Advecting Superspecies: Efficiently Modeling Transport of Organic Aerosol with a Mass-Conserving Dimensionality Reduction Method

Patrick Obin Sturm[1,3], Astrid Manders[2], Ruud Janssen[2], Arjo Segers[2], Anthony S. Wexler[1,4], Hai Xiang Lin[3,5]

[1]Air Quality Research Center, University of California, Davis, CA, USA
[2]Department of Climate, Air and Sustainability, TNO, Utrecht, the Netherlands
[3]Delft Institute of Applied Mathematics, Delft University of Technology, Delft, the Netherlands
[4]Departments of Mechanical and Aerospace Engineering, Civil and Environmental Engineering, and Land, Air and Water Resources, University of California, Davis, CA, USA
[5]Institute of Environmental Sciences, Leiden University, the Netherlands

Key Points:

- We develop a machine learning method to find a reduced-dimension set of superspecies representing tracers in a chemical transport model
- This method is designed to be physically consistent, preserving information on phase and conserving mass to numerical precision
- Advecting the superspecies reduces computation time to 56-66% of the original advection computation time

Corresponding author: Patrick Obin Sturm , posturm@ucdavis.edu

Abstract

The chemical transport model LOTOS-EUROS uses a volatility basis set (VBS) approach to represent the formation of secondary organic aerosol (SOA) in the atmosphere. Inclusion of the VBS approximately doubles the dimensionality of LOTOS-EUROS and slows computation of the advection operator by a factor of two. This complexity limits SOA representation in operational forecasts. We develop a mass-conserving machine learning (ML) method based on matrix factorization to find latent patterns in the VBS tracers that correspond to a lower-dimension set of superspecies. Tracers are reversibly compressed to superspecies before transport, and the superspecies are subsequently decompressed to tracers for process-based SOA modeling. This physically interpretable ML method conserves the total concentration and phase of the tracers throughout the process. The superspecies approach is implemented in LOTOS-EUROS and found to accelerate the advection operator by a factor of 1.5 to 1.8. Concentrations remain numerically stable over model simulation times of two weeks, including simulations at higher spatial resolutions than the ML models were trained on. Results from this case study show that this method can be used to enable detailed, process-based secondary organic aerosol representation in air quality operational forecasts in a computationally efficient manner. Beyond this case study, the physically consistent ML approach developed in this work enforces conservation laws that are essential to other Earth system modeling applications, and generalizes to other processes where computational benefit can be gained from a two-way mapping between detailed process variables and their representation in a reduced-dimensional space.

Plain Language Summary

The chemical composition of the atmosphere is a complex system involving many physical processes. Computer models can be used to improve our understanding of how these processes interact, as well as simulate hypothetical scenarios to support scientifically-informed climate and air quality policies. However, complicated models with many variables can take a lot of time to run. The LOTOS-EUROS model spends a lot of time and energy on computer simulations of the transport of chemical species, like particulate matter, by wind. Our method combines artificial intelligence with scientific fundamentals to reduce the number of variables: we model representative combinations of chemical species that are transported all at once, rather than transport each species individually. This leads to faster and cheaper simulations without loss of scientific detail or internal consistency.

1 Introduction

Vast amounts of computational resources are required to model phenomena in the earth sciences. This includes complex models of atmospheric composition that couple a large number of properties and processes (Brasseur & Jacob, 2017). Machine learning (ML) approaches are an emerging technique for decreasing the computational burden of earth system models with more efficient ML parameterizations, but have documented challenges such as unstable error growth and physical inconsistency which can happen when predicted recurrently (Kelp et al., 2018) or when interacting with other processes in the context of larger models (Rasp et al., 2018; Brenowitz & Bretherton, 2019). One approach towards ML models that can stably interact with other model processes is online training: parameter optimization of ML surrogates while running the entire model (Rasp, 2020; Kelp et al., 2022).

Other recent efforts have aimed to constrain ML approaches using scientific knowledge to ensure physically consistent results. One strategy for physically consistent ML models reposes the learning targets: rather than estimate important properties or their tendencies, instead estimate fluxes between the properties. The fluxes can then be related to tendencies in a way that balances mass, energy, or atoms (Sturm & Wexler, 2020; Yuval, O'Gorman, & Hill, 2021; Sturm & Wexler, 2022). Custom neural network architectures can also obey conservation laws by incorporating hard constraints (e.g. flux balances) in their hidden layers (Beucler et al., 2021; Sturm & Wexler, 2022): this can also improve the physical interpretability of the inner working of neural networks. Acknowledgement of the importance of physically consistent ML tools in the earth sciences (Keller & Evans, 2019; Yuval, Pritchard, et al., 2021) has informed the mass-conserving machine learning parameterization in this paper.

