservation, independently of the fixed-geometry global mass-budget

455 checks above.

6.1.3 Gaussian plume verification

For the Gaussian plume test described in Sect. 5.4, Fig. 3 shows, for the reference grid, sections in the source-aligned x - z plane of the analytical plume of Eq. (12), the TARSA pseudo-steady finite-volume solution, and the normalised (signed) error $(C_{\text{plume}} - C_{\text{FVM}}) / \max(C_{\text{FVM}})$. The two fields agree closely throughout the plume core, with the largest deviations confined to

460 the immediate source cell and to regions where concentrations approach numerical round-off.

Figure 4 presents a cell-wise comparison between the analytical and numerical solutions using $N = 5551$ evaluated grid points. The points remain close to the 1:1 line over a wide dynamic range. For the reference run, the Pearson correlation is $r = 0.982$ ($R^2 = 0.964$), with small fractional bias (FB=0.004), normalised mean square error NMSE=0.285, and figure of merit (FM) of 0.973 (definitions in Appendix H). The remaining spread at the highest concentrations is primarily due to

465 discretisation near the emission cell and to the TVD limiter, which introduces a small amount of diffusion near steep gradients in order to maintain boundedness.

Table 6. Grid sensitivity of the Gaussian plume verification. All runs use the same physical domain and flow configuration, varying only the discrete resolution.

Grid	$N_x \times N_y \times N_z$	N	r	R^2	FB	NMSE	REL _{L2}
coarse	$81 \times 9 \times 41$	2501	0.9766	0.9538	0.0072	0.3845	0.2325
reference	$121 \times 11 \times 61$	5551	0.9820	0.9643	0.0041	0.2845	0.1952
fine	$161 \times 13 \times 81$	9801	0.9850	0.9702	0.0028	0.2317	0.1740

Agreement improves with grid refinement (Table 6): relative L_2 error decreases from 0.233 to 0.195 and 0.174, while fractional bias remains below 0.01. The low observed order (approximately 0.4) reflects the singular point-source cell's con-




Figure 3. Gaussian plume verification test, shown in the source-aligned x - z plane. Left: analytical reference solution of Eq. (12). Middle: TARSA pseudo-steady finite-volume solution on the reference grid. Right: normalised signed error $(C_{\text{plume}} - C_{\text{FVM}})/\max(C_{\text{FVM}})$.




Figure 4. Cell-wise comparison between the analytical Gaussian plume and the TARSA pseudo-steady solution on the reference grid. Both axes are logarithmic. The dashed line indicates the 1:1 relation.



Table 7. Scalar structural gradient-check experiment: single-run method comparison for the objective, gradient, and runtime (single precision). Each time entry is the wall time of one complete end-to-end evaluation of J and its gradient by the given method, including all forward integrations that method requires, with no cross-method caching; timings are indicative single-run measurements, not statistical benchmarks.

method	loss		g	time (s)
analytic	3.850505×10^{-5}	-7.493795×10^{-4}		48.6707
tangent	3.850516×10^{-5}	-7.493810×10^{-4}		49.1917
adjoint	3.850516×10^{-5}	-7.493795×10^{-4}		95.3088
autodiff	3.850516×10^{-5}	-7.493803×10^{-4}		66.1279

tribution to the error norm. This is therefore a physical-consistency check; formal convergence under smooth conditions is
 470 assessed separately by MMS in Sect. 6.1.4.

6.1.4 Method of manufactured solutions (MMS) convergence

The MMS spatial-refinement test (Fig. 5) shows decreasing errors for both solver configurations from $N = 32$ to $N = 128$, with smaller errors for regularised Koren MUSCL. For its sharp Gaussian, L_1 error falls from 1.02×10^{-3} to 1.18×10^{-4} and L_2 error from 5.86×10^{-3} to 5.13×10^{-4} . The finest-grid observed orders are 1.71, 1.86, and 1.43 in L_1 , L_2 , and L_∞ . The
 475 smooth Gaussian has smaller errors but lower observed orders (1.33, 1.26, and 1.38), a difference not fully diagnosed here. Both sequences fall below the unlimited scheme’s formal second order and combine limiter activation, Strang splitting, and discrete source injection. At $N = 128$, the implicit option’s sharp-Gaussian L_2 error is 5.04×10^{-3} , compared with 5.13×10^{-4} for Koren.

Temporal refinement produces much weaker improvement. Reducing Δt from 9.58 to 2.40 s in the sharp Koren case lowers
 480 L_2 error from 6.66×10^{-4} to 5.13×10^{-4} . For the smooth case, reducing Δt from 24.88 to 6.22 s changes L_2 error from 2.96×10^{-4} to 3.07×10^{-4} , so that sequence is not monotonic. Implicit runs likewise show near-zero temporal orders. These are full-operator results dominated by spatial truncation, limiter activation, and source-time coupling, and do not measure the isolated temporal order of SSP-RK3 or Strang splitting. The test establishes error reduction under mesh refinement and lower numerical diffusion for regularised Koren MUSCL within the accuracy limits of this manufactured-source setup.

485 6.1.5 Scalar structural gradient check

Table 7 shows agreement among the four direct methods for the objective and gradient, with a relative gradient spread of about 2×10^{-6} . This quantifies the phrase “agree to single precision” used here. Figure 6 (left) plots each method’s relative difference from finite differences, $|g_{\text{FD}}(\varepsilon) - g| / \max(|g|, \varepsilon_{\text{mach}})$. Because $J(\alpha)$ is quadratic in this amplitude-scaling setup, centred differences have no leading-order truncation error. The curves mainly reflect single-precision cancellation, decreasing
 490 roughly as $1/\varepsilon$ towards 10^{-6} – 10^{-7} at the largest perturbations, rather than forming the two-branch U-curve of general nonlinear objectives.




Figure 5. Time-dependent periodic Gaussian method-of-manufactured-solutions (MMS) convergence diagnostic. Left: spatial refinement of the L_2 error as a function of grid spacing. Right: temporal refinement of the L_2 error as a function of time step. Solid lines show the default regularised Koren MUSCL configuration, dashed coloured lines show the monolithic fully implicit first-order upwind replay configuration, and grey dotted lines indicate first- and second-order reference slopes.

This check establishes consistency for a scalar control with two further limitations. Pure amplitude scaling leaves limiter slope ratios invariant: the limiter-derivative contribution, although evaluated by the tangent-linear and adjoint kernels, multiplies an analytically zero perturbation direction and is not exercised here. Distributed nonsmooth-configuration gradient experiments probe that term separately. The check also covers only the explicit solver; the implicit adjoint, whose accuracy depends on the forward GMRES tolerance, is outside its scope. Gradient scaling at operational resolutions ($100 \times 80 \times 25$ or larger) and realistic multi-parameter inversions remain future work.

6.2 Independent evaluation against ETEX-1

Figure 7 shows lowest-model-level PMCH concentrations at two times in the ETEX-1 configuration of Sect. 5.6. The simulated plume follows the observed eastward and north-eastward pathway across western and central Europe, forming a coherent elongated structure.




Figure 6. Scalar structural gradient check in the manufactured Gaussian configuration. Left: relative difference between the centred finite-difference gradient and each direct gradient evaluation (analytic, tangent-linear, adjoint, forward-mode AD) as a function of perturbation size ϵ ; because the objective is exactly quadratic in α for this setup, the decline with increasing ϵ reflects single-precision cancellation rather than a truncation–roundoff trade-off. Right: gradient-cost scaling reported as the wall time versus the problem-size metric $N_{cells} \times N_t$.




Figure 7. Lowest-model-level PMCH concentrations (nominally 1000 hPa) at 24 October 03:00 UTC (left) and 26 October 08:20 UTC (right) during the ETEX-1 experiment as simulated by TARSA. The plume propagates eastward and north-eastward from the release site in western France.



6.2.1 Station metrics: maxima, dosage, and arrival time

For each station we compute the maximum observed and modelled concentration over its retained set of quantified, quality-controlled, paired ETEX-1 intervals. Following common practice in ETEX intercomparison studies, we focus on summary time-series diagnostics rather than comparing all 3-hour interval concentrations directly, because pointwise comparisons are strongly dominated by near-zero values and non-detections around the 0.05 ng m^{-3} detection limit and are highly sensitive to small phase errors in plume timing. As an additional robust amplitude metric, we also consider the station-wise mean concentration over the same retained interval set. Figure 8 illustrates the underlying paired 3-hour interval series at a representative station (DK02, Denmark), showing the arrival, peak, and decay phases that these station-wise diagnostics summarise.

Table 8. ETEX-1 station statistics using the screening of Sect. 5.6. Amplitude and duration use all 124 retained stations; arrival and peak time use the 80 paired stations where both series reach 0.05 ng m^{-3} . FM is the scalar figure of merit for each observable, distinct from station-wise temporal overlap (FMT). Appendix H defines all metrics.

Observable	N	correlation	FM	NMSE	FB
Mean conc.	124	0.661	0.696	1.722	0.282
Dosage	124	0.652	0.682	2.108	0.307
Arrival time	80	0.939	0.939	0.028	0.082
Duration	124	0.802	0.807	0.329	0.205
Maximum	124	0.613	0.652	2.070	0.068
Peak time	80	0.950	0.948	0.019	0.084

Sampling uncertainty was quantified with 10 000 non-parametric station-bootstrap resamples and documented fixed seeds. Approximate 95 % correlation intervals are 0.90–0.96 for arrival time ($r = 0.939$, $N = 80$), 0.53–0.79 for dosage ($r = 0.652$, $N = 124$), and 0.49–0.75 for maxima ($r = 0.613$, $N = 124$). Thus timing is the most robust aspect of this evaluation, while amplitude agreement is less tightly constrained. The role of ERA5 forcing in these results is discussed in Sect. 7.

TARSA reproduces plume timing more closely than concentration amplitude (Table 8; Fig. 9). Arrival and peak-time correlations are 0.939 and 0.950 for the 80 paired hit-only stations. Across all 124 retained stations, correlations are 0.613 for maxima, 0.661 for mean concentration, 0.652 for dosage, and 0.802 for duration. Maxima span more than three orders of magnitude; their small fractional bias coexists with substantial station-level error. Mean concentration and dosage have positive biases, with amplitude errors plausibly reflecting sensitivity to plume placement and timing.

Table 9 places TARSA's dosage and arrival-time statistics alongside published Polyphemus (Quélo et al., 2007) and DREAM (Brandt et al., 2000) results. Differences in forcing, resolution, screening, and metric definitions make this a contextual comparison; the caption details the FM convention issue.

The station-wise FMT comparison (Fig. 10) complements the aggregate diagnostics with temporal concentration-history overlap. Across 11 published comparison stations, TARSA's FMT ranges from about 17 to 57 %, with the highest values at




Figure 8. ETEX-1 time series at station DK02 (Albuen, Denmark; 54.50° N, 10.58° E). Observed PMCH concentrations are reported as 3-hour interval means; TARSA values are sampled and aggregated over the corresponding 3-hour windows.




Figure 9. Comparison of TARSA against ETEX-1 station measurements. Left: plume arrival time (hours since release). Centre: maximum observed versus simulated PMCH concentration. Right: dosage (time-integrated concentration at or above 0.05 ng m⁻³). The 1:1 line is shown for reference in each panel.

525 NL1, DK5, and NL5. Agreement is spatially variable and broadly comparable to the reference models at some receptors, without uniform superiority.

ETEX-1 supports credible plume pathway and timing reconstruction, with larger uncertainty in local maxima and integrated exposure. Section 7 considers the source representation, forcing, and mixing mechanisms that may explain this difference.

6.2.2 Supplementary CAPTEX-1 field comparison.

530 As a short supporting field-release check outside Europe, we also ran the first CAPTEX release over eastern North America (Ferber et al., 1986). In this case, 208 kg of PMCH is released near Dayton, Ohio (39.80° N, 84.05° W), from 18 September 1983 17:00 UTC for 3 h, and the comparison is made against interval-averaged station observations using the same general TARSA transport core under ERA5 forcing. The configuration mirrors the refined ETEX-1 setup: ERA5 pressure-level forcing for 38–48° N, 85–70° W at the native 0.25° resolution (61 × 41 horizontal points, 25 levels), bilinearly refined by a factor of two



Table 9. Contextual comparison of ETEX-1 dosage and arrival-time statistics for TARSA and two published models. Polyphemus values are taken from [Quélo et al. \(2007\)](#). DREAM values are taken from the ETEX-1 dosage and arrival-time diagnostics reported by [Brandt et al. \(2000\)](#). Owing to differences in model configuration, meteorological forcing, resolution, station screening, and observable definitions, the table is indicative rather than a like-for-like benchmark. The FM row deserves particular caution: the TARSA values use the Dice form of Eq. (H3) (equivalently 0.52 for dosage and 0.89 for arrival time in the min/max overlap form, Appendix H), and the convention behind the published values is not always stated.

Observable	Metric	TARSA	Polyphemus	DREAM
Dosage	N	124	139	140
	correlation	0.652	0.80	0.65
	FM	0.682	0.32	0.509
	NMSE	2.108	6.27	1.137
	FB	0.307	0.93	0.043
Arrival time	N	80	139	140
	correlation	0.939	0.94	0.97
	FM	0.939	0.90	0.900
	NMSE	0.028	0.13	0.031
	FB	0.082	−0.023	0.093




Figure 10. Station-wise comparison of the figure of merit in time (FMT) for ETEX-1. TARSA is compared with the published LMD-ZT(high) results from [Idelkadi et al. \(2002\)](#) and DERMA results from [Sørensen et al. \(2007\)](#). Higher FMT indicates better temporal overlap between the observed and simulated concentration histories at a station.



in the horizontal and vertical with the model top restricted to levels below about 8 km, regularised Koren MUSCL advection
 535 with SSP-RK3 time integration and CFL ≈ 0.95 sub-stepping, a 1800 s model step, implicit vertical diffusion, and the release
 injected into the lowest model layer with the same mass-conserving bilinear stencil as in ETEX-1. Unlike ETEX-1, this case
 is not treated as a main evaluation section in the paper; its role is to provide an additional observation-space cross-check under
 a different geography and sampling workflow. For the current refined-grid baseline, the pointwise interval-average comparison
 uses 395 valid station–interval pairs and gives correlation $R = 0.760$ ($R^2 = 0.577$), bias 0.241 ng m^{-3} , RMSE 1.174 ng m^{-3} ,
 540 and a composite rank of 2.187 on the 0–4 scale used by Selvaratnam et al. (2023), defined in Appendix H. As detailed there,
 the FMS component of this rank uses a percent-agreement plume-classification convention rather than the union-based figure
 of merit in space of the classical protocols, which inflates the rank relative to the published convention; this adds to the reasons
 the cross-model comparison below is indicative only.




Figure 11. Supplementary CAPTEX-1 average-rank comparison. The TARSA CAPTEX rank from the present repository workflow is shown against the published average ranks for NAME, HYSPLIT, STILT, and FLEXPART under the pb11 and pb10 configurations reported by Selvaratnam et al. (2023). Because the TARSA value is derived from the current repository evaluation setup rather than the exact protocol of the published intercomparison — in particular, its FMS component uses the percent-agreement convention of Appendix H rather than the union-based figure of merit in space, which raises the rank — the comparison is indicative rather than a strict like-for-like benchmark.

Figure 11 shows that the TARSA CAPTEX score lies within the published model range and near the lower half of the
 545 average-rank cluster. This does not establish a controlled benchmark against those systems, but it indicates that the present
 CAPTEX-1 configuration produces skill levels that are at least of a credible magnitude relative to established long-range
 dispersion models.

6.3 Fennec dust-column remote-sensing anchor case

The Fennec results use the primary configuration of Sect. 5.7. Figure 12 shows the GRASP/POLDER reference fields during
 550 the evaluation window.

The retrieval shows spatially extensive Saharan dust loading and substantial daily changes in plume extent and intensity, providing the reference for the collocated scatter and map comparisons.




Figure 12. Daily GRASP/POLDER Level-3 fields of Insoluble_Volume_Concentration_C over the Fennec domain during 16–24 June 2011. This retrieved coarse-insoluble component is used here as a dust-dominated column proxy, not as an independent ground-truth dust measurement.




