



An N -dimensional Fortran interpolation programme (NterGeo.v2020a) for geophysics sciences – application to a back-trajectory programme (Backplumes.v2020r1) using CHIMERE or WRF outputs

Bertrand Bessagnet^{1,2}, Laurent Menut¹, and Maxime Beauchamp³

¹LMD/IPSL, École Polytechnique, Institut Polytechnique de Paris, ENS, PSL Université, Sorbonne Université, CNRS, 91128 Palaiseau, France

²Citepa, Technical Reference Center for Air Pollution and Climate Change, 42, rue de Paradis 75010 Paris, France

³IMT Atlantique, Lab-STICC UMR CNRS, 655 Avenue du Technopôle, 29280 Plouzané, France

Correspondence: Bertrand Bessagnet (bertrand.bessagnet@lmd.polytechnique.fr)

Received: 31 March 2020 – Discussion started: 5 May 2020

Revised: 24 July 2020 – Accepted: 5 November 2020 – Published: 7 January 2021

Abstract. An interpolation programme coded in Fortran for irregular N -dimensional cases is presented and freely available. The need for interpolation procedures over irregular meshes or matrixes with interdependent input data dimensions is frequent in geophysical models. Also, these models often embed look-up tables of physics or chemistry modules. Fortran is a fast and powerful language and is highly portable. It is easy to interface models written in Fortran with each other. Our programme does not need any libraries; it is written in standard Fortran and tested with two usual compilers. The programme is fast and competitive compared to current Python libraries. A normalization option parameter is provided when considering different types of units on each dimension. Some tests and examples are provided and available in the code package. Moreover, a geophysical application embedding this interpolation programme is provided and discussed; it consists in determining back trajectories using chemistry-transport or mesoscale meteorological model outputs, respectively, from the widely used CHIMERE and Weather Research and Forecasting (WRF) models.

against ground station observations. The NetCDF Operators (NCO) library (Zender, 2008) is commonly used in its recent version (v4.9.2) for horizontal and vertical interpolations to manage climate model outputs. The most frequent need is to interpolate in 3-D spatial dimension and time, and therefore in four dimensions. Fortran is extensively used for atmosphere modelling software (Sun and Grimmond, 2019; e.g. the Weather Research and Forecasting model – WRF, Skamarock et al., 2008; the Geophysical Fluid Dynamics Laboratory atmospheric component version 3 – GFDL AM3, Donner et al., 2011). More generally, geophysical models can use look-up tables of complex modules instead of a full coupling strategy between these modules, which is the case of the CHIMERE model (Mailler et al., 2017) with the embedded ISORROPIA module dealing with chemistry and thermodynamics (Nenes et al., 1998, 1999). In such a case, the look-up table can easily exceed five dimensions to approximate the model. In parallel, artificial intelligence methods are developed and can explore the behaviour of complex model outputs that requires fast interpolation methods. While more recent modern languages like Python are used in the scientific community, Fortran remains widely used in the geophysics and engineering community and is known as one of the faster languages in time execution, performing well on array handling, parallelization and, above all, portability. Some benchmarks are available on website to evaluate the

1 Introduction

Interpolation is commonly used in geophysical sciences for post-treatment processing to evaluate model performance

performance of languages on simple to complex operations (Kouatchou, 2018).

The parameterization techniques proposed to manage aerosol–droplet microphysical schemes (Rap et al., 2009) can employ either the modified Shepard interpolation method (Shepard, 1968) or the Hardy multiquadric interpolation method (Hardy, 1971, 1990), and the numerical results obtained show that both methods provide realistic results for a wide range of aerosol mass loadings. For the climate community, a comparison of six methods for the interpolation of daily European climate data is proposed by Hofstra et al. (2008); some of these methods use kriging-like methods with the capability to use co-predictors like the topography.

A Python procedure called *scipy.interpolate.griddata* is freely available (Scipy, 2014). Unfortunately, this programme is not really adapted to our problem; it could be not enough optimized for our objective as it can manage fully unstructured datasets. The goal of this paper is to present a programme to interpolate in a grid or a matrix which can be irregular (varying intervals) but structured, with the possibility to have interdependent dimensions (e.g. longitude interval edges which depend on longitude, latitude, altitude and time). We think this type of programme can be easily implemented within models or to manage model outputs for post-treatment issues. In short, the novelty of this programme is to fill the gap of interpolation issues between the treatment of very complex unstructured meshes and simple regular grids for a general dimension N .

In order to quantify the impact of such a new interpolation programme and show examples of its use, it is implemented in the Backplumes back-trajectory model, developed by the same team as the CHIMERE model (Mailler et al., 2017). This host model is well fit for this implementation, because the most important part of its calculation is an interpolation of a point in a model grid box. This paper describes (i) the methodology and the content of the interpolation programme package *NterGeo* and (ii) an application of this programme embedded in the new back-trajectory programme, *Backplumes*. These two codes are freely available (see code availability section).

2 Development of the interpolation programme

The *NterGeo* programme is fit for exploring irregular but structured grids or look-up tables defined by a unique size for each dimension, which of course can be different from one to another dimension. The space intervals can vary along a dimension and the grid interval edges in each dimension can depend on other dimensions. Two versions have been developed: (i) a version for “regular” arrays with independent dimensions and (ii) a “general” version for possible interdependent dimensions, e.g. to handle 3-D meshes which have time-varying spatial coordinates. The code does not need any libraries and is written in standard Fortran. Our interpolation




Figure 1. Description of variables for $N = 3$ with a regular grid case.

code was tested with *gfortran* (GNU Fortran project) and *ifort* (Intel). Since our programme does not include specific options and is not function compiler dependent, there is no reason to have limitations or errors with other compilers. The top shell calling script in the package provides two sets of options for “production” and “debugging” modes. Assuming the X array, the result of the function f transforming X to Y array in \mathbb{R} can be expressed as

$$Y(x_1, \dots, x_N) = f(X(x_1, \dots, x_N)). \quad (1)$$

N is the dimension of the array, and x_i is the coordinates at dimension $i \in [1, N]$ of the point X that we want to interpolate.

2.1 The programme for regular grids

A programme (*interpolation_regular.F90*) for regular grids (i.e. with independent dimensions) is available. To handle this type of grid, a classical multilinear interpolation is performed. Figure 1 shows the variables for $N = 3$ defined hereafter in the section.