Within the field of atmospheric chemistry modeling, Kelp et al. (2020) have made progress towards a stable ML neural network emulating a box model of chemistry and aerosol microphysics processes, through training parameters on the accuracy of multiple future timesteps after predicting in a lower-dimensional latent space. Kelp et al. (2020) pose a future research direction: how the low-dimensional representation of chemical species might interact with other processes, such as advection, in a larger 3-dimensional model.

The current work develops and explores a physically consistent machine learning method that compresses the high dimensional set of organic aerosol (OA) tracers to reduce the computational cost of advection in the LOTOS-EUROS chemical transport model (CTM) (Manders et al., 2017). We use unsupervised machine learning approaches to find manifold dimensions, or characteristic regimes, of organic aerosol tracers. The characteristic regimes are used to form lower-dimensional combinations of OA tracers, interpreted as superspecies, which require fewer transport calculations. These superspecies are mapped back to the full OA tracer space after the advection operator. Additional constraints are applied when compressing to and decompressing from the reduced-dimension space, to conserve mass to machine precision. We compare the linear and additive method of non-negative matrix factorization to a nonlinear and more complex neural network autoencoder, and make a model selection after evaluating several configurations based on reconstruction accuracy and physical consistency.

Organic aerosol forms an important contribution to particulate matter (Jimenez et al., 2009). OA can be emitted to the atmosphere as semi-volatile primary organic aerosol (POA) through various direct sources, including vehicle exhaust, wildfire smoke, and residential wood combustion. OA can also be formed in the atmosphere as secondary organic aerosol (SOA) through gas-phase reactions of

volatile organic compounds (VOCs), which tend to form less volatile products: semi and intermediate volatility organic compounds (siVOCs) that can partition appreciably to the particle phase. Both anthropogenic sources, like industrial activity, and biogenic sources, such as forests, emit VOC precursors to SOA. Another source of SOA is the partial evaporation of POA to siVOCs, which in turn react and partition to form SOA (Robinson et al., 2007). This SOA from evaporated and aged POA is often chemically distinct from POA, showing a higher degree of oxidation (Jimenez et al., 2009), and can be tracked separately in models. SOA can form a significant fraction of the total OA concentration (de Gouw et al., 2005; Heald et al., 2005).

Due to the large number of distinct organic species in the atmosphere, organic aerosols are rarely speciated in models, but often lumped together into volatility bins according to the magnitude of their saturation vapor pressures (Donahue et al., 2006). This modeling approach is called the volatility basis set (VBS) and accounts for the tendency of compounds to become less volatile as they are oxidized. The partitioning between gas and particle phase in each volatility bin is governed by its corresponding saturation vapor pressure and the total OA concentration. A 2D-VBS extension has also been developed that includes oxygen to carbon ratio along another dimension (Jimenez et al., 2009; Donahue et al., 2011), which can account for fragmentation of larger compounds and estimation of hygroscopicity (Jimenez et al., 2009). However, the 1D-VBS approach is most commonly applied in chemical transport models, including separate basis sets for different classes of OA precursors (Bergström et al., 2012; Hayes et al., 2015; Janssen et al., 2017; Jiang et al., 2019) Use of multiple VBS classes enables distinct properties per class and can give insight into different aerosol systems contributing to total OA. However, increased complexity and dimensionality adds a computational burden to CTM calculations. This limits the inclusion of detailed, process-based OA modeling in chemical transport models like LOTOS-EUROS v2.2.1 (Manders-Groot et al., 2021), which uses 4 VBS classes based on the configuration from Bergström et al. (2012). The VBS module in LOTOS-EUROS v2.2.1 is not used by default, and when included, significantly increases wall time of simulations. The inclusion of VBS tracers adds computation time to other operators in the model relatively more than OA-specific calculations themselves.

Most notably, the high dimensionality caused by the VBS tracers adds a computational burden to the advection operator in LOTOS-EUROS v2.2.1, which is based off of Walcek (2000). Model timing experiments in Sturm (2021) found that wall time for the advection operator can double when using the VBS module. Advection is a bulk process and does not perform OA-specific calculations. This motivates a reduced-order approach: rather than advecting each tracer separately, instead advect a smaller set of superspecies formed from combinations of the VBS tracers. We leverage the large amount of model output for the VBS tracers, and develop a mass-conserving, data-driven machine learning approach to find latent patterns in the VBS tracers that allow for a more parsimonious representation of OA in transport processes. Though demonstrated for compression of OA and related compounds during transport, the methods developed in this work generalize to other Earth system applications, enabling use of high-dimensional process models whose variables can be reversibly compressed to a physically consistent reduced-dimension representation for use in other processes.