Figure 13. Collocated column volume of insoluble coarse particles against GRASP/POLDER. Left: assimilative MERRA-2 dust reanalysis. Right: the primary TARSA forward run (Sect. 5.7), using summed DU001–DU004 without aerosol state assimilation. The two products use their respective valid collocation samples.

The primary TARSA run has positive spatial association with GRASP/POLDER ($r = 0.429$) and a positive column-volume bias of $0.094 \mu\text{m}^3 \mu\text{m}^{-2}$ (Fig. 13; Table 10). MERRA-2 has higher correlation ($r = 0.629$), larger bias, and similar RMSE; overestimation against GRASP/POLDER is also reported by Li et al. (2023). This is a contextual comparison: MERRA-2 includes aerosol data assimilation, and the two products have different valid samples (47 863 TARSA and 32 048 MERRA-2 collocations). These statistics support moderate agreement of the TARSA forward run with the retrieved proxy, without establishing a controlled ranking against MERRA-2.

Table 10. Collocated statistics for column volume concentration of insoluble coarse particles against GRASP/POLDER for the MERRA-2 assimilative dust product and the primary TARSA forward configuration. The TARSA quantity is the summed DU001–DU004 column volume from the four-day spin-up run with homogeneous Neumann lateral dust boundaries. Bias, root-mean-square error (RMSE), and mean absolute error (MAE) are in the column-volume units of the GRASP/POLDER product ($\mu\text{m}^3 \mu\text{m}^{-2}$); r is dimensionless. Each row uses the valid collocations available for the corresponding product, so the two rows are not computed on an identical sample.

Model	N	bias	RMSE	MAE	r
TARSA spin-up Neumann	47 863	0.094	0.213	0.152	0.429
MERRA-2 assimilative product	32 048	0.137	0.209	0.167	0.629

Under matched POLDER sampling, the daily triptych for 18 June 2011 (Fig. 14) shows that TARSA reproduces the broad regional dust distribution, with remaining differences in local maxima and plume sharpness.




Figure 14. Observation-weighted daily mean dust-column comparison for 18 June 2011. GRASP/POLDER, the primary TARSA spin-up Neumann run using the summed DU001–DU004 column volume, and the MERRA-2 reference product are averaged with the same hourly POLDER observation-count weights; the median observation time over the domain is 14:07 UTC.

Four process-sensitivity runs share the primary setup and collocation sample: no deposition (D), wet only (C), dry only (E), and dry plus wet (B), as summarised in Table 11.

Table 11. Fennec deposition sensitivity for the primary setup, comparing summed DU001–DU004 column volume with GRASP/POLDER on the same collocated sample. Bias, RMSE, and MAE are in $\mu\text{m}^3 \mu\text{m}^{-2}$; r is dimensionless.

Run	dry dep.	wet dep.	N	bias	RMSE	MAE	r
D	no	no	47 863	0.382	0.770	0.406	0.371
C	no	yes	47 863	0.368	0.757	0.393	0.368
E	yes	no	47 863	0.100	0.216	0.155	0.433
B	yes	yes	47 863	0.094	0.213	0.152	0.429

The process-sensitivity diagnostic indicates that dry deposition provides the dominant reduction in positive dust-column bias for this configuration. Relative to the no-deposition run (Run D), adding wet deposition alone (Run C) only slightly reduces the bias, RMSE, and MAE. Adding dry deposition (Run E) reduces the bias from 0.382 to 0.100 and the RMSE from 0.770 to 0.216, while also increasing the correlation from 0.371 to 0.433. The combined dry- and wet-deposition run (Run B) gives the lowest bias, RMSE, and MAE, whereas the dry-only run has the slightly higher correlation.

The Fennec results support regional aerosol-column evaluation under prescribed emissions, with dry deposition providing the dominant bias reduction in this configuration. Zero-gradient lateral boundaries can under-represent externally advected dust; interpretation should therefore emphasise the post-spin-up window and regions less affected by boundary inflow. The case does not independently validate dust emissions or the full aerosol process set.



7 Discussion

ERA5 forcing contributes to the skill and uncertainty of all real-case experiments. Evaluations against European wind observations report broad agreement alongside regional and seasonal biases; for example, [Molina et al. \(2021\)](#) analysed 245 HadISD stations and found systematic winter–summer differences. ETEX-1’s strong arrival-time correlation therefore reflects both the meteorological flow and TARSA’s transport, and cannot establish numerical superiority over other models. Forcing errors likewise propagate into the Fennec and CAMS simulations.

Several mechanisms could explain ETEX-1’s stronger timing than amplitude agreement. A near-point release represented on a regional grid makes maxima and dosage sensitive to source placement and unresolved plume structure. Refining ERA5 pressure levels to 49 adds no independent meteorological information, while zero-order-hold forcing may introduce timing errors at hourly transitions. The absence of terrain handling affects Alpine near-surface transport, and the direct ETEX drivers retain an advectively open 1000 hPa lower face; sensitivity to an impermeable replacement has not been isolated. These mechanisms remain hypotheses requiring targeted tests.

For Fennec, the local-stability mixing closure lacks a dedicated convective mixed-layer treatment (Appendix D). Errors in vertical dust distribution can affect sedimentation and size-dependent deposition and hence the column comparison. MERRA-2’s finer vertical discretisation and assimilation may contribute to its higher correlation with GRASP/POLDER, while its prescribed emissions, at an effective scale of order 100 km, must be remapped into TARSA’s regional grid. The present comparison does not isolate these contributions.

Native ERA5 model-level input, a convective boundary-layer option, and time-interpolated meteorology are priorities for testing these explanations. An online dust-emission model and aerosol state or emission correction would reduce reliance on prescribed MERRA-2 sources. Repeating the experiments with alternative forcing, such as MERRA-2 or JRA-3Q, would help separate meteorological uncertainty from discretisation and parameterisation effects. Scaling the gradient machinery to operational grids and realistic multi-parameter controls remains the main step towards the intended inversion and data-assimilation applications.

A further development priority is the generation of Level 4 (L4) products, understood here as model-derived analyses and forecasts combining satellite retrievals, ground-based measurements, and atmospheric transport. TARSA is intended to serve as the transport core for such products, including three-dimensional pollution maps and particulate-matter forecasts using GAPMAP and other satellite observations. Progress towards these services requires assimilation experiments with realistic satellite sampling and retrieval uncertainties, together with independent evaluation of near-surface PM_{2.5} and PM₁₀. The current dust-column comparison provides an initial methodological step; surface particulate-matter performance remains to be established.



8 Conclusions

TARSA v1.0 provides a reproducible regional Eulerian forward operator for inert-tracer and simplified aerosol studies. Its conservative finite-volume core and observation operators support controlled transport experiments, station sampling, and
605 satellite collocation under prescribed meteorology.

Closed-box tests conserve tracer mass to numerical precision, analytical plume comparisons show good agreement, and manufactured-solution errors decrease under mesh refinement. The regularised Koren limiter reduces selected gradient-error roughness; analytic, tangent-linear, adjoint, and automatic-differentiation gradients agree in the scalar emission-amplitude check, within its stated limits. ETEX-1 demonstrates stronger agreement in plume timing than concentration amplitude, sup-
610 ported by CAPTEX-1 under a different release geometry. Fennec extends the evaluation to dust columns under prescribed emissions, with moderate GRASP/POLDER agreement and a dominant bias reduction from dry deposition in the tested configuration. The supplementary CAMS replay demonstrates consistency under imposed concentration constraints.

These results establish the model's forward capabilities and identify priorities for improving vertical meteorological resolution, boundary-layer mixing, and aerosol sources. The medium-term goal is to support PM_{2.5} and PM₁₀ monitoring and
615 forecasting through L4 products, including three-dimensional pollution maps and forecasts informed by satellite observations. Within emission-inversion workflows, TARSA's forward operator and sensitivity machinery are intended to help improve anthropogenic emission inventories and quantify industrial source contributions. Realistic multi-parameter inversions, independent surface particulate-matter validation, and operational performance remain to be demonstrated.

Code and data availability. TARSA is free software distributed under the GNU Affero General Public License (AGPL-3.0-or-later). We
620 permanently archived the exact version used for all results in this paper (release v1.0.2) on Zenodo (Kuznetsov, 2026) at <https://doi.org/10.5281/zenodo.21932443>. This self-contained archive includes all experiment scripts, configuration files, cached ERA5 forcing, and station observations for both ETEX-1 and CAPTEX-1. It also features `reproduce_figures.md`, which maps every figure in the paper to the specific command needed to regenerate it. Ongoing development is hosted at <https://code.grasp-open.com/grasp-models/tarsa> (requires a GRASP-Open account), while user documentation is maintained at <https://tarsa.grasp-sas.com>.

625 External meteorological and aerosol input data were sourced from several dataset-specific repositories. We obtained ERA5 hourly pressure-level and single-level fields (Hersbach et al., 2020) from the Copernicus Climate Data Store via <https://doi.org/10.24381/cds.bd0915c6> (Hersbach et al., 2023a) and <https://doi.org/10.24381/cds.adbb2d47> (Hersbach et al., 2023b), respectively. The CAMS global reanalysis (EAC4) (Inness et al., 2019) is provided by the Copernicus Atmosphere Data Store at <https://ads.atmosphere.copernicus.eu/datasets/cams-global-reanalysis-eac4> (dataset DOI: <https://doi.org/10.24381/d58bbf47>). Additionally, MERRA-2 dust-emission fluxes and three-
630 dimensional dust mixing ratios (Gelaro et al., 2017; Randles et al., 2017) can be accessed through NASA GES DISC within the M2T1NXADG (<https://doi.org/10.5067/HM00OHQBHKTP>) and M2I3NVAER (<https://doi.org/10.5067/LTVB4GPCOTK2>) collections.



Appendix A: Advection flux details and regularised Koren limiter check

A1 First-order upwind face flux

The same face notation is used by both advection options. For an oriented face f with area S_f and unit normal $\mathbf{n}_f$, let

$$Q_f = S_f(\mathbf{u}_f \cdot \mathbf{n}_f), \quad \mathcal{G}_f = \rho_f Q_f C_f, \quad (A1)$$

where Q_f is the signed volume flux, ρ_f is the face air density, and $\mathcal{G}_f$ is the species-mass flux in the direction of $\mathbf{n}_f$. At a face shared by an owner cell P and a neighbour cell N , with $\mathbf{n}_f$ pointing from P to N , the first-order upwind flux is usually written with the positive and negative parts of Q_f . TARSA uses a smooth flux split to avoid an explicit velocity-direction branch. For a small volume-flux scale $\epsilon_Q > 0$, define

$$Q_f^{+, \epsilon} = \frac{1}{2} \left(Q_f + \sqrt{Q_f^2 + \epsilon_Q^2} \right), \quad Q_f^{-, \epsilon} = \frac{1}{2} \left(Q_f - \sqrt{Q_f^2 + \epsilon_Q^2} \right), \quad (A2)$$

and use the smoothly split upwind flux

$$\mathcal{G}_f^{\text{UP}} = \rho_f \left(Q_f^{+, \epsilon} C_P + Q_f^{-, \epsilon} C_N \right). \quad (A3)$$

As $\epsilon_Q \rightarrow 0$, $Q_f^{+, \epsilon} \rightarrow \max(Q_f, 0)$ and $Q_f^{-, \epsilon} \rightarrow \min(Q_f, 0)$, recovering the standard upwind flux. For finite ϵ_Q , the transition across flow reversal is continuously differentiable. At exact flow reversal the smoothed split gives a small diffusive contribution proportional to $\epsilon_Q(C_P - C_N)$, which vanishes with ϵ_Q and is negligible away from very small face-normal velocities. In the implementation, the smoothing is applied to the face-normal velocity $\mathbf{u}_f \cdot \mathbf{n}_f$ itself, before scaling by the face area S_f , using a single fixed velocity-scale default $\delta = 10^{-6} \text{ m s}^{-1}$; in terms of $Q_f = S_f(\mathbf{u}_f \cdot \mathbf{n}_f)$ this is equivalent to $\epsilon_Q = S_f \delta$, i.e. the effective volume-flux-scale ϵ_Q used above is proportional to the local face area rather than a single fixed volume-flux constant. This default is used identically by every advection kernel variant (plain and mass-consistent; forward, tangent, and adjoint). In the fixed-meteorology experiments reported here this split is fixed by the prescribed wind field; it becomes important mainly when differentiating with respect to wind components or parameters that can change the sign of Q_f . For a cell-wise finite-volume update the fluxes are summed with each cell's outward normal, so the same species-mass flux is removed from one cell and added to the adjacent cell with the opposite sign. When the update is written for C_p rather than for $\rho_p C_p$, the corresponding advective contribution appears as $(\rho_f / \rho_p) Q_f C_f$.

A2 Regularised Koren MUSCL reconstruction

For the positive-velocity orientation in Fig. A1, let x denote the coordinate along the one-dimensional stencil axis. With centre-to-centre spacings $\Delta x_U = x_P - x_U$ and $\Delta x_D = x_D - x_P$, define

$$g_P^- = \frac{C_P - C_U}{\Delta x_U}, \quad g_P^+ = \frac{C_D - C_P}{\Delta x_D}. \quad (A4)$$

For $g_P^+ \neq 0$, TARSA forms $r_P = g_P^- / g_P^+$ and defines the limited slope as

$$\sigma_P = \phi(r_P) g_P^+, \quad (A5)$$




Figure A1. Scheme of the regularised Koren MUSCL face reconstruction. For positive face-normal velocity, cell P is adjacent to the face on the upwind side, U is the next upstream cell, and D is the downwind cell. The stencil is drawn along a representative structured-grid axis; the lower coordinate axis marks x_U , x_P , x_f , and x_D . The same construction is applied along the corresponding grid axis for faces normal to y or z . The MUSCL state C_f^+ is obtained by extrapolating from C_P with a limited slope σ_P . The limiter maps the gradient ratio r to the slope multiplier $\phi(r)$; the regularised version replaces the two positive-ratio branch transitions of the classical Koren limiter by small C^1 transition patches and was checked to remain within the TVD admissible region. For negative face-normal velocity the same three-cell stencil is mirrored.

with $\sigma_P = 0$ if the stencil is incomplete. At the face f between P and D, the positive-velocity interface state is reconstructed as

$$C_f^+ = C_P + (x_f - x_P)\sigma_P. \quad (\text{A6})$$

The negative-velocity state C_f^- is obtained by applying the same three-cell construction after reversing the local face orientation.

The regularised limiter $\phi(r)$ uses the classical Koren branches away from the smoothed intervals and cubic-Hermite transition patches near $r = 1/4$ and $r = 5/2$. The Hermite interpolant is

$$H(r; a, b, y_a, m_a, y_b, m_b) = h_{00}(t)y_a + h_{10}(t)(b-a)m_a + h_{01}(t)y_b + h_{11}(t)(b-a)m_b, \quad (\text{A7})$$

where $t = (r - a)/(b - a)$ and

$$h_{00}(t) = 2t^3 - 3t^2 + 1, \quad h_{10}(t) = t^3 - 2t^2 + t, \quad h_{01}(t) = -2t^3 + 3t^2, \quad h_{11}(t) = t^3 - t^2. \quad (\text{A8})$$



Using $a_1 = \frac{1}{4} - \delta_1$, $b_1 = \frac{1}{4} + \delta_1$, $a_2 = \frac{5}{2} - \delta_2$, and $b_2 = \frac{5}{2} + \delta_2$, the limiter is

$$\phi(r) = \begin{cases} 0, & r \leq 0, \\ 2r, & 0 < r < a_1, \\ H\left(r; a_1, b_1, 2a_1, 2, \frac{1+2b_1}{3}, \frac{2}{3}\right), & a_1 \leq r < b_1, \\ \frac{1+2r}{3}, & b_1 \leq r < a_2, \\ H\left(r; a_2, b_2, \frac{1+2a_2}{3}, \frac{2}{3}, 2, 0\right), & a_2 \leq r < b_2, \\ 2, & r \geq b_2. \end{cases} \quad (\text{A9})$$

The default dimensionless transition widths are

$$(\delta_1, \delta_2) = (0.02, 0.05). \quad (\text{A10})$$

675 The implementation keeps the corner at $r = 0$ in its classical form; a diagnostic parameter δ_0 remains in the validation code only for compatibility with earlier tests. The limiter was checked to remain within the standard TVD admissible region,

$$0 \leq \phi(r) \leq 2, \quad \frac{\phi(r)}{r} \leq 2 \quad (r > 0), \quad \phi(1) = 1. \quad (\text{A11})$$

The local diagnostic is distributed with the model source (see the code and data availability section). It first checks the limiter and local face reconstruction directly, then runs a one-dimensional periodic passive-advection experiment on 96 cells
680 with constant positive velocity, SSP-RK3 stepping, CFL = 0.42, and 72 time steps (Fig. A2). The control parameter is the initial Gaussian source centre θ . The scalar response

$$J(\theta) = \frac{1}{2} \sum_m (C(x_m; \theta) - y_m)^2 \quad (\text{A12})$$

is evaluated at six downstream receptor points. In this diagnostic, y_m is generated from the analytical periodic translation of the same Gaussian at a shifted source position, and the analytical response supplies the reference derivative. The limiter
685 comparison includes classical Koren, regularised Koren, van Leer, first-order upwind, and a plain finite-volume PPM baseline; PPM is included only as a separate transport baseline, not as a Sweby limiter. All numerical curves use the same centred finite-difference estimate of $dJ/d\theta$ with $\varepsilon = 2.5 \times 10^{-4} \Delta x$.