For the particular case of a regular grid with independent dimensions, the result \tilde{Y} of the multilinear interpolation of the 2^N identified neighbours can be expressed as

$$\begin{aligned} \tilde{Y} = & w_N^0 \dots w_i^0 \dots w_1^0 \times Y_0(0 \dots 0 \dots 0) \\ & + w_N^0 \dots w_i^0 \dots w_1^1 \times Y_1(0 \dots 0 \dots 1) \\ & + \dots \\ & + w_N^{\delta_N} \dots w_i^{\delta_i} \dots w_1^{\delta_1} \times Y_k(\delta_N \dots \delta_i \dots \delta_1) \\ & + \dots \\ & + w_N^1 \dots w_i^1 \dots w_1^1 \times Y_{2^N-1}(1 \dots 1 \dots 1), \end{aligned} \quad (2)$$

with δ_i the binary digit equal to 0 or 1, and the weights $w_i^{\delta_i}$ for $i \in [1, N]$ defined as

$$w_i^0 = \frac{\theta_i^1 - x_i}{\theta_i^1 - \theta_i^0} w_i^1 = 1 - w_i^1. \quad (3)$$

Variable Θ_i is the list of interval edges on each dimension i and does not depend on other dimensions. $\theta_i^{\delta_i}$ indicates the bottom ($\delta_i = 0$) and top ($\delta_i = 1$) edges on each dimension $i \in [1 \dots N]$ so that $x_i \in [\theta_i^0, \theta_i^1]$. Y_k is a one-dimensional array with 2^N elements storing the value Y of the function at the identified neighbours Ψ on each dimension:

$$Y_k(\delta_N \dots \delta_i \dots \delta_1) = f\left(\Psi(\theta_N^{\delta_N}, \dots, \theta_i^{\delta_i}, \dots, \theta_1^{\delta_1})\right), \quad (4)$$

with $k \in [0, 2^N - 1]$.

The tuple $(\delta_N \dots \delta_i \dots \delta_1)$ is the binary transformation of integer k defined as $\sum_{i=0}^{N-1} (\delta_i \times 2^i)$. The coefficients $\Gamma_k = w_N^{\delta_N} \dots w_i^{\delta_i} \dots w_1^{\delta_1}$ as a product of weighting factors on each direction can be seen as a binary suite that is convenient to handle in a compacted and optimized Fortran programming strategy for the regular grid version of the code (Appendix B).

2.2 The general programme

Considering the general programme called `interpolation_general.F90`, the coordinates of edge points are stored in a one-dimensional array of $n = \prod_{i=1}^N I_i$ elements with I_i the number of edges on each dimension i . The tuple of coordinates (j_1, \dots, j_N) of an interval edge θ_k^i , with j_i the indexed coordinate on dimension i , is transformed in a one-dimensional array indexed on $k \in [1, n]$ by

$$k = \sum_{j=1}^N \left((i_j - 1) \prod_{l=0}^{j-1} I_l \right) + 1, \quad (5)$$

with $I_0 = 1$ for initialization.

Once the nearest neighbour is found, the result \tilde{Y} of the interpolation is a weighting procedure of the 2^N closest vertex using a Shepard interpolation (Shepard, 1968) based on the inverse distance calculations:

$$\tilde{Y} = \sum_{k=0}^{2^N-1} (\Gamma_k \times Y_k), \quad (6)$$

with $Y_k = f(\Upsilon_k)$ the value of the function f at neighbour Υ_k of coordinates $(\theta_k^1, \dots, \theta_k^N)$ and

$$\Gamma_k = \frac{1/d_k}{\sum_{k=0}^{2^N-1} (1/d_k)}. \quad (7)$$

The distance d_k between the point of interest of coordinates (x_1, \dots, x_N) to the neighbour $k \in [1, n]$ is calculated as

$$d_k = \left(\sum_{i=1}^N |x_i - \theta_k^i|^p \right)^{\frac{1}{p}}. \quad (8)$$

The previous formulas are valid for $d_k \neq 0$; in the case of $d_k = 0$, the procedure stops and exits, returning the exact value of the corresponding data of the nearest neighbour. For a distorted mesh or matrix, or dimensions with different units (e.g. mixing time with length), a hard-coded option (*norm=.true.* or *.false.*) is also available to normalize the intervals with an average interval Δ_i value for the calculation of distances, so that

$$d_k = \left(\sum_{i=1}^N \left(\frac{|x_i - \theta_k^i|}{|\Delta_i|} \right)^p \right)^{\frac{1}{p}}. \quad (9)$$

3 Computation strategy for the general programme

The list of input/output arguments is provided in Appendix C. In the main programme, calling the subroutine the key point is to transform first the N -dimension matrix into a 1-D array. An example of a main programme calling the subroutine is provided in the code package. The computation strategy in the subroutine can be broken down into the sequential steps as follows:

- i. Find the nearest neighbour of the input data by minimizing a distance with a simple incremental method stepping every ± 1 coordinates on each dimension (detailed later in this section).
- ii. Scan the surroundings of the nearest point within the matrix on ± 1 step on each dimension and store the corresponding block of input data to be tested. The size of the block is therefore $(1 + 2 \times 1)^N$ but can be extended to $(1 + 2 \times 2)^N$ if we increase the scanning process to ± 2 on each dimension (hard-coded option *iconf* of 1 or 2 in the declaration block).
- iii. Calculate the distance to the previously selected input data. A p -distance concept is adopted (hard-coded option *pnum* in the declaration block). The *pnum* value p should be greater than or equal to 1 to verify the Minkowski inequality and be considered as a metric.
- iv. Sort the previous block of data in ascending order and stop the sorting process when the first 2^N point is selected. The code offers the possibility to use only the first $N + 1$ neighbour (hard-coded option *neighb* in the declaration block) that is sufficient and faster in most cases.
- v. Calculate the weights and then the final result.

The first step consisting in finding the first neighbour is the trickiest and is broken down into several steps. Figure 2 displays an example in 2-D of the step-by-step procedure to find the nearest neighbour.

- i. The procedure initializes the process starting from the first point of the input data grid or taken from the last closest point if given in an argument as a non-null value.
- ii. A delta of coordinates is applied based on an average delta on each dimension to improve the initialization. This computation step of delta is externalized as it can be time consuming and should be done once for all target points at which we want to interpolate.
- iii. A test between the target value and the input data grid point coordinates determines the ± 1 steps to add on each dimension (see Fig. 2 for an example in 2-D).
- iv. If the grid point falls on the edges or outside the borders, the closest coordinates within the matrix are selected.
- v. A test on the p -distance computation between the running point and the target is performed so that if the distance calculated at iteration N_{it} is equal to the distance at iteration $N_{it} - 2$ the closest point is found.
- vi. If the distance is larger than the characteristic distance of the cell, the point is considered to be outside the borders of the input data grid. Therefore, the code allows a slight extrapolation if the target point is not too far from the borders.
- vii. At this stage, the procedure can stop if the distance to the closest vertex is 0, returning to the main programme with the exact value of the input data grid.

4 Visual example in 2-D for a regular grid

As an example to visualize the capacity of the general programme, the 2-D function used in Scipy (2014) is used to test our procedure. The function is

$$Y = f(X) = x_1 \times (1 - x_1) \times \cos(4\pi x_1) \times \sin(4\pi x_2^2), \quad (10)$$

with $x_1, x_2 \in [0, 1]$.