2 Materials and Methods

2.1 VBS approach in LOTOS-EUROS

The chemical transport model LOTOS-EUROS v2.2.1 uses a VBS scheme based off of the approach in Bergström et al. (2012) that has 4 distinct VBS classes: POA, SOA from siVOCs that are chemically aged after evaporating from semi-volatile POA emissions, (abbreviated as siSOA), and SOA from anthropogenic and biogenic gaseous VOCs abbreviated as aSOA and bSOA respectively.

Figure 1 provides an overview of the 58 tracers specific to the VBS module. Primary organic material (POM) emissions are modeled using a 9-bin VBS approach: the logarithmically distributed bins represent semi-volatile organics with effective saturation concentrations ranging from 10^{-2} to 10^{5} $\mu g\ m^{-3}$ at 298 K. The reported mass of primary emissions is distributed over the lower 4 volatility bins. As in previous work (Shrivastava et al., 2008), an additional 1.5 times this mass is distributed over the highest 5 volatility bins to represent non-reported intermediate volatility organic compounds (IVOCs). Only a fraction of the total mass remains in the particle phase: the fraction that evaporates is assumed to be semi-volatile VOCs (SVOCs), with effective saturation concentrations on the order of $1 < C^* < 10^3 \mu g\ m^{-3}$ or IVOC, with saturation concentrations on the order of $10^3 < C^* < 10^6\ \mu g\ m^{-3}$, defined at 298 K. The S/IVOCs undergo oxidation by the hydroxyl radical OH and enter the distinct siSOA VBS class. As material moves from the POA VBS to the siSOA VBS, it also moves to lower volatility bins, as shown in Figure 1. The total siSOA is represented by an 8-bin VBS using effective saturation concentrations from 10^{-2} to 10^{3} $\mu g\ m^{-3}$ (defined at 298 K). Each bin uses two tracers, one aerosol and one gas, to represent the partitioning: this results in 18 tracers for the POA VBS class and 16 tracers for the siSOA VBS class.

Formation of SOA from anthropogenic VOCs is represented with a 6-bin VBS class, defined using effective saturation concentrations of 10^{-2} to 10^{3} $\mu g\ m^{-3}$ at 298 K. This results in 12 tracers (6 in the gas phase and 6 in the particle phase). VOCs such as aromatics, alkenes and alkanes are classified in LOTOS-EUROS as anthropogenic precursors of secondary organic aerosols and upon oxidation are distributed over the 4 highest volatility bins as done by Tsimpidi et al. (2010), linearly interpolating between a low-NOx and high-NOx case as originally suggested by Lane et al. (2008).

An analogous 6-bin VBS class is used to model SOA formation from the biogenic VOCs in LOTOS-EUROS: monoterpene and isoprene. Yields from biogenic gaseous precursors are distributed over the 4 highest volatility bins according to Tsimpidi et al. (2010), with yields calculated by a branching ratio continuously dependent on NOx (Lane et al., 2008). Unlike the anthropogenic VBS class, ageing between bins is turned off for the biogenic VBS in LOTOS-EUROS v2.2.001, as in prior work (Murphy & Pandis, 2009; Tsimpidi et al., 2010, 2014; Matsui, 2017). This is informed by the low sensitivity of biogenic SOA concentration to oxidative ageing (Ng et al., 2006; Donahue et al., 2012), thought to arise from fragmentation effects that balance out functionalization effects on volatility (Murphy et al., 2012). For this reason, material never enters the 2 lowest volatility bins in LOTOS-EUROS v2.2.1, rendering the 4 corresponding tracers effectively inert. However, in LOTOS-EUROS v2.2.1 with the VBS module on, these 4 tracers are still dealt with by the model, contributing to the computational burden on processes such as advection.