The reported roughness metrics are applied to the gradient error

$$e_g(\theta) = \frac{dJ_{\text{scheme}}}{d\theta} - \frac{dJ_{\text{exact}}}{d\theta}, \quad (\text{A13})$$

690 so that smooth physical variation in the exact sensitivity is not counted as numerical noise. For the discrete sweep points θ_q , the diagnostic reports the maximum adjacent gradient-error jump, the total variation of the gradient error, and a second-difference roughness,

$$\max_q |e_{g,q+1} - e_{g,q}|, \quad \sum_q |e_{g,q+1} - e_{g,q}|, \quad \sum_q |e_{g,q+1} - 2e_{g,q} + e_{g,q-1}|. \quad (\text{A14})$$




Figure A2. Schematic setup of the one-dimensional periodic passive-advection diagnostic used for the limiter check. A Gaussian initial condition is shifted by the source-centre parameter θ on a periodic finite-volume grid, advected with constant positive velocity, and sampled at six downstream receptor points. The reference observations y_m are generated from the analytical periodic translation of the same Gaussian with a small source-centre offset.

The plotted roughness ratios are normalised by the corresponding classical Koren value.

695 In this controlled sweep (Fig. A3), the regularised Koren limiter leaves the forward response nearly unchanged relative to classical Koren: the maximum relative objective error against the analytical response is 0.38929 for classical Koren and 0.38936 for regularised Koren. Relative to classical Koren, the regularised limiter reduces the maximum adjacent gradient-error jump to about 0.84, the total variation of the gradient error to 0.89, and the second-difference gradient-error roughness to 0.77. First-order upwind gives very small second-difference roughness only through strong diffusion, with a larger maximum
700 relative objective error of 0.525 and a maximum relative gradient error of 0.861 against the analytical reference.

Appendix B: Tangent-linear, discrete-adjoint, and AD formulations

This appendix derives the sensitivity formulas used in Sects. 4 and 6.1.5, connecting multi-term least squares (Dubovik and King, 2000; Dubovik et al., 2011, 2021) to tangent-linear and discrete-adjoint propagation (Asch et al., 2016). All Jacobians are evaluated along the computed trajectory, covering both the affine upwind configuration and the trajectory-dependent regularised
705 Koren solver.

B1 Least-squares cost function

Let the observations available at model step n be related to the state by

$$\mathbf{y}_n = \mathbf{H}_n \mathbf{X}^n(\theta) + \varepsilon_n, \quad \varepsilon_n \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_n), \quad n = 0, \dots, N_t, \quad (\text{B1})$$




Figure A3. One-dimensional limiter diagnostic for the regularised Koren MUSCL reconstruction. Panels show the scalar response $J(\theta)$ for a swept Gaussian source centre, the centred finite-difference sensitivity $dJ/d\theta$, normalised gradient-error roughness metrics, and the maximum forward-response error relative to the analytical translated-Gaussian reference. The regularised Koren limiter reduces selected gradient-error roughness measures relative to classical Koren while leaving the forward response nearly unchanged. First-order upwind gives lower second-difference roughness only by stronger numerical diffusion, while PPM is included as a separate parabolic-reconstruction transport baseline.



where $\mathbf{H}_n$ is the observation operator (its Jacobian, if it is nonlinear) and $\mathbf{R}_n$ the symmetric positive-definite observation-error
710 covariance; errors at different steps are independent, and a step without observations is simply omitted from the sums below.
Denote by $\mathbf{d}_n = \mathbf{H}_n \mathbf{X}^n(\theta) - \mathbf{y}_n$ the residual, written model minus data. Let an a priori estimate θ^b of the controls be available,
with an error independent of the ε_n ,

$$\theta^b = \theta + \varepsilon_a, \quad \varepsilon_a \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_a). \quad (\text{B2})$$

In the multi-term LSM the prior is treated as one more data set, so that, by independence, the likelihood of all data given θ is a
715 product of Gaussians,

$$p(\mathbf{y}_0, \dots, \mathbf{y}_{N_t}, \theta^b | \theta) \propto \prod_{n=0}^{N_t} \exp\left(-\frac{1}{2} \mathbf{d}_n^T \mathbf{R}_n^{-1} \mathbf{d}_n\right) \exp\left(-\frac{1}{2} (\theta - \theta^b)^T \mathbf{C}_a^{-1} (\theta - \theta^b)\right) = \exp\{-\mathcal{J}(\theta)\}, \quad (\text{B3})$$

and the maximum-likelihood estimate of θ minimises the quadratic form in the exponent,

$$\mathcal{J}(\theta) = \frac{1}{2} \sum_{n=0}^{N_t} \mathbf{d}_n^T \mathbf{R}_n^{-1} \mathbf{d}_n + \frac{1}{2} (\theta - \theta^b)^T \mathbf{C}_a^{-1} (\theta - \theta^b). \quad (\text{B4})$$

Three remarks fix the conventions. (i) A smoothness constraint on the control field is a data set of the same kind, with $\theta^b =$
720 $\mathbf{0}$ and $\mathbf{C}_a^{-1} = \gamma \mathbf{D}^T \mathbf{D}$ for a finite-difference operator $\mathbf{D}$ and a weight γ ; only the relative weights of the terms affect the
minimiser (Dubovik et al., 2021). (ii) Equation (B4) is also the strong-constraint 4D-Var cost function, with $\mathbf{C}_a$ in the role of
the background covariance (Asch et al., 2016). The residual is written model minus data, as in Dubovik et al. (2008); Asch et al.
(2016) use the innovation $\mathbf{y}_n - \mathbf{H}_n \mathbf{X}^n$, so their observation terms carry the opposite sign. (iii) The scalar check of Sect. 6.1.5
is the special case of a single control α , $\mathbf{R}_n = \mathbf{I}$, and no prior term.

725 By the chain rule, the gradient of Eq. (B4) is

$$\nabla_{\theta} \mathcal{J} = \sum_{n=0}^{N_t} \mathbf{K}_n^T \mathbf{R}_n^{-1} \mathbf{d}_n + \mathbf{C}_a^{-1} (\theta - \theta^b), \quad (\text{B5})$$

$$\mathbf{K}_n = \frac{\partial \mathbf{d}_n}{\partial \theta} = \mathbf{H}_n \mathbf{S}^n, \quad \mathbf{S}^n = \frac{\partial \mathbf{X}^n}{\partial \theta},$$

where $\mathbf{K}_n$ is the Jacobian of the residual and $\mathbf{S}^n$ the $N_X \times N_{\theta}$ sensitivity matrix of the state; this is the form of Eq. (40) of
Dubovik et al. (2008). For a forward map that is affine in θ , setting Eq. (B5) to zero gives the normal equations of the LSM,
whose solution requires the assembled $N_{\theta} \times N_{\theta}$ matrix $\sum_n \mathbf{K}_n^T \mathbf{R}_n^{-1} \mathbf{K}_n + \mathbf{C}_a^{-1}$. For a large control vector, or for the trajectory-
730 dependent Koren solver, the minimiser is instead approached iteratively by Gauss–Newton or steepest-descent steps (Dubovik
et al., 2008, Eqs. 39–40), which at every iterate require only $\nabla_{\theta} \mathcal{J}$, i.e. products of $\mathbf{K}_n^T = (\mathbf{S}^n)^T \mathbf{H}_n^T$ with a vector. There
are two ways to obtain them. The tangent-linear model (Sect. B2) builds $\mathbf{S}^n$ column by column, at the cost of one forward
sensitivity integration per control, and is economical when N_{θ} is small. The adjoint (Sect. B3) evaluates the sum $\sum_n (\mathbf{S}^n)^T (\cdot)$
without forming any $\mathbf{S}^n$, at a cost independent of N_{θ} .



735 B2 Tangent-linear model

Differentiating the one-step map $\mathbf{X}^{n+1} = \mathbf{F}_n(\mathbf{X}^n, \theta)$ with respect to θ along the trajectory computed for the current θ gives the tangent-linear recursion of Eq. (11),

$$\begin{aligned} \mathbf{S}^{n+1} &= \mathbf{J}_n \mathbf{S}^n + \mathbf{B}_n, \quad n = 0, \dots, N_t - 1, \quad \mathbf{S}^0 = \frac{\partial \mathbf{X}^0}{\partial \theta}, \\ \mathbf{J}_n &= \left. \frac{\partial \mathbf{F}_n}{\partial \mathbf{X}^n} \right|_{(\mathbf{X}^n, \theta)}, \quad \mathbf{B}_n = \left. \frac{\partial \mathbf{F}_n}{\partial \theta} \right|_{(\mathbf{X}^n, \theta)}, \end{aligned} \quad (\text{B6})$$

where $\mathbf{S}^0$ is zero unless the initial condition is part of the control vector. The $N_X \times N_X$ matrix $\mathbf{J}_n$ propagates a state perturbation
740 through one model step, and the $N_X \times N_\theta$ matrix $\mathbf{B}_n$ injects a control perturbation into that step. Unrolling the recursion,

$$\mathbf{S}^n = \mathbf{J}_{0 \rightarrow n} \mathbf{S}^0 + \sum_{j=0}^{n-1} \mathbf{J}_{j+1 \rightarrow n} \mathbf{B}_j, \quad \mathbf{J}_{j \rightarrow n} \equiv \mathbf{J}_{n-1} \mathbf{J}_{n-2} \cdots \mathbf{J}_j, \quad \mathbf{J}_{n \rightarrow n} \equiv \mathbf{I}, \quad (\text{B7})$$

where $\mathbf{J}_{j \rightarrow n}$ propagates a state perturbation from step j to step n . A control acting at step j influences the state only at later steps; stacked over n , Eq. (B7) is the block lower-triangular transport matrix of Dubovik et al. (2008, Eq. 5).

In the first-order upwind replay configuration, $\mathbf{J}_n$ and $\mathbf{B}_n$ are the fixed sparse step operators of Sect. 4. In the regularised Ko-
745 ren configuration the update depends on the tracer field also through the limiter coefficients, $\mathbf{F}_n(\mathbf{X}^n, \theta) = \hat{\mathbf{F}}_n(\mathbf{X}^n, \ell_n(\mathbf{X}^n), \theta)$, so that

$$\mathbf{J}_n = \left. \frac{\partial \hat{\mathbf{F}}_n}{\partial \mathbf{X}^n} \right|_{\ell_n \text{ fixed}} + \frac{\partial \hat{\mathbf{F}}_n}{\partial \ell_n} \frac{\partial \ell_n}{\partial \mathbf{X}^n}. \quad (\text{B8})$$

The first term is the transport operator with the limiter frozen at the reference state; the second is the limiter derivative, which the C^1 regularisation of Appendix A makes continuous. Both are evaluated at the stored forward state and must be recomputed
750 whenever the optimiser changes the trajectory, as in the relinearised iteration of Dubovik et al. (2008, Sect. 2.7). Keeping only the first term gives the exact derivative of an affine surrogate, not that of the nonlinear TVD update. The tangent-linear path of Sect. 6.1.5 propagates Eq. (B6) for the single column $\mathbf{S}^n = \partial \mathbf{X}^n / \partial \alpha$.

B3 Gradient by direct differentiation and the adjoint state

Write the observation part of the gradient in Eq. (B5) as

$$755 \quad \nabla_{\theta} \mathcal{J}^o = \sum_{n=0}^{N_t} (\mathbf{S}^n)^T \mathbf{w}_n, \quad \mathbf{w}_n \equiv \mathbf{H}_n^T \mathbf{R}_n^{-1} \mathbf{d}_n, \quad (\text{B9})$$

where $\mathbf{w}_n$ is the weighted residual mapped back to state space ($\mathbf{w}_n = \mathbf{0}$ at steps without observations). Substituting Eq. (B7) and exchanging the order of summation gives

$$\nabla_{\theta} \mathcal{J}^o = (\mathbf{S}^0)^T \sum_{n=0}^{N_t} \mathbf{J}_{0 \rightarrow n}^T \mathbf{w}_n + \sum_{j=0}^{N_t-1} \mathbf{B}_j^T \sum_{n=j+1}^{N_t} \mathbf{J}_{j+1 \rightarrow n}^T \mathbf{w}_n. \quad (\text{B10})$$



Both inner sums are instances of one quantity, which we define as the adjoint state at step j ,

$$\lambda^j \equiv \sum_{n=j}^{N_t} \mathbf{J}_{j \rightarrow n}^T \mathbf{w}_n, \quad j = 0, \dots, N_t. \quad (\text{B11})$$

Its meaning follows from the tangent-linear model: a perturbation $\delta \mathbf{X}^j$ of the state at step j changes the state at every later step n by $\mathbf{J}_{j \rightarrow n} \delta \mathbf{X}^j$, and the cost function by $\delta \mathcal{J}^o = \sum_{n \geq j} \mathbf{w}_n^T \mathbf{J}_{j \rightarrow n} \delta \mathbf{X}^j = (\lambda^j)^T \delta \mathbf{X}^j$, so λ^j is the gradient of the observation term with respect to the model state at step j . Separating the $n = j$ term in Eq. (B11) and using $\mathbf{J}_{j \rightarrow n}^T = \mathbf{J}_j^T \mathbf{J}_{j+1 \rightarrow n}^T$ for $n > j$ shows that the adjoint states satisfy the backward recursion

$$\lambda^{N_t} = \mathbf{w}_{N_t}, \quad \lambda^n = \mathbf{w}_n + \mathbf{J}_n^T \lambda^{n+1}, \quad n = N_t - 1, \dots, 0, \quad (\text{B12})$$

which, written out,

$$\lambda^0 = \mathbf{w}_0 + \mathbf{J}_0^T \left[\mathbf{w}_1 + \mathbf{J}_1^T \left[\mathbf{w}_2 + \dots + \mathbf{J}_{N_t-1}^T \mathbf{w}_{N_t} \right] \right]. \quad (\text{B13})$$

Equation (B12) is the discrete adjoint of the implemented time-stepping map: it runs backward in time, applies the transposed step Jacobians, and adds the weighted residual at each observation step. The λ^n coincide with the Lagrange multipliers of the constrained formulation (Le Dimet and Talagrand, 1986; Asch et al., 2016), but no Lagrangian is needed to obtain them. Inserting Eq. (B11) into Eq. (B10) and adding the prior term, the gradient of Eq. (B4) is

$$\nabla_{\theta} \mathcal{J} = \mathbf{C}_a^{-1} (\theta - \theta^b) + (\mathbf{S}^0)^T \lambda^0 + \sum_{n=0}^{N_t-1} \mathbf{B}_n^T \lambda^{n+1}. \quad (\text{B14})$$

One gradient evaluation therefore consists of (1) one forward integration, which stores the trajectory needed to evaluate $\mathbf{J}_n$ and $\mathbf{B}_n$ and the residuals $\mathbf{d}_n$; (2) one backward sweep of Eq. (B12); and (3) the projections $\mathbf{B}_n^T \lambda^{n+1}$ onto the controls, accumulated during the sweep, plus the prior term. The cost is independent of N_{θ} ; for the affine replay configuration only the meteorological fields are needed in reverse order.