Our input data grid is a regular grid with regular intervals of 0.02 from 0 to 1 for x_1 and x_2 with therefore 51 points on each dimension. We propose to interpolate on a finer regular grid with $n = 100 \times 100, 200 \times 200$ and 300×300 points on each dimension. For these three interpolation cases, a normalized mean square error (NMSE) of the result \tilde{Y}_j for the full grid point number j can be calculated against the true value Y_j of the function as

$$\text{NMSE} = \frac{\frac{1}{n} \sum_{j=1}^n (\tilde{Y}_j - Y_j)^2}{\frac{1}{n-1} \sum_{j=1}^n (Y_j - \bar{Y}_j)^2}, \quad (11)$$

with \bar{Y}_j the mean value Y_j as $\frac{1}{n} \sum_{j=1}^n Y_j$.

Table 1. Performance for each case with $p = 1$.

Case	100×100	200×200	300×300
NMSE (%)	0.324	0.319	0.319
CPU time (s)	0.45	1.84	4.1

Table 2. Performance for the 5-D ($N = 5$) case with $p = 2$.

Number of neighbours	2^N	$N + 1$
NMSE (%)	1.570	0.870
CPU time (s)	17.32	6.00

For the three cases, the CPU time for the interpolation is evaluated and displayed in Table 1 for Machine 1 (Appendix E). As expected, the time consumption is obviously proportional to the number of points in which to interpolate. Figure 3 displays the evolution of the NMSE with the parameter p of the p -distance definition. There is a discontinuity of the NMSE from $p = 1$ to $p = 1^+$ with a slight increase with p in an asymptotic way (Fig. 4). The NMSE decreases with the number of points but a slight increase is observed from 200×200 to 300×300 .

5 Example in 5-D for a regular grid

Still using the general programme, an example in 5-D ($N = 5$) is proposed using the function

$$Y = f(X) = x_1 \times (1 - x_1) \times \cos(4\pi x_1) \times \sin(4\pi x_2) \times \cos(4\pi x_3) \times \sin(4\pi x_4) \times \cos(4\pi x_5), \quad (12)$$

with $x_1, x_2, x_3, x_4, x_5 \in [0, 1]$. The input data grid is a regular grid of $I_i = 35$ interval edges on each dimension $i \in [1, 5]$ with $35^5 = 52\,521\,875$ grid points. The goal is to find the results on a coarse grid of nine elements on each dimension with $9^5 = 59\,049$ grid points. This case is an opportunity to test the influence of the number of neighbours in calculating the result. In our case, the parameter p of the p distance is set to $p = 2$. The interpolation seems to provide better performance of the NMSE for our function with less neighbours (case $N + 1$) and obviously with a lower CPU time (Table 2). This could certainly depend on the type of function to interpolate.

Another test with the 5-D case is performed to test the influence of the normalization as defined in Eq. (9) (flag *norm*) by defining an irregular grid still with $35^5 = 52\,521\,875$ input data points but with (i) random intervals values and (ii) one dimension depending on another. The definition of the input grid is defined in Appendix D and provided in the code package. With a similar order of magnitude of consumed CPU time, the normalization *norm=True* produces a NMSE of 0.499 % compared to the NMSE of 0.822 % for *norm=False*.

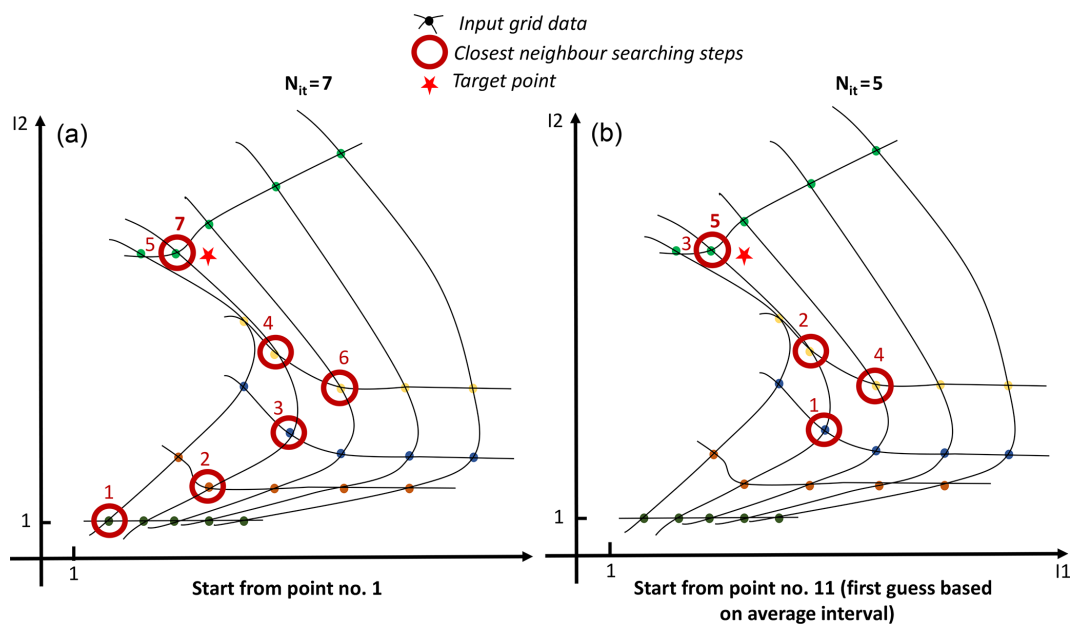


Figure 2. Real example in 2-D of the step-by-step procedure to find the nearest neighbour of a target point for an irregular but structured 5 × 5 grid (a) when starting the process from the first point of the grid on the lowest left corner and (b) when starting with a first guess based on an average delta computed for each dimension.

There is then an added value of using such a normalization with comparable CPU time consumption (rising from 2.68 to 3.44 s for our case).

6 Comparison with Python for a regular grid

The code has been tested against the Python procedure `scipy.interpolate.griddata`, freely available from Scipy (2014), for the following function:

$$Y = f(X) = x_1 \times (1 - x_1) \times \cos(4\pi x_1) \times \sin(4\pi x_2) \times \cos(4\pi x_3), \quad (13)$$

with $x_1, x_2, x_3 \in [0, 1]$.

The input data grid is a regular grid of $I_i = 35$ interval edges on each dimension $i \in [1, 5]$ with $35^3 = 42\,875$ grid points. The goal is to find the results on a coarse grid of nine elements on each dimension with $9^3 = 729$ grid points. A case in 3-D has been used for this test because the Python library was not able to work with very large datasets (*overflow error*), while our programme could work perfectly. Here, `scipy.interpolate.griddata` is used with the bilinear interpolation option, while our method is configured with $p = 2$.

Table 3 clearly shows how the Fortran code is faster compared to the Python library. However, the bilinear interpolation method seems to provide a higher accuracy than the inverse distance method embedded in our programme. Nevertheless, the error produced by our method looks acceptable.

Table 3. Comparison of performance between our code for a 3-D case with the grid data Python library. Machine 2 is used (Appendix F).