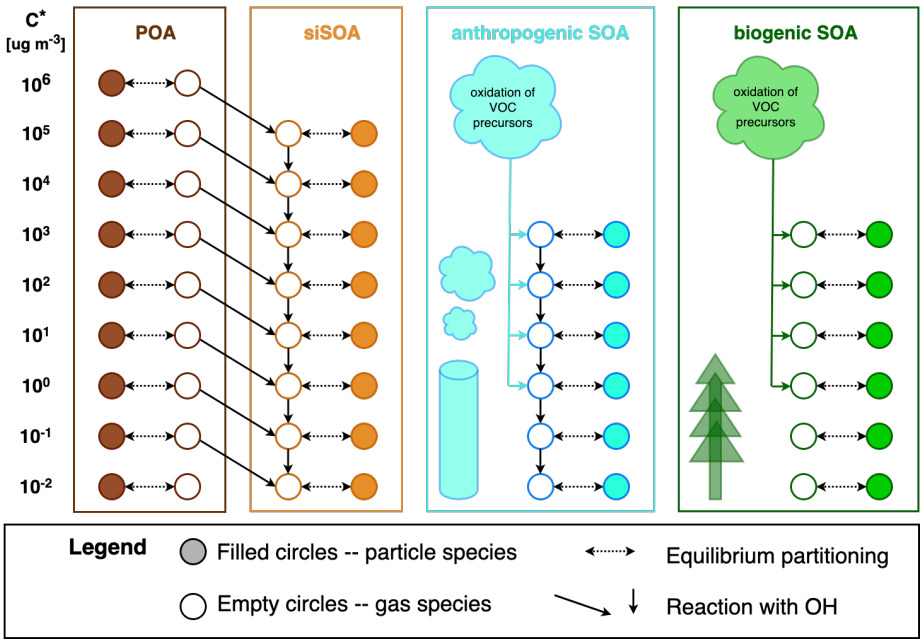

Figure 1. Schematic representation of the VBS approach in LOTOS-EUROS v2.2.1, including the 4 VBS classes with 58 tracers, and their thermodynamic and chemical relationships. This diagram was inspired by the schematic in Shrivastava et al. (2008).

2.2 Tracer compression methods

A method for tracer compression for transport in the GEOS-Chem global CTM is given by Liao et al. (2007), where various oxidation products are lumped together by phase and class, and assumed to behave similarly in transport. The relative compositions are stored separately and used to distribute the lumped tracers back to individual products after transport. Another approach for OA tracers given by Matsui (2017) compresses VBS tracers in a global aerosol model from 106 to 26 (a compression factor of approximately 4) by using fewer volatility bins. This effectively lowers the bin resolution and combines material across a wider range of saturation vapor pressures. Analogously, Matsui (2017) converts between high-resolution and low-resolution bins in a sectional aerosol model for use in processes not directly related to aerosols. An example of tracer compression for advection in a 2D-VBS is given by Zhao et al. (2020) who sum tracers along the O:C axis, resulting in a 1D-VBS for decreased dimensionality in advection.

One compression technique for advection of 1D-VBS tracers could be developed based on partitioning, where the compressed tracers themselves contain all the information needed to decompress to the VBS tracer space without losing resolution. The total concentration for each volatility bin could be advected, as well as total OA concentration, reducing the 58 phase-specific tracers to 29 combined phase tracers, and an additional tracer to keep track of total organic aerosol concentration. After advection, total OA along with the saturation vapor

concentration determines the partitioning between phase in each volatility bin. This compression is not "lossy", meaning no information is lost on compression itself. However, such a strategy yields a compression factor of only approximately 2, about halving the number of tracers. We seek a compression technique that can reduce the number of tracers further, which may be lossy but leverages a large amount of representative model output, using machine learning algorithms optimized to accurately reconstruct VBS distributions.

2.3 Model Configuration

To find latent patterns for a reduced order representation of the 58 VBS tracers, we use LOTOS-EUROS version 2.2.1 (Manders-Groot et al., 2021; Manders et al., 2017) with the optional VBS module. The model was used in its default configuration using 5 levels, the first one being a 25 m surface layer, the second layer reaching the top of the mixing layer, and the other three layers being reservoir layers up to 5 km altitude. The horizontal domain covers 15°W to 35°E and 35-70°N on a lonxlat grid of 0.5x0.25°. This grid is termed the 'MACC' grid, referring to the grid used in the project 'Monitoring Atmospheric Composition and Change, a predecessor of the current CAMS (Copernicus Atmospheric Monitoring Service). Meteorology is taken from ECMWF IFS 12-hour operational forecasts, using hourly surface values and 3-hour 3-D fields interpolated to hourly values. For gas-phase chemistry, a condensed and slighty modified version of CBM-IV is used (Gery et al., 1989). Wet deposition includes in-cloud and below-cloud scavenging as described in Seinfeld and Pandis (2006), deposition of gases is calculated using DEPAC (Zanten et al., 2010), and deposition of particles follows Zhang (2001). The model includes tree-specific biogenic isoprene and terpene emissions as described in Beltman et al (2013) using a high-resolution tree-species database (Köble & Seufert, 2001) that are combined with land cover data from CORINE2000 (EEA, 2005). Anthropogenic emissions are CAMS emissions for 2015 (CAMS-regional air pollutants as delivered in 2018) with a bottom-up estimation for residential wood combustion emissions, providing the best estimate of organic carbon emissions (Denier van der Gon et al., 2015). Wildfire emissions are taken from the MACC global fire assimilation system (Kaiser et al., 2012). Initial and boundary conditions for most species were taken from CAMS-near real-time, for organic matter no boundary conditions were used.