Three special cases identify Eq. (B14) with familiar results.

- *Initial condition as control*, $\theta = \mathbf{X}^0$: then $\mathbf{S}^0 = \mathbf{I}$, $\mathbf{B}_n = \mathbf{0}$, and $\nabla \mathcal{J} = \mathbf{C}_a^{-1} (\mathbf{X}^0 - \mathbf{X}^{0,b}) + \lambda^0$. This is the strong-constraint 4D-Var gradient of Asch et al. (2016, Sect. 2.4.5); only the end point λ^0 of the sweep is used.
- *Source field at every step as control*, $\theta = (\mathbf{s}_0, \dots, \mathbf{s}_{N_t-1})$, entering as $\mathbf{X}^{n+1} = \mathbf{F}_n(\mathbf{X}^n + \Delta t \mathbf{s}_n)$ (Dubovik et al., 2008, Eq. 2): then $\mathbf{S}^0 = \mathbf{0}$ and $\partial \mathcal{J}^o / \partial \mathbf{s}_n = \mathbf{B}_n^T \lambda^{n+1} = \Delta t \mathbf{J}_n^T \lambda^{n+1}$, which is the second summand of Eq. (B12). The source gradient is obtained by storing, rather than discarding, the intermediate products of the same sweep; this is the backward integration of Dubovik et al. (2008, Eq. 41). In TARSA the source enters the split step of Eq. (9) after the advection and before the final diffusion half-step, and the stored quantity is the intermediate adjoint at that point (Sect. B4).
- *Scalar emission amplitude α of Sect. 6.1.5*: then $\mathbf{S}^0 = \mathbf{0}$ and $\mathbf{B}_n = \partial_{\alpha} \mathbf{F}_n$ is a single column, and the adjoint entry of Table 7 is $d\mathcal{J}/d\alpha = \sum_n (\partial_{\alpha} \mathbf{F}_n)^T \lambda^{n+1}$.



B4 Matrix-free evaluation, automatic differentiation, and verification

Neither $\mathbf{J}_n^T$ nor $\mathbf{H}_n^T$ is assembled. A model step is a composition of substeps, $\mathbf{F}_n = \mathbf{f}_{n,m_n} \circ \dots \circ \mathbf{f}_{n,1}$ (Strang halves, CFL-limited advection substeps, SSP-RK3 stages, and source and sink updates), so $\mathbf{J}_n^T$ is the product of the substep transposes in reverse order. The backward step $\lambda^{n+1} \mapsto \mathbf{J}_n^T \lambda^{n+1}$ of Eq. (B12) and the control projection $\mathbf{B}_n^T \lambda^{n+1}$ of Eq. (B14) are implemented together as the reverse pass

$$\begin{aligned} \mu &\leftarrow \lambda^{n+1}; \\ \text{for } j = m_n, \dots, 1 : \quad \mathbf{g} &\leftarrow \mathbf{g} + \left(\frac{\partial \mathbf{f}_{n,j}}{\partial \theta} \right)^T \mu, \quad \mu \leftarrow \left(\frac{\partial \mathbf{f}_{n,j}}{\partial \mathbf{X}} \right)^T \mu, \end{aligned} \quad (\text{B15})$$

where $\mathbf{g}$ accumulates the control gradient over all steps; on exit, $\mu = \mathbf{J}_n^T \lambda^{n+1}$. Each substep transpose is that of the implemented discrete stencil, including its boundary operations. For a one-dimensional upwind substep with Courant number ν , the forward update $X_i \leftarrow (1 - \nu)X_i + \nu X_{i-1}$ has the adjoint $\mu_i \leftarrow (1 - \nu)\mu_i + \nu\mu_{i+1}$, which carries sensitivity upstream; this is the exact transpose of the discrete operator, not a rerun of the forward scheme with reversed winds, i.e. a discrete rather than a continuous adjoint (Sandu et al., 2005; Liu and Sandu, 2008). Note that $\mathbf{B}_n$ and $\mathbf{B}_n^T$ are different, generally rectangular, mappings: the former injects a control perturbation into the state, the latter gathers state-space sensitivities into control space.

Automatic differentiation (AD) produces the same derivatives of the executed discrete program. Forward-mode AD carries dual numbers through the code and thereby propagates Eq. (B6) for one direction; reverse-mode AD applied to the scalar $\mathcal{J}$ generates the recursion of Eq. (B12) (Giering and Kaminski, 1998; Griewank and Walther, 2008). AD differentiates the program path actually executed, so limiter branches, boundary-condition branches, and non-smooth parameterisations determine the local derivative returned, and reverse mode needs the same storage or checkpointing of the forward trajectory as a hand-coded adjoint. In TARSA, the tangent-linear and adjoint kernels are hand-coded transposes of the forward kernels, and both AD routes of Sect. 3.9 are forward mode; the autodiff column of Table 7 is therefore an independent implementation of the tangent-linear path, and reverse-mode AD is not used in this paper.

Two consistency checks apply. The dot-product identity

$$\langle \mathbf{J}_n \delta \mathbf{X}, \mu \rangle = \langle \delta \mathbf{X}, \mathbf{J}_n^T \mu \rangle \quad (\text{B16})$$

must hold to round-off for arbitrary test vectors, and likewise for $\mathbf{H}_n$; it is checked for the transport kernels in the unit tests distributed with the code. The gradient test compares $\mathbf{v}^T \nabla_{\theta} \mathcal{J}$ with the centred difference $[\mathcal{J}(\theta + \varepsilon \mathbf{v}) - \mathcal{J}(\theta - \varepsilon \mathbf{v})] / (2\varepsilon)$, which is affected by truncation error at large ε and by round-off at small ε ; it is the finite-difference reference of Sect. 6.1.5, where the analytic, tangent-linear, adjoint, and AD gradients are compared with it and with one another.

Appendix C: Supplementary CAMS replay consistency experiment

This supplementary CAMS replay tests consistency under prescribed initial and boundary concentrations. It evaluates short-range evolution of imposed three-dimensional fields, without independently validating emissions, chemistry, or aerosol realism.



C1 Setup

The CAMS experiment uses the global reanalysis EAC4 data set from the Copernicus Atmosphere Monitoring Service (Inness et al., 2019) over a regional domain covering the Iberian Peninsula, the western Mediterranean, and north-western Africa. The domain is 30°–45° N, 10° W–5° E, at the native ERA5 0.25° pressure-level resolution used throughout this paper (a nominal 61 × 61 × 25 grid), forced with ERA5 and CAMS fields prepared for 7–10 July 2008 (94 hourly output snapshots, i.e. close to four days). TARSA is initialised from the CAMS three-dimensional tracer fields and uses CAMS throughout the simulation as time-varying lateral boundary conditions and lower-boundary concentrations. No additional emissions are prescribed inside the domain. The model is then integrated forward under ERA5 meteorology using the monolithic fully implicit backward-Euler configuration.

For long-lived CO, the replay probes preservation of large-scale structure over an advective residence time; aerosol tracers additionally exercise sedimentation and wet deposition over the following days.

C2 Results

Figure C1 summarises representative vertical-profile comparisons for selected species. For CO, the TARSA and CAMS profiles agree closely through much of the troposphere, indicating that the large-scale vertical structure is preserved well over the replay window. For aerosol species, larger discrepancies appear near the surface and in the upper troposphere, which is consistent with the simplified treatment of microphysics, ageing, and deposition in the present TARSA configuration and with differences between CAMS and ERA5 meteorology.

Figure C2 and Table C1 summarise gridpoint agreement. CO reaches $R^2 = 0.9002$ in linear space and 0.9274 in log space. These pooled statistics include all grid cells at all 94 hourly snapshots, including the initial field that equals CAMS by construction; early times and boundary-influenced cells therefore inherit agreement from the imposed constraints. No persistence or other trivial baseline is evaluated, so the transport contribution to this agreement is not isolated. Log-space sample sizes are smaller because non-positive concentrations are excluded before transformation.

Across aerosols, the degree of agreement depends on both the species and the metric space. Dust bins 1–2 remain highly consistent with CAMS, with $R^2 \approx 0.90$ in linear space and $R^2 \approx 0.92$ in log space. For carbonaceous aerosols, linear-space agreement is lower (black carbon $R^2 = 0.7457$, organic matter $R^2 = 0.6609$), while the log-space statistics improve substantially (black carbon $R^2 = 0.8917$, organic matter $R^2 = 0.8768$), indicating that the main discrepancies are concentrated in the high-concentration tail. Sea-salt bin 1 shows the opposite behaviour: strong linear-space agreement ($R^2 = 0.8933$) but weaker log-space agreement ($R^2 = 0.7050$), consistent with sensitivity to near-surface winds and small-scale variability.

The replay preserves the large-scale CO structure and produces species-dependent aerosol departures. Missing aerosol processes and forcing differences are plausible contributors, but this constrained experiment does not isolate their effects from transport errors.




Figure C1. Vertical-profile comparison between TARSA and CAMS for carbon monoxide (top), fine dust (second panel), black carbon (third panel), and organic matter (bottom).




Figure C2. Scatter comparison of TARSA-predicted concentrations versus CAMS reanalysis for several species: carbon monoxide (top left), black carbon (top right), fine dust (middle left), coarse dust bin 2 (middle right), organic matter (bottom left), and sea-salt bin 1 (bottom right).



Table C1. Summary statistics against CAMS for the species shown in Fig. C2. “Linear” denotes metrics computed in linear space, while “log” denotes metrics computed after log transformation.

Species	Linear space			Log space			Notes
	n	r	R^2	n	r	R^2	
CO	8744350	0.9488	0.9002	8398297	0.9630	0.9274	inert tracer (replay; no chemistry)
Black carbon	8744234	0.8635	0.7457	8396661	0.9443	0.8917	no aging/partitioning
Dust bin 1 (fine)	8744270	0.9492	0.9009	8397558	0.9579	0.9175	fine-mode transport captured
Dust bin 2 (coarse)	8744297	0.9474	0.8976	8397525	0.9574	0.9165	increased spread vs CAMS
Organic matter (OM)	8744211	0.8130	0.6609	8396677	0.9364	0.8768	no secondary formation
Sea-salt bin 1	8744307	0.9451	0.8933	8397910	0.8397	0.7050	sensitive to near-surface winds

Appendix D: Vertical diffusivity formulation

TARSA diagnoses turbulent transport from local stability using a first-order closure applied at every model layer, regardless of position relative to the diagnosed boundary-layer height. The vertical exchange coefficient K_z is defined from Monin–Obukhov-type similarity for unstable conditions ($Ri < 0$), from a near-neutral limit of the same shear-driven closure for $|Ri|$ close to zero, and from a Louis–Tiedtke–Geleyn closure for stable conditions ($Ri > 0$), closely following the simplified ECMWF treatment adopted in the main text.

D1 Local-stability closure above and within the boundary layer

Monin–Obukhov similarity for unstable layers ($Ri < 0$)

For statically unstable layers ($Ri < 0$), the exchange coefficient is written as

$$K_z = \frac{l^2}{\Phi_M(\zeta) \Phi_H(\zeta)} \left| \frac{\partial U}{\partial z} \right|, \quad (\text{D1})$$

where l is the mixing length, Φ_M and Φ_H are the Businger–Dyer stability functions in the forms $\Phi_M(\zeta) = (1 - 16\zeta)^{-1/4}$ and $\Phi_H(\zeta) = (1 - 16\zeta)^{-1/2}$ (Dyer, 1974), and ζ is approximated here by the local Richardson number Ri following the simplified similarity approach of Nieuwstadt (1984). In the present implementation the asymptotic mixing length for this unstable free-shear regime is set to $\lambda = 150$ m. This closure is applied identically to every unstable layer, whether situated within or above the diagnosed boundary-layer height: TARSA does not implement a separate convective mixed-layer parameterisation for the unstable regime (Sect. 7).



Near-neutral layers ($|Ri| < 0.05$)

865 Layers with $|Ri| < 0.05$ are treated as near-neutral and use the neutral limit of the same shear-driven closure, without either stability correction:

$$K_z = l^2 \left| \frac{\partial U}{\partial z} \right|, \quad (D2)$$

using the same asymptotic mixing length $\lambda = 150$ m as the unstable branch.

Revised Louis scheme for stable layers ($Ri > 0$)

For stable layers, the exchange coefficient is diagnosed as

$$870 \quad K_z = \left| \frac{\partial U}{\partial z} \right| l^2 f_H(Ri), \quad (D3)$$

where f_H denotes the Louis–Tiedtke–Geleyn stability function for heat and scalars (Louis, 1982), revised following Beljaars (1995); Beljaars and Viterbo (1998); Viterbo et al. (1999) to maintain a realistic level of mixing in stable boundary layers:

$$f_H(Ri) = \frac{1}{1 + 2b Ri (1 + d Ri)^{1/2}}. \quad (D4)$$

875 Because tracer mixing ratios are scalars, TARSA uses the heat form of the revised scheme rather than the momentum form $f_M(Ri) = [1 + 2b Ri (1 + d Ri)^{-1/2}]^{-1}$ (ECMWF, 2024a). The two functions differ by the factor $(1 + d Ri)$ inside the stability term, so that f_M/f_H grows from about 1.1 at $Ri = 0.2$ to about 1.9 at $Ri = 1$; applying the momentum function to a scalar would overestimate mixing in strongly stable layers by that factor.

The numerical values used by TARSA for Eq. (D4) are listed in Table D1.

Table D1. Constants used in the Louis–Tiedtke–Geleyn stability function for stable layers.

Parameter	Value	Units	Source
b	5	1	(ECMWF, 2024a; Louis, 1982; Beljaars, 1995; Viterbo et al., 1999)
d	1	1	(ECMWF, 2024a; Louis, 1982; Beljaars, 1995; Viterbo et al., 1999)

The mixing length is bounded using the Blackadar form

$$880 \quad \frac{1}{l} = \frac{1}{\kappa z} + \frac{1}{\lambda}, \quad (D5)$$

885 where κ is the von Kármán constant, set to $\kappa = 0.4$. In stable boundary layers, TARSA sets the asymptotic length scale λ to 10 % of the diagnosed boundary-layer height, taken directly from the ERA5 boundary-layer-height field (b1h); in stable layers above the boundary layer, a constant value $\lambda = 30$ m is used. This reflects the reduced turbulence length scale under shallow stable conditions while permitting somewhat deeper mixing in weakly stable layers. Consistent with the vertical-coordinate convention of Sect. 2.3, the height z entering Eq. (D5), and the height compared against the ERA5 boundary-layer height



when selecting λ in the stable regime, are evaluated from the archived geopotential without subtracting the surface elevation, i.e. as height above mean sea level rather than above ground. Over elevated terrain this weakens the near-surface κz limit and biases the stable-regime λ selection towards the above-boundary-layer value; a terrain-following formulation of this pathway is planned for a future release. TARSA applies no additional lower or upper bound to K_z itself in any regime; internally, only the inputs to Eqs. (D1)–(D2) and (D3) are floored and capped for numerical robustness, namely the squared velocity-difference term entering the shear (floored at $10^{-5} \text{ m}^2 \text{ s}^{-2}$), the resulting shear magnitude $|\partial U / \partial z|$ (capped at 1000 s^{-1}), and the Richardson number of Sect. D2 (clamped to $|Ri| \leq 10^3$).