	Our code with $N + 1$ neighbours	Our code with 2^N neighbours	Python
NMSE (%)	0.627	1.03	0.326
CPU time (s)	0.04	0.04	19.49

7 Geophysics application

7.1 The Backplumes model

In order to test this new interpolation programme, it is implemented in a back-trajectory model called Backplumes. This model was already used in some studies such as Mailler et al. (2016) and Flamant et al. (2018). Backplumes is open source and is available on the CHIMERE website. Backplumes calculates back trajectories from a starting point and a starting date. It is different from other back-trajectory models, such as the Hybrid Single-Particle Lagrangian Integrated Trajectory (HYSPLIT) (Stein et al., 2015), Stochastic Time-Inverted Lagrangian Transport (STILT) (Lin et al., 2003; Nehr Korn et al., 2010) and FLEXible PARTicle dispersion model (FLEXPART) (Pisso et al., 2019), because it launches hundreds of particles and plots all trajectories as outputs. Thus, the answer is complementary compared to the other models: the output result is all possible trajectories and not only the most likely.




Figure 3. Interpolation results for the three cases. Figures were generated with the Generic Mapping Tools (Wessel et al., 2019).

An advantage of Backplumes for the WRF and CHIMERE users is that the code is dedicated to directly read output results of these models. Being developed by the CHIMERE developer teams, the code is completely homogeneous with CHIMERE in terms of numerical libraries. Another advantage is that the code is very fast and calculates hundreds of trajectories in a few minutes. Using the wind fields of WRF or CHIMERE, and running on the same grid, the results of back trajectories are fully consistent with the simulations done by the models.

Backplumes is dedicated to calculate transport but not chemistry: only passive air particles (or tracers) are released. But a distinction could be made between gaseous or particulate tracers: for the latter, a settling velocity is calculated to have a more realistic trajectory. The model is easy to use and light because a small set of meteorological parameters is required. These meteorological parameters are described in Table G1 for WRF and CHIMERE.

The first step of the calculation is to choose a target location as a starting point. The user must select a date, longitude, latitude and altitude, obviously included in the mod-

elled domain and during the modelled period. From this starting point, the model will calculate trajectories back in time. The number of trajectories is up to the user and may vary from one to several hundreds of tracers.

At each time step and for each trajectory, the position of the air mass is estimated by subtracting its pathway travelled as longitude $\Delta\lambda$, latitude $\Delta\phi$ and altitude Δz to the current position. To do so, all necessary variables are interpolated with the NterGeo.v2020a interpolation programme described in the previous section. The calculation is described in Appendix G.

In order to respect the Courant–Friedrichs–Lewy (CFL) number, a sub-time step may be calculated. If the input data are provided hourly (as in many regional models), the meteorological variables are interpolated between the two consecutive hours to obtain refined input data.

The goal of Backplumes is to estimate all possible back trajectories. Then, starting from one unique point, it is necessary to add a pseudo-turbulence in the calculation of the altitude. Depending on the vertical position of the tracer, several hypotheses are made. Two parameters are checked for




Figure 4. Evolution of performance based on the NMSE for the three cases as a function of the parameter p of the p -distance computation.

each tracer and each time step: (i) the boundary layer height enables us to know if the tracer is in the boundary layer or above in the free troposphere, (ii) the surface sensible heat fluxes enable to know if the atmosphere is stable or unstable.

When the tracer is diagnosed in the boundary layer, there are two cases: the boundary layer is stable or unstable. If the boundary layer is stable, $Q_0 < 0$, the tracer stays in the boundary layer at the same altitude. The new vertical position of the tracer is

$$z_{t-1} = z_t. \quad (14)$$

If the boundary layer is unstable, $Q_0 > 0$, the tracer is considered in the convective boundary layer and may be located at every level in this boundary layer for the time before this. Therefore, a random function is applied to reproduce a potential vertical mixing.

$$z_{t-1} = \text{Rand} \times \bar{h} \quad (15)$$

The random function “Rand” calculates a coefficient between 0 and 1 to represent stochastic vertical transport of the tracer.

It is considered that 15 mn is representative of a well-mixed convective layer (Stull, 1988). If the time step is larger than 15 mn, the random function is applied. But if the time step is less than 15 mn, the vertical mixing is reduced to the vicinity of the current position of the tracer. In this case, we have

$$z_{t-1} = \text{Rand} \times \Delta z \times [z_t], \quad (16)$$

where $\Delta z = \frac{1}{2} (z_t^{k-1} + z_t^{k+1})$ and k is the vertical model level corresponding to z_t .

In the free troposphere, the evolution of the tracer is considered to be influenced by the vertical wind component. A random function is applied to estimate its possible vertical motion with values between 0 and $w/2 \text{ m s}^{-1}$, representative of all possible values of vertical wind speed in the troposphere, Stull (1988). The vertical variability of the tracer’s position in the free troposphere is calculated by diagnosing the vertical velocity as

$$z_{t-1} = z_t - (0.5 + \text{Rand})w \frac{3600}{\Delta t}. \quad (17)$$

7.2 Examples of back-trajectory computations

An example is presented for the same case and the WRF and CHIMERE models. The difference between the two models is the number of vertical levels (35 for WRF and 20 for CHIMERE, from the surface to 200 hPa). The online modelling system WRF-CHIMERE is used, meaning that the horizontal grid is the same (a large domain including Europe and Africa and with $\Delta x = \Delta y = 60 \text{ km}$). The wind field is the same for both models: CHIMERE directly uses the wind field calculated by WRF. The boundary layer height is different between the two models, with WRF using the Hong et al. (2006) schemes and CHIMERE using the Troen and Mahrt (1986) scheme. The surface sensible heat flux is the same between the two models, with CHIMERE using the flux calculated by WRF. WRF has more vertical model levels than CHIMERE; thus, meteorological fields are interpolated from WRF to CHIMERE. It impacts the horizontal and vertical wind fields.

Figure 5 presents the results of back trajectories launched on 10 August 2013 at 12:00 UTC. The location is at longitude 10° E and latitude 25° N , with an altitude of 0 m a.g.l. This location is of no scientific interest but is in the middle of the domain, in order to have the longer trajectories. The complete duration of trajectories represents 10 d back in time. Overall, 120 trajectories are launched at the same position and time. They are randomly mixed when they are in the boundary layer to represent the mixing and the diffusion.

The most important part of the plume comes from the north of the starting point. For this main plume, the calculation is similar between the two models. Another large part of Backplumes is modelled at the east of the starting point. However, this fraction is mainly modelled with WRF but not with CHIMERE, where only a few trajectories are diagnosed. One possible explanation may be found by analysing the vertical transport of the trajectories.