Short simulations of 14 days in the last two weeks of February and July 2018 were used with 5 days of spin-up, the subsequent 5 days for training and the last 4 days for evaluation. Evaluation of the simulations with observations is outside the scope of the present paper, as the model is regularly evaluated in model validation reports, as well as CAMS ensemble and model evaluations.

2.4 Linear Approach

A linear approach could be used to project the tracer space into a lower dimensional subspace allowing linear combinations of the tracers to be passed to the advection operator. Principal component analysis is a common linear projection method but is mean-centered and can lead to negative values, which are less readily interpretable as concentrations. Non-negative matrix factorization (NMF), also called positive matrix factorization, is an unsupervised machine learning algorithm chosen in applications where values must remain non-negative, for example pixel values in image compression (Lee & Seung, 1999) or concentrations in the physical sciences (Paatero & Tapper, 1994). Given a matrix of non-negative data $\mathbf{V} \in \mathbb{R}^{m \times n}$ with m dimensions and n data points, the NMF algorithm returns two non-negative approximate factors of $\mathbf{V}$ according to an objective function

$$\underset{\mathbf{W},\mathbf{H}}{argmin} \|\mathbf{V} - \mathbf{WH}\| \qquad s.t. \quad \mathbf{W}, \mathbf{H} \geq 0 \tag{1}$$

where $\mathbf{W} \in \mathbb{R}_{\geq 0}^{m \times r}$ is a mapping from the m dimensional space to a lower dimensional latent space with r features, and $\mathbf{H} \in \mathbb{R}_{\geq 0}^{r \times n}$ is the latent space representation of each data point. The inequality is interpreted as an element-wise constraint. We use the Frobenius norm in the objective function, which is the default NMF norm in the scikit-learn Python package (Pedregosa et al., 2011).

The NMF algorithm operates on a data matrix, handling batches of observations all at once. For our application, compression of current concentrations of VBS tracers $\vec{v} \in \mathbb{R}^m$ to a lower dimensional space needs to happen with each new time step. For the purpose of speeding computations, it might be counterproductive to perform the NMF algorithm online in every time step. If $\mathbf{W}$ is optimized using equation 1 on sufficiently representative training data, it can be used to decompress a set of superspecies $\vec{h} \in \mathbb{R}^r$ to a decompressed set of tracers $\vec{v}_{dec} \in \mathbb{R}^m$ to approximate $\vec{v}$. However, we still need to obtain the superspecies vector $\vec{h}$. Given a sufficiently representative $\mathbf{W}$, we can use its Moore-Penrose pseudoinverse $\mathbf{W}^+ \in \mathbb{R}^{r \times m}$ to compress a new set of tracers $\vec{v}$ to a corresponding set of new superspecies $\vec{h}$. $\mathbf{W}^+$ may have negative elements for $r > 1$ (more than one superspecies, or degree of freedom), theoretically yielding negative superspecies values or decompressed tracer values. This potential limitation is quantified in section 3.2. Instead of a Moore-Penrose pseudoinverse, a positive-valued compression matrix $\mathbf{B} \in \mathbb{R}_{\geq 0}^{r \times m}$ can be obtained by similar non-negative matrix factorization methods, using the objective function:

$$\underset{\mathbf{B}}{argmin} \|\mathbf{H} - \mathbf{BV}\|_F^2 \qquad s.t. \quad \mathbf{B} \geq 0 \tag{2}$$

The full approach to obtain non-negative compression and decompression matrices then becomes

1. Given tracer data $\mathbf{V}$, find $\mathbf{H}$, $\mathbf{W}$ such that $\mathbf{V} - \mathbf{WH}$ is minimized.
2. Given tracer data $\mathbf{V}$, and using $\mathbf{H}$ from the previous step, find $\mathbf{B}$ such that $\mathbf{H} - \mathbf{BV}$ is minimized.
3. Use $\mathbf{B}$ to compress subsequent observations of VBS tracers $\vec{v}$ to a non-negative vector of superspecies $\vec{h}$, and $\mathbf{W}$ to decompress $\vec{h}$ to the original tracer space $\vec{v}_{dec}$.