D2 Richardson number estimation

Following the ECMWF IFS documentation (ECMWF, 2024a), the Richardson number at interface $k + \frac{1}{2}$ is estimated from the vertical shear and moist static stability:

$$|\Delta U|_{k+\frac{1}{2}}^2 = (u_k - u_{k+1})^2 + (v_k - v_{k+1})^2, \quad (\text{D6})$$

$$\left(\frac{\Delta s_v}{c_p T} \right)_{k+\frac{1}{2}} = \frac{2(s_k - s_{k+1})}{(s_k - gz_k) + (s_{k+1} - gz_{k+1})} + \varepsilon(q_k - q_{k+1}), \quad (\text{D7})$$

$$Ri_{k+\frac{1}{2}} = (gz_k - gz_{k+1}) \frac{\left(\frac{\Delta s_v}{c_p T} \right)_{k+\frac{1}{2}}}{|\Delta U|_{k+\frac{1}{2}}^2}, \quad (\text{D8})$$

where $s = gz + c_p T$. With this definition, Eq. (D7) can be rewritten in the more transparent form

$$\left(\frac{\Delta s_v}{c_p T} \right)_{k+\frac{1}{2}} = \frac{2(T_k - T_{k+1})}{T_k + T_{k+1}} + \frac{2g(z_k - z_{k+1})}{c_p(T_k + T_{k+1})} + \varepsilon(q_k - q_{k+1}), \quad (\text{D9})$$

in which the first term is the layer temperature contrast and the second is the dry-adiabatic correction associated with the layer thickness. These expressions are used diagnostically to classify local stability and to determine which branch of the K_z formulation is applied in each layer.

Appendix E: Aerosol-process parameterisations

TARSA v1.0 includes aerosol processes that can be represented as linear operators acting on the transported tracer: hygroscopic growth used diagnostically to update effective particle properties, gravitational settling, dry deposition, and wet deposition. Nonlinear microphysical processes such as nucleation, coagulation, and condensational growth are not solved prognostically.

E1 Hygroscopic growth

Following the κ -Köhler framework of Petters and Kreidenweis (2007), the wet radius of aerosol mode k is diagnosed as

$$r_{\text{wet},k} = r_{\text{dry},k} \text{gf}(RH, \kappa_k), \quad (\text{E1})$$



with growth factor

$$gf(RH, \kappa_k) = \left[1 + \kappa_k \frac{RH}{1 - RH} \right]^{1/3}. \quad (E2)$$

Relative humidity is diagnosed from temperature, pressure, and specific humidity:

$$RH = \frac{e}{e_s}, \quad e_s = 6.112 \exp \left(\frac{17.67 T_C}{T_C + 243.5} \right), \quad e = \frac{qp}{0.622 + 0.378q}, \quad (E3)$$

915 where $T_C = T - 273.15$ is temperature in °C and p is expressed in hPa.

The corresponding wet density is

$$\rho_{\text{wet},k} = \frac{\rho_{\text{dry},k} V_{\text{dry},k} + \rho_{\text{water}} (V_{\text{wet},k} - V_{\text{dry},k})}{V_{\text{wet},k}}, \quad (E4)$$

with

$$V_{\text{dry},k} = \frac{4}{3} \pi r_{\text{dry},k}^3, \quad V_{\text{wet},k} = \frac{4}{3} \pi r_{\text{wet},k}^3, \quad \rho_{\text{water}} \approx 1000 \text{ kg m}^{-3}. \quad (E5)$$

920 In TARSA this hygroscopic-growth step is diagnostic; it does not introduce additional prognostic variables.

E2 Sedimentation

The gravitational settling velocity is parameterised in the Stokes regime (Seinfeld and Pandis, 2016) as

$$w_{\text{sed},k} = \frac{2}{9} \frac{\rho_k g r_k^2}{\mu} C_F, \quad (E6)$$

925 where ρ_k and r_k are the effective particle density and radius, $\mu = 1.458 \times 10^{-6} T^{3/2} / (T + 110.4)$ is the dynamic viscosity of air from Sutherland's law, and C_F is the Cunningham slip correction:

$$C_F = 1 + \frac{\lambda_{\text{air}}}{r_k} \left[1.257 + 0.4 \exp \left(-\frac{1.1 r_k}{\lambda_{\text{air}}} \right) \right]. \quad (E7)$$

Here λ_{air} is the mean free path of air molecules, typically about $0.066 \mu\text{m}$ at standard conditions (Alvarez et al., 2024); the subscript distinguishes it from the scavenging coefficients λ of the wet-deposition subsection below.

930 For some applications TARSA also provides a simplified option in which w_{sed} is prescribed as a constant for each aerosol type using values drawn from the ECMWF aerosol specification (ECMWF, 2024b). The constants used in this option are summarised in Table E1.

E3 Dry deposition

Dry deposition acts as a first-order sink in the lowest model layer,

$$F_{\text{dry}} = -\frac{V_d}{\Delta z_1} C_{\text{ground}}, \quad (E8)$$



Table E1. Constant sedimentation velocities used in the simplified species-dependent option. Species abbreviations: OM, organic matter; BC, black carbon; NI1/NI2, nitrate bins 1–2; AM, ammonium; SOA, secondary organic aerosol.

Species	w_{sed} (m s ⁻¹)
Sea salt 1	2.4×10^{-5}
Dust 1	6.9×10^{-5}
Dust 2	1.982×10^{-4}
Dust 3	1.962×10^{-3}
OM/BC/NI1/AM/SOA	2.0×10^{-4}
Nitrate 2	1.33×10^{-3}

935 where V_d is the dry-deposition velocity, Δz_1 is the thickness of the lowest model layer, and C_{ground} is the tracer mixing ratio in the lowest model level; the division by Δz_1 converts the surface deposition flux into the mixing-ratio tendency required by Eq. (1). Following Zhang et al. (2001); Zhang and He (2014), TARSA uses a size- and surface-dependent parameterisation in which

$$V_d = \begin{cases} a_1 u_* & \text{for } D_p \leq 2.5 \mu\text{m}, \\ S_v (b_1 u_* + b_2 u_*^2 + b_3 u_*^3) \exp \left(K_1 \left(\frac{LAI}{LAI_{\text{max}}} - 1 \right) \right) R & \text{for } 2.5 \mu\text{m} < D_p \leq 10 \mu\text{m}, \\ S_v (d_1 u_* + d_2 u_*^2 + d_3 u_*^3) \exp \left(K_2 \left(\frac{LAI}{LAI_{\text{max}}} - 1 \right) \right) R & \text{for } D_p > 10 \mu\text{m}, \end{cases} \quad (\text{E9})$$

940 with

$$K_1 = c_1 u_* + c_2 u_*^2 + c_3 u_*^3, \quad (\text{E10})$$

$$K_2 = e_1 u_* + e_2 u_*^2 + e_3 u_*^3. \quad (\text{E11})$$

Here u_* is the friction velocity, LAI is the leaf area index, and R is a rebound-reduction factor for super-coarse particles:

$$R = \exp(-R_0 \sqrt{S_t}), \quad (\text{E12})$$

945 where the Stokes number is parameterised as

$$S_t = \frac{w_{\text{sed}} u_*}{g A_{\text{co}}} \quad \text{over vegetated surfaces}, \quad (\text{E13})$$

$$S_t = \frac{w_{\text{sed}} u_*^2}{g \nu} \quad \text{over smooth surfaces}, \quad (\text{E14})$$

with $\nu = \mu/\rho$ the kinematic viscosity of air and A_{co} the characteristic collector radius.

The expressions above parameterise the surface deposition velocity V_{ds} of Zhang and He (2014). Two simplifications relative to their bulk formulation are deliberate. First, the aerodynamic-resistance term of their bulk velocity, $V_d = V_g + (R_a + V_{ds}^{-1})^{-1}$, is not applied, so the V_d used here corresponds to their surface component alone and overestimates the bulk value where



aerodynamic resistance is appreciable. Second, the gravitational-settling contribution V_g is excluded from V_d because settling is treated separately as an explicit sedimentation flux (previous subsection), which avoids double-counting. In addition, the coarse and giant polynomials can return negative values at small friction velocities (below about 0.13 m s^{-1} for the two presets used in this paper); the implementation floors the parameterised velocity at zero in all branches.

The fine-particle branch $a_1 u_*$ gives V_d directly in m s^{-1} . For the coarse and giant branches, $S_v = 10^{-2}$ converts the cm s^{-1} output of the Zhang and He (2014) polynomial regressions (with u_* in m s^{-1}) to m s^{-1} . The numerical constants used in the dry-deposition configurations reported in this paper, including their coefficient units, are given in Table E2; the remaining implementation constants (S_v , R_0 , and the rebound-activation threshold $u_{*,R}$) are collected in Table E3. The Fennec dust simulations use the desert preset over land and the water preset over ocean, selected from the ERA5 land-sea mask. Both are smooth-surface presets, so LAI , LAI_{max} , and A_{co} do not enter those reported calculations.

Table E2. Dry-deposition coefficients used in the reported Fennec dust simulations; all values are taken from Zhang and He (2014). Coefficients are in the Zhang and He convention described in the text: the fine-particle branch gives velocity in m s^{-1} , while the coarse and giant branches require the conversion factor $S_v = 10^{-2}$. The Zhang and He PM_{10+} LAI-adjustment coefficients are denoted e_i here, corresponding to f_i in their notation.

Surface preset	Parameters	Values	Units
Water	a_1	6.9×10^{-3}	1
Water	b_1, b_2, b_3	$2.6 \times 10^{-1}, -1.3 \times 10^0, 3.0 \times 10^0$	$\text{cm m}^{-1}; \text{cm s m}^{-2}; \text{cm s}^2 \text{ m}^{-3}$
Water	c_1, c_2, c_3	0, 0, 0	$\text{s m}^{-1}; \text{s}^2 \text{ m}^{-2}; \text{s}^3 \text{ m}^{-3}$
Water	d_1, d_2, d_3	$-8.7 \times 10^{-1}, -5.5 \times 10^0, 9.9 \times 10^1$	$\text{cm m}^{-1}; \text{cm s m}^{-2}; \text{cm s}^2 \text{ m}^{-3}$
Water	e_1, e_2, e_3	0, 0, 0	$\text{s m}^{-1}; \text{s}^2 \text{ m}^{-2}; \text{s}^3 \text{ m}^{-3}$
Desert	a_1	4.3×10^{-3}	1
Desert	b_1, b_2, b_3	$3.9 \times 10^{-1}, -3.3 \times 10^0, 8.8 \times 10^0$	$\text{cm m}^{-1}; \text{cm s m}^{-2}; \text{cm s}^2 \text{ m}^{-3}$
Desert	c_1, c_2, c_3	0, 0, 0	$\text{s m}^{-1}; \text{s}^2 \text{ m}^{-2}; \text{s}^3 \text{ m}^{-3}$
Desert	d_1, d_2, d_3	$-7.3 \times 10^0, 4.6 \times 10^1, 9.4 \times 10^1$	$\text{cm m}^{-1}; \text{cm s m}^{-2}; \text{cm s}^2 \text{ m}^{-3}$
Desert	e_1, e_2, e_3	0, 0, 0	$\text{s m}^{-1}; \text{s}^2 \text{ m}^{-2}; \text{s}^3 \text{ m}^{-3}$

Table E3. Additional constants in the dry-deposition implementation. Rebound is disabled over water in the reported configurations and is enabled over the desert preset only for $u_* \geq u_{*,R}$.

Parameter	Value	Units	Source
S_v	10^{-2}	m cm^{-1}	Coarse/giant velocity conversion (Zhang and He, 2014)
R_0	0.1	1	(ECMWF, 2024b; Bergametti et al., 2018; Zhang et al., 2001)
$u_{*,R}$	0.335	m s^{-1}	(ECMWF, 2024b; Bergametti et al., 2018)

In the configurations reported here, dry deposition is applied only to aerosol cases for which the option is enabled. It is active in the Fennec dust-column evaluation and in the associated process-sensitivity diagnostic, where the five mineral-dust bins use



the size-dependent form above. It is disabled for the inert-tracer tests and for the supplementary CAMS replay experiment
965 reported in Appendix C, where lower-boundary concentrations are prescribed directly.

E4 Wet deposition

Wet removal is parameterised as a first-order decay,

$$\frac{dC}{dt} = -\Lambda C, \quad \Lambda = \lambda_{\text{in}} + \lambda_{\text{bc}}, \quad (\text{E15})$$

where λ_{in} and λ_{bc} are the in-cloud and below-cloud scavenging coefficients.

970 E4.1 Scavenging-regime determination.

Following Stohl et al. (2005) and Grythe et al. (2017), TARSA diagnoses an effective sub-grid cloud water quantity

$$PCW = \text{CTWC} \frac{F}{cc}, \quad (\text{E16})$$

$$F = \max \left[0.05, cc \cdot \frac{I_l f_{rl}(I_l) + I_c f_{rc}(I_c)}{I_l + I_c} \right], \quad (\text{E17})$$

where cc is cloud cover, I_l and I_c are large-scale and convective precipitation rates, CTWC is the column-integrated cloud
975 water content, and the precipitation-rate-dependent area-correction factors f_{rl} and f_{rc} are listed in Table E4. The corresponding sub-grid precipitation rate is

$$I_s = \frac{I_l + I_c}{F}. \quad (\text{E18})$$

At each horizontal grid point, layers with $PCW > 0$ are treated as in-cloud layers. Layers below the lowest cloudy layer but still under precipitating conditions are treated as below-cloud scavenging layers. In all remaining layers, $\Lambda = 0$.

Table E4. Correction factors used for the precipitating-area fraction, following Grythe et al. (2017). Precipitation rates are in mm h^{-1} .

Factor	$I_s \leq 1$	$1 < I_s \leq 3$	$3 < I_s \leq 8$	$8 < I_s \leq 20$	$20 < I_s$
f_{rl}	0.50	0.65	0.80	0.90	0.95
f_{rc}	0.40	0.55	0.70	0.80	0.90

980 E4.2 In-cloud scavenging.

The in-cloud scavenging coefficient follows Grythe et al. (2017):

$$\lambda_{\text{in}} = F_{\text{nuc}} \frac{I_s \rho_{\text{water}}}{PCW} r_{\text{icl}}, \quad (\text{E19})$$

$$F_{\text{nuc}} = \alpha(T) CCN_{\text{eff}} + (1 - \alpha(T)) IN_{\text{eff}}, \quad (\text{E20})$$



where I_s is here expressed in ms^{-1} and PCW (in kg m^{-2}) is the sub-grid precipitating cloud water defined above, so that
 985 $I_s \rho_{\text{water}} / PCW$ is the cloud-water washout rate of the precipitating cloud fraction, where

$$\alpha(T) = \begin{cases} 0, & T \leq 253.15 \text{ K}, \\ \left(\frac{T - 253.15}{20.00} \right)^2, & 253.15 < T < 273.15 \text{ K}, \\ 1, & T \geq 273.15 \text{ K}. \end{cases} \quad (\text{E21})$$

Here $r_{\text{icl}} = 0.0062 \text{ m}^3 \text{ kg}^{-1}$ stores the dimensionless cloud-water replenishment constant of Grythe et al. (2017) divided by the density of water: $\rho_{\text{water}} r_{\text{icl}} = 6.2$, close to the value $\text{ic}_r = 6.1$ that Grythe et al. (2017) obtained by empirical tuning against ECMWF cloud fields. CCN_{eff} and IN_{eff} are aerosol-type-dependent cloud-condensation-nuclei (CCN) and ice-nuclei (IN) activation efficiencies, prescribed per aerosol type with the values recommended by Grythe et al. (2017): mineral dust, which
 990 is an inefficient CCN, uses $CCN_{\text{eff}} = 0.15$ and $IN_{\text{eff}} = 0.02$ (the values validated for dust by Groot Zwaafink et al., 2016); black carbon uses $CCN_{\text{eff}} = 0.9$ and $IN_{\text{eff}} = 0.1$; and highly soluble aerosol such as sulfate uses $CCN_{\text{eff}} = IN_{\text{eff}} = 0.9$. Aerosol types without an explicit specification fall back to the generic implementation default $CCN_{\text{eff}} = 0.9$, $IN_{\text{eff}} = 0.1$. The five Fennec dust bins (Sect. 5.7) use the mineral-dust values, so that in-cloud nucleation scavenging of dust is weak and
 995 its wet removal is dominated by below-cloud scavenging.

E4.3 Below-cloud scavenging.