Figure 6 presents all plumes displayed in the previous figure but projected along the same time–altitude axis. The differences between the two Backplumes results are mainly due to the calculation of the boundary layer height. When WRF

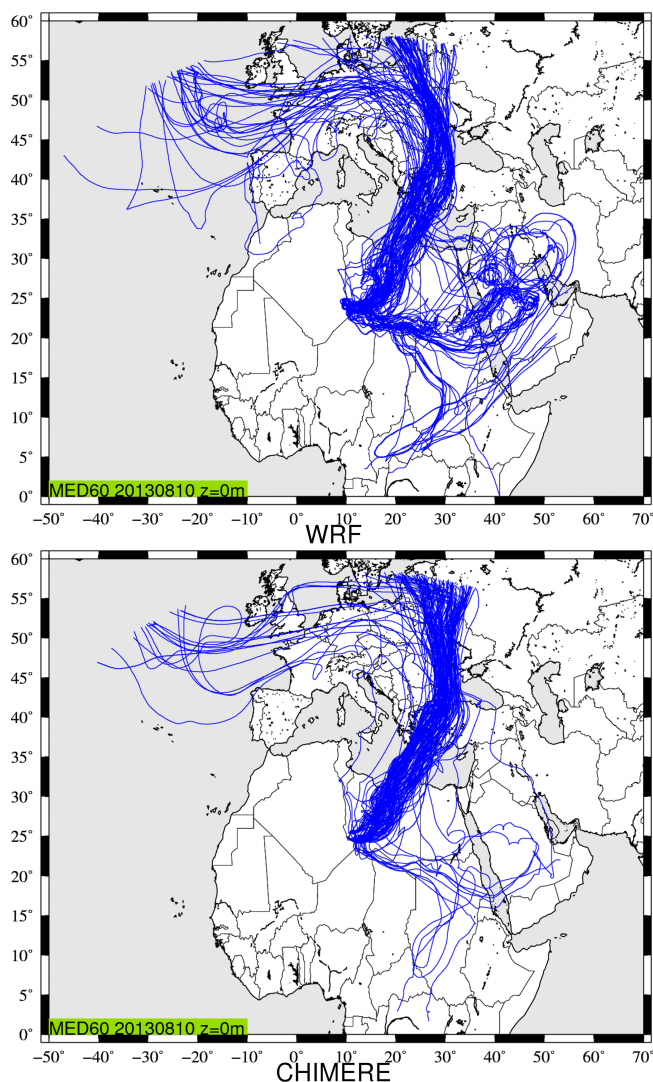


Figure 5. Back trajectories calculated using CHIMERE and WRF modelled meteorological fields. The starting point is at longitude 10° E and latitude 25° N, with an altitude of 0 m a.g.l. on 10 August 2013 at 12:00 UTC. It corresponds to a case studied during the Chemistry-Aerosol Mediterranean Experiment (ChArMEx) campaign (Menut et al., 2015).

diagnoses an altitude of ≈ 3000 m, CHIMERE diagnoses ≈ 2000 m, leading to different direction and wind speed. Then, this implies a split of the plumes with WRF but not with CHIMERE. This illustrates the sensitivity of the result to the driver model. But, in both cases, the answer in our case is clearly that the main contributions of the air masses located at the starting point are mainly coming from the north-east, crossing Tunisia and then the Mediterranean Sea and Europe. The main difference between the two calculations is the eastern part of the plume, which is more intense with WRF than CHIMERE.

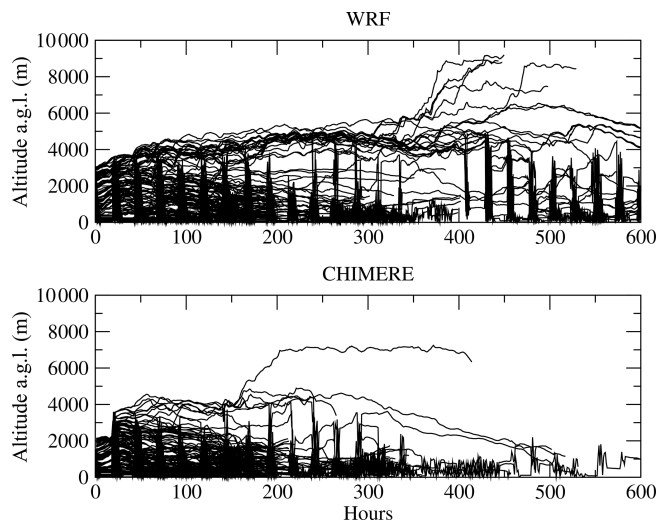


Figure 6. Projection of all back trajectories on a single time–altitude axis.

8 Conclusions

A new interpolation programme written in Fortran has been developed to interpolate on N -dimensional matrices. It has been evaluated for several dimension cases up to $N = 5$. The code is fast compared to similar Python routines and highly portable in existing geophysical codes. The interpolation programme works for any dimension N above 2 and is designed to work with irregular but structured grids (characterized by a size for each dimension) or look-up tables. Already used in its “regular” version in CHIMERE, the “general” programme has been tested on a new real application which calculates air mass back trajectories from two widely used atmospheric models: CHIMERE and WRF. This interpolation programme can be used for any application in geophysics and engineering sciences and also to explore large structured matrices.

Appendix A: List of frequently used abbreviations

AGL	Above ground level
CFL	Courant–Friedrichs–Lewy number
CHIMERE	National French CTM
CTM	Chemistry-transport model
CPU	Central processing unit
NMSE	Normalized mean square error
PBL	Planetary boundary layer
PSFC	Surface pressure
WRF	Weather Research and Forecasting model

Appendix B: Binary strategy

This piece of code shows the strategy to optimize the computation of weights for the “regular case”. The idea is to minimize the number of operations to benefit from the calculation at each dimension. A non-optimized loop would require $2^N - 1$ multiplications, while the optimized loop requires only $2^{N+1} - 4$ multiplications for the weight calculations. Then, for large values of $N \gg 2$, the ratio of required operations between the non-optimized and the optimized loop is $\approx (N - 1)/2$.

```
! ...
nn=2**ndim                ! ndim is the dimension of the case study
pn=2**(ndim-1)            !
! Loop to convert k in binary
do k=0,nn-1
  do j=0,ndim-1
    if (btest(k,j)) then
      ibin(j,k)=1
    else
      ibin(j,k)=0
    endif
  enddo
enddo
...
! Main optimized loop (2 loops sequence) to calculate the weighting factors benefiting from
! the previous iteration on the main dimension i
do i=1,ndim                ! Loop 1
  ni=2**i
  pi=ni/2
  do k=0,ni-1              ! Loop 2
    delta=ibin(i-1,k)
    if (i.ne.1) then
      if (delta.eq.0) then
        ws(k)=weight(k)
        weight(k)=weight(k)*w(i,delta) ! where w is the weight on each dimension
      else
        weight(k)=ws(k-pi)*w(i,delta)
      endif
    else
      weight(k)=w(i,delta)
    endif
  enddo
enddo
```

Appendix C: Code design

Note that *avedelta* and *maxdelta* arrays have been externalized to optimize the calculations. In the code package, an independent programme is available to calculate these arrays to be implemented in the user's main programme. The programme is written in Fortran *double precision* ingesting the following arguments:

```

subroutine interpolation_general(&
ndim,      &! Input : Int
maxdim,    &! Input : Int
kdim,      &! Input : Array 1D, Int
vect,      &! Input : Array 2D, Real
vtable,    &! Input : Array 1D, Real
table,     &! Input : Array 1D, Real
avedelta,  &! Input : Array 1D, Real
maxdelta,  &! Input : Array 1D, Real
table,     &! Input : Array 1D, Real
resu,      &! Output : Real
inei,      &! Output : Integer
neighbours,&! Output/Input : Array 2D, Int
weights,   &! Output : Array 1D, Real
found      &! Output : Logical
)

```

Some hard-coded variables can be tested by the user to improve the results. They have been tested and some results are described in this paper. A recompilation is necessary if the user changes these values.

```

logical ,parameter      :: norm=.false.    ! Normalize or not by the average delta on each
    dimension
logical ,parameter      :: verbose=.false. ! Level of message writing (.true. for debug)
logical ,parameter      :: neighb=.true.   ! .true. only find up to ndim+1 closest neighbours to
    be faster
                                     ! .false. find up to 2\^NDIM closest neighbours
integer ,parameter      :: iconf=1        ! Number of cell to account for before and after
                                     ! the closest point, iconf=2 can be tested not more
real (kind=iprec) ,parameter :: pnum=2.0d+00 ! p-distance parameter

```

Table C1. Description of subroutine arguments.