The compression and decompression matrices $\mathbf{B}$ and $\mathbf{W}$ are optimized for each VBS class, to avoid mixing different classes of OA that have different properties (e.g. molar mass). An important hyperparameter of this approach is r, the size of the latent space (number of superspecies). This can be chosen by constructing an elbow plot of error metrics with varying r, while also considering compression factor.

2.5 Nonlinear Approach

We investigate whether a more complicated model than the pair of non-negative matrices is appropriate for compressing VBS tracers. Given the recent success of artificial neural networks (NNs) in emulating models of atmospheric phenomena (Beucler et al., 2021; Kelp et al., 2020), we construct a neural network autoencoder that can reversibly compress the VBS tracers to a latent space. Neural networks are connected networks of artificial neurons: each neuron calculates a linear combination of its input, adds a bias scalar, and feeds this result to a (usually non-linear) activation function (Marsland, 2014). Neurons performing this operation

on the same input in parallel are designated as a layer within the neural network. Neural networks can have multiple such layers: vector output from neuron layers that are not final output of the NN are called hidden layers. A neural network autoencoder attempts to replicate the identity function via compression, where hidden layers compress the input to the NN to a smaller latent space of size r. For our application, the activation function chosen for each neuron is a rectified linear unit that outputs the maximum of its input and zero. This choice of activation function constrains output of both the hidden layer and the NN output to their respective positive half-spaces. In other words, like the non-negative compression/decompression matrices in section 2.4, this activation function ensures concentrations will not go below zero.

While matrix multiplication to a lower-dimensional space is also part of the linear approach in section 2.4, the neural network adds complexity in its parameter space via multiple layers with weight parameters, as well as bias and activation functions between layers of neurons. This model should be chosen if it significantly outperforms the linear method using the same size r. As the NNs are compared directly to the linear method, one NN per VBS class is chosen.

Training a neural network involves optimizing the coefficients of the linear combination and bias scalar for each perceptron through local minimization methods, often gradient descent. To prevent overfitting of the NNs, dropout layers are used to temporarily remove some neurons during training, and training of NNs is stopped if predictions are not improved on a set of validation data after a certain number of passes through the training data. The neural network models are constructed and trained using the Keras library (Chollet et al., 2015) using a Tensorflow backend (Abadi et al., 2016).

2.6 Physically Consistent Models: Conserving Mass and Phase

Sections 2.4 and 2.5 developed methods to ensure non-negativity of both the compressed superspecies and decompressed tracers. This section refines the linear method to encode other physical information: concentration and phase.

An advantage of the linear method is that the direction of the decompressed tracer space is invariant to scaling of the superspecies space. In other words, the concentration of superspecies can be adjusted without changing the relative volatility distribution of the decompressed tracers. We can use a scaling factor after compression to ensure that the total concentration of superspecies is equal to the total concentration of the tracers for each VBS class. Similarly, after decompression, we can ensure that the total concentration of decompressed tracers is equal to the total concentrations of superspecies. This ensures that compression and decompression neither add nor remove mass. The scaling factor s_{com} after using $\mathbf{B}$ to compress tracers $\vec{v}$ to the superspecies vector $\vec{h}$ is

$$s_{com} = \frac{\sum_{i=1}^{m} v_i}{\sum_{j=1}^{r} h_j} \tag{3}$$

After decompression to $\vec{v}_{dec}$ using $\mathbf{W}$, the decompressed tracers can be scaled using a factor s_{dec}, where

$$s_{dec} = \frac{\sum_{j=1}^{r} h_j}{\sum_{i=1}^{m} v_{dec,i}} \tag{4}$$

Despite conserving total concentration of all tracers, the concentration of total organic aerosol (TOA) may not be conserved due to errors in the mass distribution over volatility bins after decompression. A variation of this method to conserve TOA instead of total concentration, as well as an alternative way to conserve total concentration only using **W** from NMF, is explored in Sturm (2021). However, the compromise of conserving TOA versus total concentration is avoidable by adding another cross section: creating compression and decompression matrices **B** and **W** for each phase as well as VBS class, e.g. one transformation for all biogenic gaseous VBS tracers and a separate transformation for all biogenic particle tracers.