The default below-cloud scavenging option used in the present TARSA configuration is the precipitation-intensity power-law scheme. For cloud-free layers below the lowest precipitating-cloud layer, the below-cloud coefficient is

$$\lambda_{\text{bc}} = \begin{cases} C_{\text{rain}} \alpha_{\text{rain}} I_s^{\beta_{\text{rain}}}, & T \geq T_{\text{rain}} \text{ and } I_s > I_{\text{min}}, \\ C_{\text{ice}} \alpha_{\text{snow}} I_s^{\beta_{\text{snow}}}, & T < T_{\text{rain}} \text{ and } I_s > I_{\text{min}}, \\ 0, & I_s \leq I_{\text{min}}, \end{cases} \quad (\text{E22})$$

1000 where I_s is the sub-grid precipitation intensity in mm h^{-1} , $T_{\text{rain}} = 273.15 \text{ K}$ separates rain from snow, and $I_{\text{min}} = 0.01 \text{ mm h}^{-1}$ suppresses removal under negligible precipitation. The aerosol-dependent factors C_{rain} and C_{ice} are multiplicative scaling factors and are set to unity in the default aerosol definitions. The default power-law coefficients are listed in Table E5. Legacy diameter-dependent Laakso–Kyrö and Wang et al. parameterisations are retained as optional below-cloud kernels in the code, but they are not the default scheme described here.

1005 E5 Aerosol types and bulk properties

Table E6 summarises the dry densities, characteristic modal radii, geometric standard deviations, and hygroscopicity parameters adopted from the ECMWF IFS aerosol specification (ECMWF, 2024b). These values are used to diagnose effective bulk properties for the aerosol tracers transported in TARSA.



Table E5. Default parameters used by the power-law below-cloud scavenging scheme. The precipitation intensity I_s is expressed in mm h^{-1} and the resulting λ_{bc} is in s^{-1} .

Hydrometeor regime	α	β	C	T criterion	$I_{\min}$
Rain	1.0×10^{-5}	0.8	$C_{\text{rain}} = 1$	$T \geq 273.15 \text{ K}$	0.01 mm h^{-1}
Snow	3.0×10^{-5}	0.8	$C_{\text{ice}} = 1$	$T < 273.15 \text{ K}$	0.01 mm h^{-1}

Table E6. Physical properties and hygroscopicity parameters of the aerosol species used in TARSA, summarised from the ECMWF IFS aerosol specification (ECMWF, 2024b).

Aerosol type	Density [kg m^{-3}]	r_{mod} [μm]	σ	κ
Sea salt (80% RH)	1183	0.1992, 1.992	1.9, 2.0	1.12
Dust	2610	0.05, 0.42, 0.79, 16.2	2.2, 1.18, 1.93, 1.53	0.00
Black carbon	1000	0.0118	2.0	0.00 / 0.05
Organic matter	1300	0.09	1.6	0.06 / 0.12
Sulfates	1760	0.11	1.6	0.60
Nitrate fine	1730	0.0355	2.0	0.67
Nitrate coarse	1400	0.199, 1.992	1.9, 2.0	0.67
Ammonium	1760	0.0355	2.0	0.60
SOA	1800	0.09	1.6	0.12
Passive tracer	0.0	0.0	–	0.0

Appendix F: ERA5 input fields required by TARSA

1010 For reproducibility it is useful to summarise the ERA5 fields required by the current TARSA configuration. The model uses hourly ERA5 pressure-level data for the three-dimensional dynamical and thermodynamic state, combined with a small set of single-level fields for surface and boundary-layer diagnostics; the operational rationale for preferring the pressure-level product over model-level ERA5-complete requests is given in Sect. 2.3.

1015 The 25 pressure levels requested from the CDS reanalysis-era5-pressure-levels dataset, and used directly as the TARSA vertical layers in the native configurations (Sect. 3.1), are (in hPa): 1000, 975, 950, 925, 900, 875, 850, 825, 800, 775, 750, 700, 650, 600, 550, 500, 450, 400, 350, 300, 250, 200, 150, 100, 50. The ETEX refinement is described in Sect. 5.6.

F1 ERA5 variables for advection–diffusion processes

Table F1 lists the ERA5 fields consumed by the advection–diffusion core and the vertical-mixing diagnosis.

1020 These fields are used to reconstruct the three-dimensional wind field in geometric-height coordinates, convert archived geopotential to level height, compute air density and virtual temperature, and diagnose the vertical turbulent diffusivity K_z .



Table F1. ERA5 variables required for the advection–diffusion core and vertical mixing.

Symbol	Description	ERA5 name	Level type	Units
u_x	Eastward wind	u	Pressure	m s^{-1}
u_y	Northward wind	v	Pressure	m s^{-1}
ω	Vertical velocity	w	Pressure	Pa s^{-1}
T	Air temperature	t	Pressure	K
q	Specific humidity	q	Pressure	kg kg^{-1}
Φ	Geopotential	z	Pressure	$\text{m}^2 \text{s}^{-2}$
p_s	Surface pressure	sp	Single	Pa
PBLH	Boundary-layer height	blh	Single	m

F2 ERA5 variables for the dry-deposition model

Table F2 lists the additional single-level fields required when the dry-deposition option is active.

Table F2. ERA5 variables required for dry-deposition parameterisation.

Variable	Name	ERA5 name	Level type	Units
u_*	Friction velocity	friction_velocity	Single	m s^{-1}
LAI	Leaf area index (high vegetation)	leaf_area_index_high_vegetation	Single	$\text{m}^2 \text{m}^{-2}$
LSM	Land-sea mask	land_sea_mask	Single	0–1

These variables enter the empirical dry-deposition velocity formulation via surface roughness and vegetation effects.

F3 ERA5 variables for the wet-deposition parameterisation

1025 Table F3 lists the cloud and precipitation fields required when the wet-deposition option is active.

Table F3. ERA5 variables required for wet-deposition parameterisation.

Variable	Name	ERA5 name	Level type	Units
cc	Fraction of cloud cover	fraction_of_cloud_cover	Pressure	0–1
$CLWC$	Cloud liquid water content	specific_cloud_liquid_water_content	Pressure	kg kg^{-1}
$CIWC$	Cloud ice water content	specific_cloud_ice_water_content	Pressure	kg kg^{-1}
I_l	Large-scale precipitation	large_scale_precipitation	Single	m h^{-1}
I_c	Convective precipitation	convective_precipitation	Single	m h^{-1}

From these inputs TARSA derives the precipitating-area fraction, effective precipitation intensity, and total cloud water content used in the wet-scavenging coefficients.



Appendix G: Fully implicit first-order solver option

This appendix documents the fully implicit backward-Euler configuration of TARSA, which treats first-order upwind advection and vertical diffusion implicitly and yields a global sparse linear system per time step. This option is used for the MMS comparison (Sect. 6.1.4) and the supplementary CAMS replay (Appendix C).

G1 Fully implicit time discretisation

Let C_p^n denote the tracer mixing ratio in cell p at time t^n , and C_p^{n+1} the value at $t^{n+1} = t^n + \Delta t$. Using the face notation from Sect. 3.4, but orienting each face flux with the outward normal of cell p , define $Q_{p,f} = S_f \mathbf{u}_f \cdot \mathbf{n}_{p,f}$. After dividing the conservative species-mass balance by the cell density ρ_p , the fully implicit first-order upwind step is written as

$$V_p \frac{C_p^{n+1} - C_p^n}{\Delta t} + \sum_{f \in \partial p} \frac{\rho_{p,f}}{\rho_p} Q_{p,f} C_{p,f}^{\text{UP},n+1} = \frac{\rho_{\text{top}}}{\rho_p} K_{z,\text{top}} \left(\frac{\partial C^{n+1}}{\partial z} \right)_{\text{top}} S_{\text{top}} - \frac{\rho_{\text{bot}}}{\rho_p} K_{z,\text{bot}} \left(\frac{\partial C^{n+1}}{\partial z} \right)_{\text{bot}} S_{\text{bot}} + V_p E_p, \quad (\text{G1})$$

where the compact product $Q_{p,f} C_{p,f}^{\text{UP},n+1}$ denotes the corresponding first-order upwind face contribution from Appendix A1; in the differentiable implementation this product is evaluated through the smooth positive–negative flux split rather than by a hard sign branch. Here $\rho_{p,f}$ is the air density at face f , and all fluxes are evaluated at the new time level $n + 1$.

G2 Global sparse linear system

Collecting contributions from all cells yields a sparse linear system for the unknown state vector:

$$\hat{\mathbf{A}}_{n+1} \mathbf{X}^{n+1} = \mathbf{b}, \quad (\text{G2})$$

with

$$\hat{\mathbf{A}}_{n+1} = \mathbf{M} + \mathbf{A}_{\text{adv}} + \mathbf{A}_{\text{diff}} + \mathbf{A}_{\text{sink}}, \quad \mathbf{M} = \text{diag} \left(\frac{V_p}{\Delta t} \right). \quad (\text{G3})$$

This mass matrix corresponds to the locally density-normalised mixing-ratio equation above; multiplying each row by ρ_p recovers the conservative species-mass form. Here $\mathbf{A}_{\text{adv}}$ is the first-order upwind advection operator assembled from the density-ratio-weighted face-flux coefficients, $\mathbf{A}_{\text{diff}}$ is the vertical diffusion operator, and $\mathbf{A}_{\text{sink}}$ collects linear sink terms such as deposition where applicable. The right-hand side is

$$b_p = \frac{V_p}{\Delta t} C_p^n + V_p E_p + \sum_{f \in \partial p \cap \partial \Omega} \Phi_{p,f}^{\text{BC}}, \quad (\text{G4})$$

where $\Phi_{p,f}^{\text{BC}}$ encodes prescribed boundary contributions.

Because the vertical diffusion operator couples only adjacent levels within each (i, j) column, it is block-tridiagonal in isolation. Horizontal first-order upwind advection introduces sparse off-block couplings between neighbouring columns, yielding a large but highly structured sparse matrix.



G3 Boundary-condition treatment

1055 Boundary conditions enter the system through the upwind face values $C_{p,f}^{\text{UP}}$, the boundary terms $\Phi_{p,f}^{\text{BC}}$, and dedicated Dirichlet rows:

- The bottom layer ($k = 1$) is handled as a strong Dirichlet row: the lowest-layer state is pinned to the prescribed time-varying lower-boundary field (in the CAMS replay experiment, the CAMS lower-boundary concentrations), and interior rows couple to it only through their own off-diagonal coefficients.
- 1060 – The top layer ($k = N_z$) is likewise a strong Dirichlet row, pinned to the prescribed top-layer value (zero in the configuration used here), consistent with the top treatment of the vertical-diffusion solve described in Sect. 3.6.
- Lateral inflow boundaries are treated using ghost-cell values $C_{\text{in}}(x, t)$, which impose Dirichlet inflow conditions whenever the face velocity points into the domain.
- Lateral outflow boundaries require no special treatment: the upwind state is simply the last interior value, corresponding
1065 to zero-gradient outflow.

For the forward configurations considered here, no additional boundary fluxes are prescribed, so beyond the Dirichlet rows and lateral ghost contributions only emissions enter the right-hand side. The resulting system can be written compactly as

$$\hat{\mathbf{A}}_{n+1} \mathbf{X}_{n+1} = \mathbf{M}(\mathbf{X}_n + \Delta t \mathbf{E}_n), \quad (\text{G5})$$

where $\mathbf{E}_n$ is the vector of cell-centred source terms. In practice the matrix is solved with a sparse iterative solver (restarted
1070 GMRES, restart length 50, at most 300 iterations, relative tolerance 10^{-6} , preconditioned with an incomplete-LU factorisation with drop tolerance 10^{-2} ; if the ILU factorisation fails the preconditioner falls back to a Jacobi diagonal, and if GMRES fails the solver falls back to a direct sparse solve). No explicit non-negativity clipping is applied within the time stepping in order to preserve linearity of the forward update.

Appendix H: Definitions of the evaluation metrics used in Sect. 6

1075 This appendix gives the exact definitions of the summary statistics reported for the Gaussian, ETEX-1, and CAPTEX-1 comparisons in Sect. 6, matching the corresponding Gaussian, ETEX-1, and CAPTEX-1 validation scripts distributed with the model source (see the code and data availability section). Unless stated otherwise, M_i and R_i denote the model and reference (observed) values of a paired sample of size N , following the model–reference notation introduced in Sect. 6, and an overbar denotes the sample mean, e.g. $\bar{M} = N^{-1} \sum_i M_i$.

1080 H1 Correlation, coefficient of determination, and root-mean-square error

The Pearson correlation coefficient r and the root-mean-square error, already introduced in Sect. 6, are

$$r = \text{corr}(M, R), \quad \text{RMSE} = \sqrt{(M - R)^2}. \quad (\text{H1})$$



Throughout this paper, R^2 denotes the squared Pearson correlation, $R^2 = r^2$.

H2 Fractional bias and normalised mean square error

1085 For a paired sample $\{M_i, R_i\}_{i=1}^N$, the fractional bias and normalised mean square error are

$$\text{FB} = \frac{2(\bar{M} - \bar{R})}{\bar{M} + \bar{R}}, \quad \text{NMSE} = \frac{\overline{(M - R)^2}}{\bar{M} \bar{R}}, \quad (\text{H2})$$

identically implemented for the Gaussian, ETEX-1, and CAPTEX-1 comparisons. FB is positive for model overprediction, is bounded in $[-2, 2]$ for positive means, and both statistics evaluate to an undefined (not-a-number) result whenever the corresponding denominator is zero or non-finite. The relative- L_2 column (REL_{L_2}) of Table 6 is computed by the same script

1090 as $\text{REL}_{L_2} = \|M - R\|_2 / \|R\|_2$ over the same evaluated points.

H3 Figure of merit (FM) and figure of merit in time (FMT)

Two related but distinct overlap statistics are used. The scalar figure of merit, FM, used for every observable row of Table 8 and for the Gaussian-plume scatter statistic quoted in Sect. 6, is evaluated once across the N paired station (or grid-point) values of a single derived quantity, e.g. station-wise dosage or maximum concentration,

$$1095 \quad \text{FM} = \frac{2 \sum_i \min(M_i, R_i)}{\sum_i M_i + \sum_i R_i} \in [0, 1], \quad (\text{H3})$$

with $\text{FM} = 1$ for perfect agreement. Equation (H3) is the Dice–Sørensen overlap form. It maps one-to-one onto the min/max (intersection-over-union) overlap form used in several ETEX evaluation studies (cf. Mosca et al., 1998), $\text{FM}_{\min/\max} = \sum_i \min(M_i, R_i) / \sum_i \max(M_i, R_i)$, through $\text{FM}_{\min/\max} = \text{FM} / (2 - \text{FM})$; the Dice form is always the larger of the two. Published FM values do not always state which form is used, which is one of the definitional differences flagged in the caption of Table 9. The figure of merit in time, FMT, instead measures the temporal overlap of the model and observed concentration time series at a single station s , expressed as a percentage,

$$1100 \quad \text{FMT}_s = 100 \frac{\sum_t \min(M_{s,t}^+, R_{s,t}^+)}{\sum_t \max(M_{s,t}^+, R_{s,t}^+)}, \quad (\text{H4})$$

shown in Fig. 10, where $X^+ = \max(X, 0)$ clips both series to non-negative values before forming the ratio and the sum runs over that station's 3-hourly intervals with finite paired observations. FMT_s is undefined when the clipped denominator is zero or non-finite.

1105

H4 Dosage, arrival time, duration, and peak time

For each ETEX-1 station s , let $\mathcal{T}_s \subseteq \{1, \dots, 30\}$ denote the retained set of quantified, quality-controlled intervals for which a paired model value is available, and let C_t denote either the modelled or observed 3-hour interval concentration on that set. The concentrations are compared against the common detection threshold $C_{\text{thr}} = 0.05 \text{ ng m}^{-3}$ (Sect. 5.6). The first official



1110 observation window is 23 October 15:00–18:00 UTC. Writing $t_c(t)$ for the calendar midpoint of interval t , $t_{\text{ref}} = 23$ October 16:00 UTC for the release start, and $\Delta t = 3$ h for the interval length,

$$\text{Dosage: } D = \sum_{t \in \mathcal{T}_s: C_t \geq C_{\text{thr}}} C_t \Delta t, \quad (\text{H5})$$

$$\text{Arrival time: } \tau_{\text{arr}} = t_c(t^*) - t_{\text{ref}}, \quad t^* = \min\{t \in \mathcal{T}_s : C_t \geq C_{\text{thr}}\}, \quad (\text{H6})$$

$$\text{Duration: } T_{\text{dur}} = \Delta t \# \{t \in \mathcal{T}_s : C_t \geq C_{\text{thr}}\}, \quad (\text{H7})$$

$$1115 \quad \text{Peak time: } \tau_{\text{peak}} = t_c(t_{\text{max}}) - t_{\text{ref}}, \quad t_{\text{max}} = \min\left\{t \in \mathcal{T}_s : C_t = \max_{u \in \mathcal{T}_s} C_u\right\}, \quad (\text{H8})$$

each evaluated independently for the model and for the observations on the same retained paired set. Arrival and peak time use the representative interval midpoint because the three-hour means do not resolve the event time within a window; when maxima are tied, the earliest interval is used. Peak time is reported only when the corresponding series reaches C_{thr} , and the paired peak-time comparison requires hits in both series. Duration counts the retained intervals at or above threshold and
1120 multiplies by the fixed interval length; it is a cumulative time at or above threshold rather than the span between the first and last qualifying interval when qualifying intervals are separated by an intervening below-threshold interval. Station-selection and threshold conventions otherwise follow Sect. 5.6 and the caption of Table 8.