Variable	Type	Description	Array dimension
<i>ndim</i>	Integer	Dimension $N > 1$	nd
<i>maxdim</i>	Integer	Total number of elements of the input table $n = \prod_{i=1}^N I_i$ with I_i the number of elements in each dimension i	nd
<i>kdim</i>	Integer 1-D array	Array of the number of elements I_i on each dimension i	(0 : N)
<i>vect</i>	Real 2-D array	Array storing the list of edges on each dimension in a one-dimensional array	(1 : N , 1 : n)
<i>vtable</i>	Real 1-D array	Coordinate values of the point at which to interpolate data	(1 : N)
<i>table</i>	Real 1-D array	Values for the list of known points <i>vect</i> (input grid data)	(1 : n)
<i>avedelta</i>	Real 1-D array	Inverse of average intervals on each dimension N	(N)
<i>maxdelta</i>	Real 1-D array	Maximum intervals on each dimension N	(N)
<i>resu</i>	Real	Result of interpolation for <i>vtable</i>	nd
<i>inei</i>	Integer	Number of neighbours	nd
<i>neighbours</i>	Real 2-D array	Array of neighbour coordinates	(1 : 2^N , 1 : n)
<i>weights</i>	Real 1-D array	Weight for each neighbour	(1 : 2^N)
<i>found</i>	Logical	Returns <i>true</i> or <i>false</i> , respectively, if the result is found or not found if the point is outside the bounds	nd

Appendix D: Irregular structured grid example in 5-D

Below is an example of a 5-D array input grid data with irregular intervals with the last dimension (5) depending on dimension (1).

```

! Definition of main dimensions
isize=9                !Outputgrid size
npoints=35             !Inputgrid size
ndim=5                 !Dimension of the example
allocate(kdim(0:ndim)) !Number of element per dimension array
kdim(0)=1              !Fake dimension for computation issues
kdim(1)=npoints
kdim(2)=npoints
kdim(3)=npoints
kdim(4)=npoints
kdim(5)=npoints
! Main array allocation
allocate(delta(ndim))
allocate(dstart(ndim))
allocate(vect(ndim,kdim(1),kdim(2),kdim(3),kdim(4),kdim(5)))
allocate(rando(ndim,1:npoints-1)) !Random variable
maxdim=kdim(1)*kdim(2)*kdim(3)*kdim(4)*kdim(5)
!
! Definition of the input grid "vect" with delta, dstart as:
do j=1,ndim
  delta(j)=1.d+00/dfloat(npoints-1)/(dfloat(j)**3)
  dstart(j)=1.d+00/dfloat(npoints-1)/2.d+00/(dfloat(j)**3)
enddo
do j=1,ndim
  do i=1,npoints-1
    rando(j,i)=dfloat(int(dlog(rand()*10.d+00+1.d+00))+2) !Random interval on each dimension
  enddo
  rando(j,:)=dfloat(npoints-1)*rando(j,:)/sum(rando(j,:)) !Normalisation of random intervals
enddo
vect=0.0d+00 ! Initialisation
! "vect" computation
do i=2,kdim(1)
  vect(1,i,:,:,,:)= vect(1,i-1,:,:,,:)+rando(1,i-1)*delta(1)
enddo
do j=2,kdim(2)
  vect(2,,:,j,:,:,)= vect(2,,:,j-1,:,:,)+rando(2,j-1)*delta(2)
enddo
do k=2,kdim(3)
  vect(3,:::,k,:,:,)= vect(3,:::,k-1,:,:,)+rando(3,k-1)*delta(3)
enddo
do l=2,kdim(4)
  vect(4,:::,l,:,:,)= vect(4,:::,l-1,:,:,)+rando(4,l-1)*delta(4)
enddo
do m=2,kdim(5)
  do i=1,kdim(1)
    vect(5,i,:,:,m)=vect(5,i,:,:,m-1)+rando(5,m-1)*delta(5)+&
      &delta(5)*dfloat(i-1)/dfloat(kdim(1)-1) !Dim 5 depends on Dim. 1
  enddo
enddo
! ....

```

Appendix E: Characteristics of Machine 1

- Architecture: x86_64
- CPU op-mode(s): 32-bit, 64-bit
- Byte order: Little Endian
- CPU(s): 64
- Online CPU(s) list: 0–63
- Thread(s) per core: 2
- Core(s) per socket: 8
- Socket(s): 4
- NUMA node(s): 8
- Vendor ID: AuthenticAMD
- CPU family: 21
- Model: 1
- Model name: AMD Opteron Processor 6276
- Stepping: 2
- CPU MHz: 2300.000
- CPU max MHz: 2300.0000
- CPU min MHz: 1400.0000
- Bogomips: 4599.83
- Virtualization: AMD-V
- L1d cache: 16 KB
- L1i cache: 64 KB
- L2 cache: 2048 KB
- L3 cache: 6144 KB
- Memory block size: 128 MB
- Total online memory: 128 GB
- Total offline memory: 0 B
- Linux version 3.10.0-1062.12.1.el7.x86_64 (mock-build@kbuilder.bsys.centos.org) (gcc version 4.8.5 20150623 (Red Hat 4.8.5-39))

Appendix F: Characteristics of Machine 2

- Architecture: x86_64
- CPU op-mode(s): 32-bit, 64-bit
- Byte order: Little Endian
- CPU(s): 96
- Online CPU(s) list: 0–47
- Offline CPU(s) list: 48–95
- Thread(s) per core: 1
- Core(s) per socket: 24
- Socket(s): 2
- NUMA node(s): 2
- Vendor ID: GenuineIntel
- CPU family: 6
- Model: 85
- Model name: Intel Xeon Platinum 8168 CPU at 2.70 GHz
- Stepping: 4
- CPU MHz: 2701.000
- CPU max MHz: 2701.0000
- CPU min MHz: 1200.0000
- Bogomips: 5400.00
- Virtualization: VT-x
- L1d cache: 32 KB
- L1i cache: 32 KB
- L2 cache: 1024 KB
- L3 cache: 33 792 KB
- NUMA node0 CPU(s): 0–23
- NUMA node1 CPU(s): 24–47
- Memory block size: 128 MB
- Total online memory: 190.8 GB
- Linux version 3.10.0-957.41.1.el7.x86_64 (mockbuild@x86-vm-26.build.eng.bos.redhat.com) (gcc version 4.8.5 20150623 (Red Hat 4.8.5-36))

Appendix G: The WRF and CHIMERE model parameters used

Parameters are the three-dimensional wind components, the boundary layer height \bar{h} , the surface sensible heat flux Q_0 and the altitude of each model layer. The wind components are used for the horizontal and vertical transport. The boundary layer height is used to define the vertical extent of the possible mixing, and the surface sensible heat flux is used to know if the current modelled hour corresponds to a stable or unstable surface layer (for when the tracer is close to the surface).