2.7 Implementing the ML Parameterization into a CTM

We select the most promising method (section 3.3) based on accuracy, compression factor, and physical consistency. This method is then built into a version of LOTOS-EUROS v2.2.1, extended to include the ML superspecies parameterization. Additional tracers for superspecies are added to the model and adopt the characteristics of their respective tracers. Subroutines were added to the VBS module to load the ML parameterizations, as well as perform the compression and decompression operations. These subroutines are then called in the driver program:

1. The initialization subroutine loads offline-optimized ML parameterizations before the time loop starts.
2. Within the time loop, directly before the call to the advection operator, the compression subroutine is called to map VBS tracers to superspecies concentrations, overwriting the current superspecies values. The advection operator is resolved using superspecies instead of the VBS tracers, which are skipped.
3. Within the time loop, directly after the call to the advection operator, the decompression routine is called to transform superspecies into VBS tracers, overwriting previous VBS tracer values which were not included in advection calculations.

3 Results: Model Development and Selection

3.1 Compression Factor and Accuracy

To obtain a sense of error obtained by a maximum compression factor and the simplest model, we use NMF with a single superspecies ($r = 1$) per VBS class to obtain a decompression matrix (in this case a vector) **W** and calculate its pseudoinverse $\mathbf{W}^+$ to be used for compression. This compression strategy is evaluated on reconstruction accuracy of test model output of the entire domain and time period, using average bias and root mean square error (RMSE). While bias is an indicator of the total material that is introduced or removed artificially by compression, RMSE is an absolute metric that indicates how accurately the reconstructed VBS tracers reproduce the volatility distribution. Table 1 shows both reconstruction error metrics for the tracer set of each class, as well as the reconstruction bias and RMSE's of total organic aerosol concentration (TOA) and total organic material (TOM) from summing across VBS classes. The mean

concentrations for each VBS class, as well as TOA and TOM, are included for comparison.

Table 1. Test reconstruction error metrics using the NMF/Pseudoinverse approach with 1 superspecies per VBS class.

	Mean [$\mu g\, m^{-3}$]	RMSE [$\mu g\, m^{-3}$]	Bias [$\mu g\, m^{-3}$]
aVOC	0.0043	0.0021	-3.9×10^{-6}
bVOC	0.0262	0.0061	2.9×10^{-4}
POA	0.0558	0.0441	-0.0021
siSOA	0.0153	0.0205	6.4×10^{-5}
TOA	0.386	0.266	0.094
TOM	1.61	0.0978	-0.0328

Moreover, by the use of a single superspecies ($r = 1$) the tracers pass through a linear transformation of rank 1: the concentration distribution over the volality bins will always have the same shape, with gridcells and different time steps differing only in magnitude, as scaled by h. This means any spatiotemporal variability of the distribution shape will be lost after passing through a single-dimensional superspecies space. More complexity is needed to capture variation in volatility distribution. This motivates larger matrices that have more degrees of freedom r, which comes at the cost of compression factor: the inevitable trade-off between accuracy and computational speed is a canonical problem in atmospheric chemistry modeling. Figure 2 visualizes the effect of compression extent on accuracy, using $\mathbf{W}^+$ to convert to superspecies and $\mathbf{W}$ to map back to tracers. Reconstruction accuracy is reported for the set of tracers in each class (both particle and gas) as well as TOA (total organic aerosol, calculated by summing the concentrations of particle tracers across classes).

Figure 2 shows RMSE monotonically decreasing with increasing number of superspecies, with diminishing returns after 3 superspecies. More superspecies to advect will increase the computational burden of the advection operator in LOTOS-EUROS without a substantial improvement in RMSE or bias. In light of the desire to maximize compression factor, the two elbow plots indicate that 3 superspecies strikes a good balance between dimension reduction and accuracy. Using 3 superspecies per class ranges from a compression factor of 4 (the aVOC and bVOC basis sets) to 6 (the POA basis set) with a significant improvement in accuracy from 2 superspecies and minimal improvement in accuracy when using 4 or more superspecies.

Improved accuracy with number of superspecies comes from the increased degrees of freedom, as each subsequent column of $\mathbf{W}$ adds another basis direction. Each column of $\mathbf{W}$, when normalized, can also be interpreted as a superspecies of unit concentration, and its elements corresponding to composition of VBS tracers. Each superspecies can also be interpreted as a different regime of organic aerosol, found through a data-driven method. Multiple superspecies can be combined in different amounts, corresponding to their concentrations, to form other distributions.