H5 CAPTEX-1 composite rank

For the supplementary CAPTEX-1 comparison, the composite 0–4 rank statistic used by Selvaratnam et al. (2023), quoted in
1125 Sect. 6.2 and shown in Fig. 11, is

$$\text{Rank} = R^2 + \text{clip}\left(1 - \frac{|\text{FB}|}{2}, 0, 1\right) + \text{clip}\left(\frac{\text{FMS}}{100}, 0, 1\right) + \text{clip}\left(1 - \frac{\text{KS}}{100}, 0, 1\right), \quad (\text{H9})$$

where $\text{clip}(x, 0, 1)$ clamps x to $[0, 1]$, and $R^2 = r^2$ and FB follow Eq. (H2), computed on the pointwise paired concentration sample ($N = 395$ valid station–interval pairs for the value quoted in Sect. 6.2). KS is the Kolmogorov–Smirnov statistic (in percent), the maximum absolute difference between the empirical cumulative distribution functions of the pooled (unpaired)
1130 model and observed concentration values. FMS is a plume-classification agreement statistic (in percent): at each distinct observation time, every reporting station is classified as plume-present or plume-absent by comparison with $C_{\text{thr}} = 0.05 \text{ ng m}^{-3}$ (applied as a background-clipping tolerance to both series ahead of a strict > 0 test), FMS for that time is the percentage of stations for which the model and observed classifications agree, and the reported FMS is the unweighted mean of this percentage over all observation times, i.e. each time step receives equal weight regardless of how many stations reported at that time.
1135 This percent-agreement convention differs from the union-based figure of merit in space of the classical evaluation protocols (Mosca et al., 1998), which counts only stations in the union of the predicted and observed plumes and does not credit stations where both series are below threshold; because both-absent stations dominate outside the plume, the percent-agreement value is systematically higher than the union-based one, and the rank in Eq. (H9) is correspondingly higher than a rank computed under the union convention. The implementation also reports two alternative FMS/Rank variants under different background-clipping
1140 conventions; the value quoted in the main text and shown in Fig. 11 uses the clipped ("non-trivial") convention described above.



Appendix I: Minimal quickstart workflow

Listing 1 prepares or reuses ERA5 inputs for a small Etna domain on 1–2 April 2024, injects a point source, runs the explicit solver, and writes NetCDF output. The tutorial directory supplies cached inputs when available.

Listing 1. Minimal TARSA quickstart example illustrating the high-level workflow from ERA5 input preparation to saved NetCDF output.

```
1145 1: using TARSA
      2:
      3:
      4: # Optional: set Copernicus CDS credentials here instead of ~/.cdsapirc
      5: # ENV["CDSAPI_URL"] = "https://cds.climate.copernicus.eu/api"
1150 6: # ENV["CDSAPI_KEY"] = "<uid>:<api-token>"
      7:
      8: levels_file, surface_file, _, _ = TARSA.prepare_input_data(
      9:     34.0, 40.0, # lat_min, lat_max
     10:     12.0, 18.0, # lon_min, lon_max
1155 11:     2024,      # year
     12:     4,        # month
     13:     [1, 2];   # days
     14:     download_cams=false,
     15:     data_dir=joinpath(@__DIR__, "data"),
1160 16: )
     17:
     18: input = TARSA.load_input_data(levels_file, surface_file)
     19:
     20: emissions = TARSA.build_point_source_emissions(
1165 21:     input,
     22:     15.0, # source_lon
     23:     37.75; # source_lat
     24:     rate_kg_s=8.0, injection_height=2500.0, sigma_z=600.0,
     25: )
1170 26:
     27: mixing_ratio = TARSA.run!(TARSA.build_simulation(
     28:     input,
     29:     emissions;
     30:     solver=TARSA.ExplicitSolver(cfl=0.7, dt=3600.0),
1175 31:     lateral_bc=:neumann0,
     32:     bottom_bc=:neumann0,
     33: ))
     34:
     35: output_file = joinpath(@__DIR__, "out", "quickstart_etna.nc")
1180 36: rm(output_file; force=true)
     37: TARSA.save_results(
     38:     output_file,
     39:     input,
     40:     mixing_ratio,
1185 41:     emissions;
     42:     template_file=levels_file,
     43:     variables=["concentration"],
     44:     save_emissions=false,
1190 45: )
```



Author contributions. Konstantin Kuznetsov designed the model formulation, implemented the TARSA code, performed the experiments, analysed the results, and wrote the manuscript. Oleg Dubovik contributed to the conceptual framing of the study and to scientific discussion of the model framework. Anton Lopatin contributed to the evaluation of the model against GRASP retrievals. David Fuertes and Elena Lind contributed to scientific discussion and manuscript review.

1195 *Competing interests.* The authors declare that they have no competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

1200 *Acknowledgements.* The authors thank Pavel Litvinov (CSO, GRASP Earth) and Abhina Kumar Behera (Research Scientist, GRASP Earth) for their valuable feedback. The authors are grateful to Mian Chin (NASA Goddard Space Flight Center) for advice on atmospheric transport and aerosol processes and for constructive feedback that improved the model description and validation methodology. The authors also thank Juan Carlos Antuña-Sánchez and Juan Gómez López (GRASP Earth) for their support with computational resources and engineering infrastructure.

Financial support. This work was supported by the European Space Agency (ESA) under ESA Contract No. 4000145377/24/NL/GM/jxh for the BASS project, and by the Centre National d'Études Spatiales (CNES) under Contract No. 220692/00 for the DECOGAP project.



1205 References

- Alexe, M. and Sandu, A.: An investigation of discrete adjoints for flux-limited numerical schemes, in: Proceedings of the 45th annual ACM Southeast Conference, pp. 373–378, 2007.
- Alvarez, C. A., Gunn, A., Swann, C., Trimble, S. M., Ewing, R. C., and Lapôte, M. G.: Direct measurements of dust settling velocity under low-density atmospheres using time-resolved particle image velocimetry, *Geophysical Research Letters*, 51, e2024GL109958, <https://doi.org/10.1029/2024GL109958>, 2024.
- 1210 Asch, M., Bocquet, M., and Nodet, M.: Data Assimilation: Methods, Algorithms, and Applications, Fundamentals of Algorithms, Society for Industrial and Applied Mathematics, Philadelphia, PA, <https://doi.org/10.1137/1.9781611974546>, 2016.
- Beljaars, A.: The impact of some aspects of the boundary layer scheme in the ECMWF model, in: Proc. Seminar on Parameterization of Sub-Grid Scale Physical Processes, 1995.
- 1215 Beljaars, A. and Viterbo, P.: The role of the boundary layer in a numerical weather prediction model. Clear and Cloudy Boundary Layers, AAM Holtslag and PG Duynkerke, Eds, 1998.
- Bergametti, G., Marticorena, B., Rajot, J.-L., Foret, G., Alfaro, S., and Laurent, B.: Size-resolved dry deposition velocities of dust particles: In situ measurements and parameterizations testing, *Journal of Geophysical Research: Atmospheres*, 123, 11–080, <https://doi.org/10.1029/2018JD028964>, 2018.
- 1220 Bey, I., Jacob, D. J., Yantosca, R. M., Logan, J. A., Field, B. D., Fiore, A. M., Li, Q., Liu, H. Y., Mickley, L. J., and Schultz, M. G.: Global modeling of tropospheric chemistry with assimilated meteorology: model description and evaluation, *Journal of Geophysical Research: Atmospheres*, 106, 23 073–23 095, <https://doi.org/10.1029/2001JD000807>, 2001.
- Botchorishvili, R., Elbern, H., and Janelidze, T.: Adjoint for advection schemes on the sphere in ICON model, arXiv preprint arXiv:1806.06117, <https://doi.org/10.48550/arXiv.1806.06117>, 2018.
- 1225 Brandt, J., Christensen, J. H., Frohn, L. M., and Zlatev, Z.: Numerical modelling of transport, dispersion, and deposition: Validation against ETEX-1, ETEX-2 and Chernobyl, *Environmental Modelling & Software*, 15, 521–531, [https://doi.org/10.1016/S1364-8152\(00\)00035-9](https://doi.org/10.1016/S1364-8152(00)00035-9), 2000.
- Carruthers, D. J., Holroyd, R. J., Hunt, J. C. R., Weng, W. S., Robins, A. G., Apsley, D. D., Thompson, D. J., and Smith, F. B.: UK-ADMS: A new approach to modelling dispersion in the earth's atmospheric boundary layer, *Journal of Wind Engineering and Industrial*
- 1230 *Aerodynamics*, 52, 139–153, [https://doi.org/10.1016/0167-6105\(94\)90044-2](https://doi.org/10.1016/0167-6105(94)90044-2), 1994.
- Chaboureaud, J.-P., Flamant, C., Dauhut, T., Kocha, C., Lafore, J.-P., Lavaysse, C., Marnas, F., Mokhtari, M., Pelon, J., Reinales Martínez, I., Schepanski, K., and Tulet, P.: Fennec dust forecast intercomparison over the Sahara in June 2011, *Atmospheric chemistry and physics*, 16, 6977–6995, <https://doi.org/10.5194/acp-16-6977-2016>, 2016.
- Cimorelli, A. J., Perry, S. G., Venkatram, A., Weil, J. C., Paine, R. J., Wilson, R. B., Lee, R. F., Peters, W. D., and Brode, R. W.: AERMOD: A dispersion model for industrial source applications. Part I: General model formulation and boundary layer characterization, *Journal of Applied Meteorology*, 44, 682–693, <https://doi.org/10.1175/JAM2227.1>, 2005.
- 1235 Dubovik, O. and King, M. D.: A flexible inversion algorithm for retrieval of aerosol optical properties from Sun and sky radiance measurements, *Journal of Geophysical Research: Atmospheres*, 105, 20 673–20 696, <https://doi.org/10.1029/2000JD900282>, 2000.
- Dubovik, O., Lapyonok, T., Kaufman, Y. J., Chin, M., Ginoux, P., Kahn, R. A., and Sinyuk, A.: Retrieving global aerosol sources from
- 1240 satellites using inverse modeling, *Atmospheric Chemistry and Physics*, 8, 209–250, <https://doi.org/10.5194/acp-8-209-2008>, 2008.



- Dubovik, O., Herman, M., Holdak, A., Lapyonok, T., Tanré, D., Deuzé, J., Ducos, F., Sinyuk, A., and Lopatin, A.: Statistically optimized inversion algorithm for enhanced retrieval of aerosol properties from spectral multi-angle polarimetric satellite observations, *Atmospheric Measurement Techniques*, 4, 975–1018, <https://doi.org/10.5194/amt-4-975-2011>, 2011.
- Dubovik, O., Lapyonok, T., Litvinov, P., Herman, M., Fuertes, D., Ducos, F., Torres, B., Derimian, Y., Huang, X., Lopatin, A.,
1245 Chaikovsky, A., Aspetsberger, M., and Federspiel, C.: GRASP: a versatile algorithm for characterizing the atmosphere, *SPIE Newsroom*,
<https://doi.org/10.1117/2.1201408.005558>, 2014.
- Dubovik, O., Fuertes, D., Litvinov, P., Lopatin, A., Lapyonok, T., Dubovik, I., Xu, F., Ducos, F., Chen, C., Torres, B., Derimian, Y., Li, L.,
Herreras-Giralda, M., Herrera, M., Karol, Y., Matar, C., Schuster, G. L., Espinosa, R., Puthukkudy, A., Li, Z., Fischer, J., Preusker, R.,
Cuesta, J., Kreuter, A., Cede, A., Aspetsberger, M., Marth, D., Bindreiter, L., Hangler, A., Lanzinger, V., Holter, C., and Federspiel, C.:
1250 A comprehensive description of multi-term LSM for applying multiple a priori constraints in problems of atmospheric remote sensing:
GRASP algorithm, concept, and applications, *Frontiers in Remote Sensing*, 2, 706 851, <https://doi.org/10.3389/frsen.2021.706851>, 2021.
- Dyer, A. J.: A review of flux-profile relationships, *Boundary-Layer Meteorology*, 7, 363–372, <https://doi.org/10.1007/BF00240838>, 1974.
- ECMWF: IFS Documentation CY49R1 - Part IV: Physical Processes, chap. 4, ECMWF, <https://doi.org/10.21957/c731ee1102>, 2024a.
- ECMWF: IFS Documentation CY49R1 - Part VIII: Atmospheric Composition, chap. 8, ECMWF, <https://doi.org/10.21957/d13af18259>,
1255 2024b.
- Ferber, G. J., Heffter, J. L., Draxler, R. R., Lagomarsino, R. J., Thomas, F. L., and Dietz, R. N.: Cross-Appalachian Tracer Experiment
(CAPTEX '83) Final Report, Tech. Rep. NOAA Tech. Memo. ERL ARL-142, NOAA Air Resources Laboratory, 1986.
- Flemming, J., Huijnen, V., Arteta, J., Bechtold, P., Beljaars, A., Blechschmidt, A.-M., Diamantakis, M., Engelen, R. J., Gaudel, A., Inness,
A., Jones, L., Josse, B., Katragkou, E., Marecal, V., Peuch, V.-H., Richter, A., Schultz, M. G., Stein, O., and Tsikerdekis, A.: Tropospheric
1260 chemistry in the Integrated Forecasting System of ECMWF, *Geoscientific Model Development*, 8, 975–1003, <https://doi.org/10.5194/gmd-8-975-2015>, 2015.
- Fuertes, D., Dubovik, O., Remer, L. A., Martins, J. V., Kleidman, R. G., and Federspiel, C.: GRASP-AirPhoton Multi-Angle Polarimeters
(GAPMAP): Calibrated high information content measurements for atmospheric and surface products from a commercial constellation
of CubeSats, in: 4S Symposium 2022, Vilamoura, Portugal, [https://airdrive.eventsair.com/eventsairwesteuprod/production-atpi-public/](https://airdrive.eventsair.com/eventsairwesteuprod/production-atpi-public/cc816beb2ac8457f959951ce63947ebf)
1265 [cc816beb2ac8457f959951ce63947ebf](https://airdrive.eventsair.com/eventsairwesteuprod/production-atpi-public/cc816beb2ac8457f959951ce63947ebf), 16–20 May 2022, 2022.
- Gelaro, R., McCarty, W., Suárez, M. J., Todling, R., Molod, A., Takacs, L., Randles, C. A., Darmenov, A., Bosilovich, M. G., Reichle,
R., Wargan, K., Coy, L., Cullather, R., Draper, C., Akella, S., Buchard, V., Conaty, A., da Silva, A. M., Gu, W., Kim, G.-K., Koster, R.,
Lucchesi, R., Merkova, D., Nielsen, J. E., Partyka, G., Pawson, S., Putman, W., Rienecker, M., Schubert, S. D., Sienkiewicz, M., and Zhao,
B.: The Modern-Era Retrospective Analysis for Research and Applications, Version 2 (MERRA-2), *Journal of Climate*, 30, 5419–5454,
1270 <https://doi.org/10.1175/JCLI-D-16-0758.1>, 2017.
- Giering, R. and Kaminski, T.: Recipes for adjoint code construction, *ACM Transactions on Mathematical Software*, 24, 437–474,
<https://doi.org/10.1145/293686.293695>, 1998.
- Grell, G. A., Peckham, S. E., Schmitz, R., McKeen, S. A., Frost, G., Skamarock, W. C., and Eder, B.: Fully coupled “online” chemistry
within the WRF model, *Atmospheric Environment*, 39, 6957–6975, <https://doi.org/10.1016/j.atmosenv.2005.04.027>, 2005.
- 1275 Griewank, A. and Walther, A.: *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, Society for Industrial and
Applied Mathematics, Philadelphia, PA, 2nd edn., <https://doi.org/10.1137/1.9780898717761>, 2008.
- Groot Zwaafink, C. D., Grythe, H., Skov, H., and Stohl, A.: Substantial contribution of northern high-latitude sources to mineral dust in the
Arctic, *Journal of Geophysical Research: Atmospheres*, 121, 13 678–13 697, <https://doi.org/10.1002/2016JD025482>, 2016.