Backplumes calculates the back trajectories using longitude, latitude and altitude in metres. In the case of input data with vertical levels in pressure coordinates, the altitude is calculated from pressure levels (Pielke, 1984). This is the case of the WRF model and the calculation is done as follows.

The altitude is computed as

$$p^* = p_{\text{surf}} - p_{\text{top}}, \quad (\text{G1})$$

where p_{surf} (PSFC) is the surface pressure and p_{top} is the top pressure of the model domain. If p_{top} is constant over the whole domain, p_{surf} and thus p^* are dependent on the first level grid.

$$z_0 = \frac{\Phi(1) + \Phi'(1)}{g}, \quad (\text{G2})$$

where Φ is the geopotential (PHB) and Φ' (PH) its perturbation at vertical level k . g is the acceleration of gravity, $g = 9.81 \text{ m s}^{-2}$. For each vertical level k , the layer thickness Δz and the cell top altitude z_k are estimated as

$$\begin{aligned} d_m &= \log \left(\frac{p^* \eta_M - p_{\text{top}}}{p^* \eta_M + p_{\text{top}}} \right) \\ d_u &= \log \left(\frac{p^* \eta_M - p_{\text{top}}}{p^* \eta_F + p_{\text{top}}} \right) \\ z_1 &= \frac{\Phi(k) + \Phi'(k)}{g} \\ z_2 &= \frac{\Phi(k+1) + \Phi'(k+1)}{g} \\ \Delta z &= (z_2 - z_1) \frac{d_u}{d_m} \\ z(k) &= z_1 + \Delta z - z_0, \end{aligned} \quad (\text{G3})$$

where η_M is its value on full (w) levels (ZNU) and η_F is the η value on half (mass) levels (ZNU). The layer thicknesses are space and time dependent; this calculation is performed for all trajectories and all time steps.

Table G1. List of parameters read by the Backplumes programme to calculate trajectories.

Parameter	Model variable name
WRF model	
Longitude, latitude	XLONG, XLAT
Parameters for altitude	P_TOP, ZNU, ZNW, P, PB PH, PHB, PSFC
Wind components	U, V, W
Q_0	HFX
\bar{h}	PBLH
CHIMERE model	
Longitude, latitude	lon, lat
Altitude	hlay
Wind	winz, winm, winw
Q_0	sshf
\bar{h}	hgt

The new position of a tracer back in time is calculated as follows:

$$\begin{aligned} \phi_{\text{rad}} &= \phi \frac{\pi}{180} \\ \Delta x &= u \frac{3600}{\Delta t} \\ \Delta y &= v \frac{3600}{\Delta t} \\ \Delta \lambda &= \frac{\Delta x}{R \cos(\phi_{\text{rad}})} \frac{180}{\pi} \\ \Delta \phi &= \frac{\Delta y}{R} \frac{180}{\pi}, \end{aligned} \quad (\text{G4})$$

with the wind speed is provided in m s^{-1} on an hourly basis, and R is the Earth's radius as $R = 6371 \text{ km}$. The new position for one tracer is thus

$$\begin{aligned} \lambda_{t-1} &= \lambda_t - \Delta \lambda \\ \phi_{t-1} &= \phi_t - \Delta \phi. \end{aligned} \quad (\text{G5})$$

Code availability. The current versions of the models are freely available. The exact version of the model used to produce the results used in this paper is archived on Zenodo for NterGeo at <https://doi.org/10.5281/zenodo.3733278> (Bessagnet, 2020) under the GNU General Public License v3.0 or later, as are input data and scripts to run the model and produce the plots for all the simulations presented in this paper. The Backplumes model is an open-source code and is available on the CHIMERE model website: <https://www.lmd.polytechnique.fr/~menut/backplumes.php> (last access: 21 December 2020, Menut, 2020).

Author contributions. BB developed the code. LM and BB co-developed and implemented the code in the Backplumes.v2020r1 model. MB supported BB for the developments.

Competing interests. The authors declare that they have no conflict of interest.

Financial support. This research was funded by the DGA (French Directorate General of Armaments; grant no. 2018 60 0074) in the framework of the NETDESA project.

Review statement. This paper was edited by Juan Antonio Añel and reviewed by two anonymous referees.

References

- Bessagnet, B.: A *N*-dimensional Fortran Interpolation Program (NterGeo) for Geophysics Sciences (Version 2020v1), Zenodo, <https://doi.org/10.5281/zenodo.3733278>, 2020.
- Donner, L. J., Wyman, B. L., Hemler, R. S., Horowitz, L. W., Ming, Y., Zhao, M., Golaz, J.-C., Ginoux, P., Lin, S.-J., Schwarzkopf, M. D., Austin, J., Alaka, G., Cooke, W. F., Delworth, T. L., Freidenreich, S. M., Gordon, C. T., Griffies, S. M., Held, I. M., Hurlin, W. J., Klein, S. A., Knutson, T. R., Langenhorst, A. R., Lee, H.-C., Lin, Y., Magi, B. I., Malyshev, S. L., Milly, P. C. D., Naik, V., Nath, M. J., Pincus, R., Ploshay, J. J., Ramaswamy, V., Seman, C. J., Shevliakova, E., Sirutis, J. J., Stern, W. F., Stouffer, R. J., Wilson, R. J., Winton, M., Wittenberg, A. T., and Zeng, F.: The Dynamical Core, Physical Parameterizations, and Basic Simulation Characteristics of the Atmospheric Component AM3 of the GFDL Global Coupled Model CM3, *J. Climate*, 24, 3484–3519, <https://doi.org/10.1175/2011JCLI3955.1>, 2011.
- Flamant, C., Deroubaix, A., Chazette, P., Brito, J., Gaetani, M., Knippertz, P., Fink, A. H., de Coetlogon, G., Menut, L., Colomb, A., Denjean, C., Meynadier, R., Rosenberg, P., Dupuy, R., Dominutti, P., Duplissy, J., Bourrianne, T., Schwarzenboeck, A., Ramonet, M., and Totems, J.: Aerosol distribution in the northern Gulf of Guinea: local anthropogenic sources, long-range transport, and the role of coastal shallow circulations, *Atmos. Chem. Phys.*, 18, 12363–12389, <https://doi.org/10.5194/acp-18-12363-2018>, 2018.
- Hardy, R.: Multivariate equations of topography and other irregular surfaces, *J. Geophys. Res.*, 71, 1905–1915, 1971.
- Hardy, R.: Theory and applications of the multiquadric-biharmonic method 20 years of discovery 1968–1988, *Computers and Mathematics with Applications*, 19, 163–208, [https://doi.org/10.1016/0898-1221\(90\)90272-L](https://doi.org/10.1016/0898-1221(90)90272-L), 1990.
- Hofstra, N., Haylock, M., New, M., Jones, P., and Frei, C.: Comparison of six methods for the interpolation of daily, European climate data, *J. Geophys. Res.-Atmos.*, 113, D21110, <https://doi.org/10.1029/2008JD010100>, 2008.
- Hong, S. Y., Noh, Y., and Dudhia, J.: A new vertical diffusion package with an explicit treatment of entrainment processes, *Mon. Weather Rev.*, 134, 2318–2341, <https://doi.org/10.1175/MWR3199.1>, 2006.
- Kouatchou, J.: NASA Modeling Guru: Basic Comparison of Python, Julia, Matlab, IDL and Java (2018 Edition), available at: <https://modelingguru.nasa.gov/docs/DOC-2676> (last access: 11 March 2020), 2018.
- 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 upstream influence of atmospheric observations: The Stochastic Time-Inverted Lagrangian Transport (STILT) model, *J. Geophys. Res.-Atmos.*, 108, 4493, <https://doi.org/10.1029/2002JD003161>, 2003.
- Mailler, S., Menut, L., di Sarra, A. G., Becagli, S., Di Iorio, T., Bessagnet, B., Briant, R., Formenti, P., Doussin, J.-F., Gómez-Amo, J. L., Mallet, M., Rea, G., Siour, G., Sferlazzo, D. M., Traversi, R., Udisti, R., and Turquety, S.: On the radiative impact of aerosols on photolysis rates: comparison of simulations and observations in the Lampedusa island during the ChArMEEx/ADRIMED campaign, *Atmos. Chem. Phys.*, 16, 1219–1244, <https://doi.org/10.5194/acp-16-1219-2016>, 2016.
- Mailler, S., Menut, L., Khvorostyanov, D., Valari, M., Couvidat, F., Siour, G., Turquety, S., Briant, R., Tuccella, P., Bessagnet, B., Colette, A., Létinois, L., Markakis, K., and Meleux, F.: CHIMERE-2017: from urban to hemispheric chemistry-transport modeling, *Geosci. Model Dev.*, 10, 2397–2423, <https://doi.org/10.5194/gmd-10-2397-2017>, 2017.
- Menut, L.: BACKPLUMES program v2020r1 April 2020, IPSL-LMD-CNRS, available at: <https://www.lmd.polytechnique.fr/~menut/backplumes.php>, last access: 21 December 2020.
- Menut, L., Rea, G., Mailler, S., Khvorostyanov, D., and Turquety, S.: Aerosol forecast over the Mediterranean area during July 2013 (ADRIMED/CHARMEX), *Atmos. Chem. Phys.*, 15, 7897–7911, <https://doi.org/10.5194/acp-15-7897-2015>, 2015.
- Nehrkorn, T., Eluszkiewicz, J., Wofsy, S., Lin, J., Gerbig, C., Longo, M., and Freitas, S.: Coupled weather research and forecasting-stochastic time-inverted lagrangian transport (WRF-STILT) model, *Meteorol. Atmos. Phys.*, 107, 51–64, <https://doi.org/10.1007/s00703-010-0068-x>, 2010.
- Nenes, A., Pilinis, C., and Pandis, S.: ISORROPIA: A new thermodynamic model for inorganic multicomponent atmospheric aerosols, *Aquatic Geochem.*, 4, 123–152, 1998.
- Nenes, A., Pandis, S. N., and Pilinis, C.: Continued development and testing of a new thermodynamic aerosol module for urban and regional air quality models, *Atmos. Environ.*, 33, 1553–1560, 1999.
- Pielke Sr., R. A.: Mesoscale Meteorological Modeling, Academic Press, 1984.

- Pisso, I., Sollum, E., Grythe, H., Kristiansen, N. I., Casiani, M., Eckhardt, S., Arnold, D., Morton, D., Thompson, R. L., Groot Zwaaftink, C. D., Evangeliou, N., Sodemann, H., Haimberger, L., Henne, S., Brunner, D., Burkhardt, J. F., Fouilloux, A., Brioude, J., Philipp, A., Seibert, P., and Stohl, A.: The Lagrangian particle dispersion model FLEX-PART version 10.4, *Geosci. Model Dev.*, 12, 4955–4997, <https://doi.org/10.5194/gmd-12-4955-2019>, 2019.
- Rap, A., Ghosh, S., and Smith, M. H.: Shepard and Hardy Multiquadric Interpolation Methods for Multicomponent Aerosol–Cloud Parameterization, *J. Atmos. Sci.*, 66, 105–115, <https://doi.org/10.1175/2008JAS2626.1>, 2009.
- Scipy, C.: Interpolate unstructured D-dimensional data, available at: <https://docs.scipy.org/doc/scipy-0.14.0/reference/generated/scipy.interpolate.griddata.html> (last access: 21 December 2020), 2014.
- Shepard, D.: A Two-Dimensional Interpolation Function for Irregularly-Spaced Data, in: *Proceedings of the 1968 23rd ACM National Conference, Association for Computing Machinery, New York, NY, USA*, 517–524, <https://doi.org/10.1145/800186.810616>, 1968.
- Skamarock, W. C., Klemp, J. B., Dudhia, J., Gill, D. O., Barker, D. M., Duda, M., Huang, X., Wang, W., and Powers, J. G.: A description of the Advanced Research WRF Version 3, NCAR Tech. Note, 1–125, 2008.
- 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, *B. Am. Meteorol. Soc.*, 96, 2059–2077, <https://doi.org/10.1175/BAMS-D-14-00110.1>, 2015.
- Stull, R.: *an introduction to boundary layer meteorology*, Kluwer Academic Publishers Group, 1988.
- Sun, T. and Grimmond, S.: A Python-enhanced urban land surface model SuPy (SUEWS in Python, v2019.2): development, deployment and demonstration, *Geosci. Model Dev.*, 12, 2781–2795, <https://doi.org/10.5194/gmd-12-2781-2019>, 2019.
- Troen, I. and Mahrt, L.: A simple model of the atmospheric boundary layer: Sensitivity to surface evaporation, *Bound.-Lay. Meteorol.*, 37, 129–148, 1986.
- Wessel, P., Luis, J. F., Uieda, L., Scharroo, R., Wobbe, F., Smith, W. H. F., and Tian, D.: The Generic Mapping Tools Version 6, *Geochem. Geophys. Geosy.*, 20, 5556–5564, <https://doi.org/10.1029/2019GC008515>, 2019.
- Zender, C. S.: Analysis of self-describing gridded geoscience data with netCDF Operators (NCO), *Environ. Model. Softw.*, 23, 1338–1342, <https://doi.org/10.1016/j.envsoft.2008.03.004>, 2008.