- Grythe, H., Kristiansen, N. I., Groot Zwaafink, C. D., Eckhardt, S., Ström, J., Tunved, P., Krejci, R., and Stohl, A.: A new
1280 aerosol wet removal scheme for the Lagrangian particle model FLEXPART v10, *Geoscientific Model Development*, 10, 1447–1466,
<https://doi.org/10.5194/gmd-10-1447-2017>, 2017.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Sim-
mons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren,
P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R. J.,
1285 Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Vil-
laume, S., and Thépaut, J.-N.: The ERA5 global reanalysis, *Quarterly Journal of the Royal Meteorological Society*, 146, 1999–2049,
<https://doi.org/10.1002/qj.3803>, 2020.
- Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I., Schepers,
D., Simmons, A., Soci, C., Dee, D., and Thépaut, J.-N.: ERA5 hourly data on pressure levels from 1940 to present, Copernicus Climate
1290 Change Service (C3S) Climate Data Store (CDS), <https://doi.org/10.24381/cds.bd0915c6>, accessed 1 June 2026, 2023a.
- Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I., Schepers, D.,
Simmons, A., Soci, C., Dee, D., and Thépaut, J.-N.: ERA5 hourly data on single levels from 1940 to present, Copernicus Climate Change
Service (C3S) Climate Data Store (CDS), <https://doi.org/10.24381/cds.adbb2d47>, accessed 1 June 2026, 2023b.
- Huang, C., Sebastian, A. S., and Viswanathan, V.: Learning second-order TVD flux limiters using differentiable solvers, *arXiv preprint*
1295 *arXiv:2503.09625*, <https://doi.org/10.48550/arXiv.2503.09625>, 2025.
- Idelkadi, A., Hourdin, F., and Issartel, J. P.: Validation of LMD-ZT model with ETEX-1 experiment, in: *Air Pollution Modelling and Simu-
lation: Proceedings Second Conference on Air Pollution Modelling and Simulation, APMS'01 Champs-sur-Marne, April 9–12, 2001*, pp.
50–53, Springer, 2002.
- Inness, A., Ades, M., Agustí-Panareda, A., Barré, J., Benedictow, A., Blechschmidt, A.-M., Dominguez, J. J., Engelen, R., Eskes, H., Flem-
1300 ming, J., Huijnen, V., Jones, L., Kipling, Z., Massart, S., Parrington, M., Peuch, V.-H., Razinger, M., Remy, S., Schulz, M., and Suttie, M.:
The CAMS reanalysis of atmospheric composition, *Atmospheric Chemistry and Physics*, 19, 3515–3556, <https://doi.org/10.5194/acp-19-3515-2019>, 2019.
- Jöckel, P., von Kuhlmann, R., Lawrence, M. G., Steil, B., Brenninkmeijer, C. A. M., Crutzen, P. J., Rasch, P. J., and Eaton, B.: On a funda-
mental problem in implementing flux-form advection schemes for tracer transport in three-dimensional general circulation and chemistry
1305 transport models, *Quarterly Journal of the Royal Meteorological Society*, 127, 1035–1052, <https://doi.org/10.1002/qj.49712757318>, 2001.
- Kochkov, D., Yuval, J., Langmore, I., Norgaard, P., Smith, J., Mooers, G., Klöwer, M., Lottes, J., Rasp, S., Düben, P., Hatfield, S., Battaglia,
P., Sanchez-Gonzalez, A., Willson, M., Brenner, M. P., and Hoyer, S.: Neural general circulation models for weather and climate, *Nature*,
632, 1060–1066, <https://doi.org/10.1038/s41586-024-07744-y>, 2024.
- Koren, B.: A robust upwind discretization method for advection, diffusion and source terms, in: *Numerical Methods for Advection-Diffusion
1310 Problems*, edited by Vreugdenhil, C. B. and Koren, B., vol. 45 of *Notes on Numerical Fluid Mechanics*, pp. 117–138, Vieweg, Braun-
schweig/Wiesbaden, 1993.
- Kuznetsov, K.: TARSA v1.0: a lightweight Eulerian forward model for aerosol transport, Zenodo [software],
<https://doi.org/10.5281/zenodo.21932443>, 2026.
- Le Dimet, F.-X. and Talagrand, O.: Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects,
1315 *Tellus A*, 38, 97–110, <https://doi.org/10.3402/tellusa.v38i2.11706>, 1986.



- Li, L., Dubovik, O., Derimian, Y., Schuster, G. L., Lapyonok, T., Litvinov, P., Ducos, F., Fuertes, D., Chen, C., Li, Z., Lopatin, A., Torres, B., and Che, H.: Retrieval of aerosol components directly from satellite and ground-based measurements, *Atmospheric Chemistry and Physics*, 19, 13 409–13 443, <https://doi.org/10.5194/acp-19-13409-2019>, 2019.
- Li, L., Derimian, Y., Chen, C., Zhang, X., Che, H., Schuster, G. L., Fuertes, D., Litvinov, P., Lapyonok, T., Lopatin, A., Matar, C., Ducos, F., Karol, Y., Torres, B., Gui, K., Zheng, Y., Liang, Y., Lei, Y., Zhu, J., Zhang, L., Zhong, J., Zhang, X., and Dubovik, O.: Climatology of aerosol component concentrations derived from multi-angular polarimetric POLDER-3 observations using GRASP algorithm, *Earth System Science Data*, 14, 3439–3469, <https://doi.org/10.5194/essd-14-3439-2022>, 2022.
- Li, L., Che, H., Su, X., Zhang, X., Gui, K., Zheng, Y., Zhao, H., Zhao, H., Liang, Y., Lei, Y., Zhang, L., Zhong, J., Wang, Z., and Zhang, X.: Quantitative evaluation of dust and black carbon column concentration in the MERRA-2 reanalysis dataset using satellite-based component retrievals, *Remote Sensing*, 15, 388, <https://doi.org/10.3390/rs15020388>, 2023.
- Lin, J. C., Gerbig, C., Wofsy, S. C., Andrews, A. E., Daube, B. C., Davis, K. J., and Grainger, C. A.: A near-field tool for simulating the up-stream influence of atmospheric observations: the Stochastic Time-Inverted Lagrangian Transport (STILT) model, *Journal of Geophysical Research: Atmospheres*, 108, 4493, <https://doi.org/10.1029/2002JD003161>, 2003.
- Liu, Z. and Sandu, A.: On the properties of discrete adjoints of numerical methods for the advection equation, *International journal for numerical methods in fluids*, 56, 769–803, <https://doi.org/10.1002/fld.1547>, 2008.
- Louis, J. F.: A short history of the operational PBL parameterization at ECMWF, in: Workshop on Planetary Boundary Layer Parameterization, ECMWF, England, 1982, 1982.
- Mallet, V., Quélo, D., Sportisse, B., Ahmed de Biasi, M., Debry, É., Korsakissok, I., Wu, L., Roustan, Y., Sartelet, K., Tombette, M., and Foudhil, H.: Technical Note: The air quality modeling system Polyphemus, *Atmospheric Chemistry and Physics*, 7, 5479–5487, <https://doi.org/10.5194/acp-7-5479-2007>, 2007.
- Molina, M. O., Gutiérrez, C., and Sánchez, E.: Comparison of ERA5 surface wind speed climatologies over Europe with observations from the HadISD dataset, *International Journal of Climatology*, 41, 4864–4878, <https://doi.org/10.1002/joc.7103>, 2021.
- Mosca, S., Graziani, G., Klug, W., Bellasio, R., and Bianconi, R.: A statistical methodology for the evaluation of long-range dispersion models, *Atmospheric Environment*, 32, 4307–4324, [https://doi.org/10.1016/S1352-2310\(98\)00179-4](https://doi.org/10.1016/S1352-2310(98)00179-4), 1998.
- Nieuwstadt, F. T. M.: The turbulent structure of the stable, nocturnal boundary layer, *Journal of the Atmospheric Sciences*, 41, 2202–2216, [https://doi.org/10.1175/1520-0469\(1984\)041<2202:TTSOTS>2.0.CO;2](https://doi.org/10.1175/1520-0469(1984)041<2202:TTSOTS>2.0.CO;2), 1984.
- Nodop, K., Connolly, R., and Girardi, F.: The field campaigns of the European Tracer Experiment (ETEX), *Atmospheric Environment*, 32, 4095–4108, [https://doi.org/10.1016/s1352-2310\(98\)00190-3](https://doi.org/10.1016/s1352-2310(98)00190-3), 1998.
- Ou, Y., Li, L., Li, Z., Zhang, Y., Dubovik, O., Derimian, Y., Chen, C., Fuertes, D., Xie, Y., Lopatin, A., Ducos, F., and Peng, Z.: Spatio-Temporal Variability of Aerosol Components, Their Optical and Microphysical Properties over North China during Winter Haze in 2012, as Derived from POLDER/PARASOL Satellite Observations, *Remote Sensing*, 13, 2682, <https://doi.org/10.3390/rs13142682>, 2021.
- Petters, M. D. and Kreidenweis, S. M.: A single parameter representation of hygroscopic growth and cloud condensation nucleus activity, *Atmospheric Chemistry and Physics*, 7, 1961–1971, <https://doi.org/10.5194/acp-7-1961-2007>, 2007.
- Quélo, D., Krysta, M., Bocquet, M., Isnard, O., Minier, Y., and Sportisse, B.: Validation of the Polyphemus platform on the ETEX, Chernobyl and Algeciras cases, *Atmospheric environment*, 41, 5300–5315, <https://doi.org/10.1016/j.atmosenv.2007.02.035>, 2007.
- Randles, C. A., da Silva, A. M., Buchard, V., Colarco, P. R., Darmenov, A., Govindaraju, R., Smirnov, A., Holben, B., Ferrare, R., Hair, J., Shinozuka, Y., and Flynn, C. J.: The MERRA-2 aerosol reanalysis, 1980 onward. Part I: System description and data assimilation evaluation, *Journal of Climate*, 30, 6823–6850, <https://doi.org/10.1175/JCLI-D-16-0609.1>, 2017.



- Ryder, C. L., Highwood, E. J., Rosenberg, P. D., Trembath, J., Brooke, J. K., Bart, M., Dean, A., Crosier, J., Dorsey, J., Brindley, H., Banks, J., Marsham, J. H., McQuaid, J. B., Sodemann, H., and Washington, R.: Optical properties of Saharan dust aerosol and contribution from the coarse mode as measured during the Fennec 2011 aircraft campaign, *Atmospheric Chemistry and Physics*, 13, 303–325, <https://doi.org/10.5194/acp-13-303-2013>, 2013.
- Sandu, A., Daescu, D. N., Carmichael, G. R., and Chai, T.: Adjoint sensitivity analysis of regional air quality models, *Journal of Computational Physics*, 204, 222–252, <https://doi.org/10.1016/j.jcp.2004.10.011>, 2005.
- Scire, J. S., Strimaitis, D. G., and Yamartino, R. J.: A User's Guide for the CALPUFF Dispersion Model (Version 5), Tech. rep., Earth Tech, Inc., Concord, MA, https://www.calpuff.org/calpuff/download/CALPUFF_UsersGuide.pdf, 2000.
- Seibert, P. and Frank, A.: Source-receptor matrix calculation with a Lagrangian particle dispersion model in backward mode, *Atmospheric Chemistry and Physics*, 4, 51–63, <https://doi.org/10.5194/acp-4-51-2004>, 2004.
- Seinfeld, J. H. and Pandis, S. N.: *Atmospheric chemistry and physics: from air pollution to climate change*, John Wiley & Sons, 2016.
- Selvaratnam, V., Thomson, D. J., and Webster, H. N.: Validation of the atmospheric dispersion model NAME against long-range tracer release experiments, *Journal of Applied Meteorology and Climatology*, 62, 1165–1174, <https://doi.org/10.1175/JAMC-D-23-0021.1>, 2023.
- Sofiev, M., Vira, J., Kouznetsov, R., Prank, M., Soares, J., and Genikhovich, E.: Construction of the SILAM Eulerian atmospheric dispersion model based on the advection algorithm of Michael Galperin, *Geoscientific Model Development*, 8, 3497–3522, <https://doi.org/10.5194/gmd-8-3497-2015>, 2015.
- Sørensen, J. H., Baklanov, A., and Hoe, S.: The Danish emergency response model of the atmosphere (DERMA), *Journal of environmental radioactivity*, 96, 122–129, <https://doi.org/10.1016/j.jenvrad.2007.01.030>, 2007.
- Stein, A. F., Draxler, R. R., Rolph, G. D., Stunder, B. J. B., Cohen, M. D., and Ngan, F.: NOAA's HYSPLIT atmospheric transport and dispersion modeling system, *Bulletin of the American Meteorological Society*, 96, 2059–2077, <https://doi.org/10.1175/BAMS-D-14-00110.1>, 2015.
- Stohl, A., Forster, C., Frank, A., Seibert, P., and Wotawa, G.: The Lagrangian particle dispersion model FLEXPART version 6.2, *Atmospheric Chemistry and Physics*, 5, 2461–2474, <https://doi.org/10.5194/acp-5-2461-2005>, 2005.
- Sykes, R. I., Parker, S. F., Henn, D. S., Cerasoli, C. P., and Santos, L. P.: PC-SCIPUFF Version 1.2PD Technical Documentation, Tech. Rep. ARAP Report No. 718, Titan Corporation, Titan Research and Technology Division, ARAP Group, Princeton, NJ, 1998.
- Thomas, L. H.: *Elliptic Problems in Linear Difference Equations over a Network*, Tech. rep., Watson Scientific Computing Laboratory, Columbia University, New York, 1949.
- van Leer, B.: Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method, *Journal of Computational Physics*, 32, 101–136, [https://doi.org/10.1016/0021-9991\(79\)90145-1](https://doi.org/10.1016/0021-9991(79)90145-1), 1979.
- Viterbo, P., Beljaars, A., Mahfouf, J.-F., and Teixeira, J.: The representation of soil moisture freezing and its impact on the stable boundary layer, *Quarterly Journal of the Royal Meteorological Society*, 125, 2401–2426, <https://doi.org/10.1002/qj.49712555904>, 1999.
- Woodfield, J., Weller, H., and Cotter, C. J.: New limiter regions for multidimensional flows, *Journal of Computational Physics*, 515, 113–286, <https://doi.org/10.1016/j.jcp.2024.113286>, 2024.
- Zhang, L. and He, Z.: An empirical algorithm estimating dry deposition velocity of fine, coarse and giant particles, *Atmospheric Chemistry and Physics*, 14, 3729–3737, <https://doi.org/10.5194/acp-14-3729-2014>, 2014.
- Zhang, L., Gong, S., Padro, J., and Barrie, L.: A size-segregated particle dry deposition scheme for an atmospheric aerosol module, *Atmospheric environment*, 35, 549–560, [https://doi.org/10.1016/S1352-2310\(00\)00326-5](https://doi.org/10.1016/S1352-2310(00)00326-5), 2001.



Zhang, X., Li, L., Chen, C., Chen, X., Dubovik, O., Derimian, Y., Gui, K., Zheng, Y., Zhao, H., Zhang, L., Guo, B., Wang, Y., Holben, B., Che, H., and Zhang, X.: Validation of the aerosol optical property products derived by the GRASP/Component approach from multi-angular polarimetric observations, *Atmospheric Research*, 263, 105 802, <https://doi.org/10.1016/j.atmosres.2021.105802>, 2021.