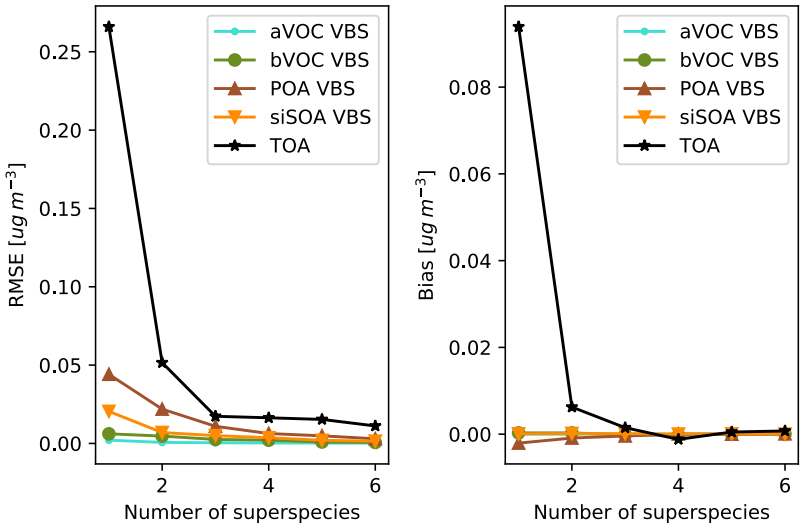

Figure 2. Relationship between the number of superspecies and the RMSE and bias for the 4 VBS classes, as well as TOA. There are diminishing returns in accuracy after 3 superspecies per VBS class.

3.2 Physical Consistency of Results

3.2.1 Motivating Non-negative Constraints

Section 2.4 raised the theoretical possibility of obtaining negative concentrations when using the pseudoinverse $\mathbf{W}^+$ to compress tracers into superspecies. Negative elements in $\mathbf{W}^+$ can lead to negative superspecies. Negative superspecies concentrations are not directly a problem, as the current advection scheme in LOTOS-EUROS v2.2.1 is based on that of Walcek (2000), which is able to handle negative tracer values. However, using the non-negative $\mathbf{W}$ to decompress negative superspecies concentrations back to the tracer space can lead to negative tracer values. Here, we quantify this limitation in practice using 3 superspecies.

Negative concentrations that are extremely small in magnitude can be approximated as zero. This tolerance can of course be set to a threshold, for example $-1 \times 10^{-8}\ \mu g\ m^{-3}$. However, using the test data of the POA VBS as an example, there are over 4.7 million cases in the test data where a POA VBS tracer is below $-1 \times 10^{-8}\ \mu g\ m^{-3}$, which is more than 19% of the 24 million values in the test data for the POA VBS.

One could choose a more relative, less arbitrary tolerance: for instance, all concentrations that are more negative than the magnitude of the corresponding bias for each VBS. These "significantly negative" concentrations would be negative even after an additive bias correction. For the POA VBS, there were 855,083 such concentrations, about 3.5% of the total test data. Using this relative tolerance, other VBS classes showed even larger proportions of "significantly negative" concentrations: 4.2%, 5.6%, and 7.0% respectively for the siSOA, aSOA, and bSOA VBS classes (for the anthropogenic VBS and siSOA VBS, which had positive biases, the tolerance was chosen to be the negative magnitude of the corresponding bias).

Using the pseudoinverse $\mathbf{W}^+$ for compressing VBS tracers can result in a number of significantly negative values when using 3 superspecies per VBS class, which motivates the development of non-negative compression strategies. For each VBS class, we find a positive compression matrix $\mathbf{B}$ to replace $\mathbf{W}^+$, according to the objective function and constraints in equation 2.

We compare this matrix factorization approach with a neural network autoencoder approach for each VBS class. We construct and train a 5-layer neural network autoencoder with rectified linear unit activation functions in the superspecies and output layers to ensure non-negativity of both superspecies and decompressed VBS tracers. In other hidden layers, a sigmoidal activation function, hyperbolic tangent, is used. In training, a dropout rate of 0.1 is used for every layer except for the superspecies layer. Figure 3.2.1 shows the general layer and node architecture of both anthropogenic and biogenic NN autoencoders: the autoencoders for the POA and siSOA VBS classes vary only in the dimensionality of their input and output layers (18 and 16 values, as opposed to 12). For the autoencoder of each VBS class, the center superspecies layer is chosen to have 3 values: the value of this hyperparameter is chosen for comparison to the linear matrix factorization approach. Both the matrix factorization and autoencoder approaches are evaluated on how well they reconstruct the tracers after decompression.

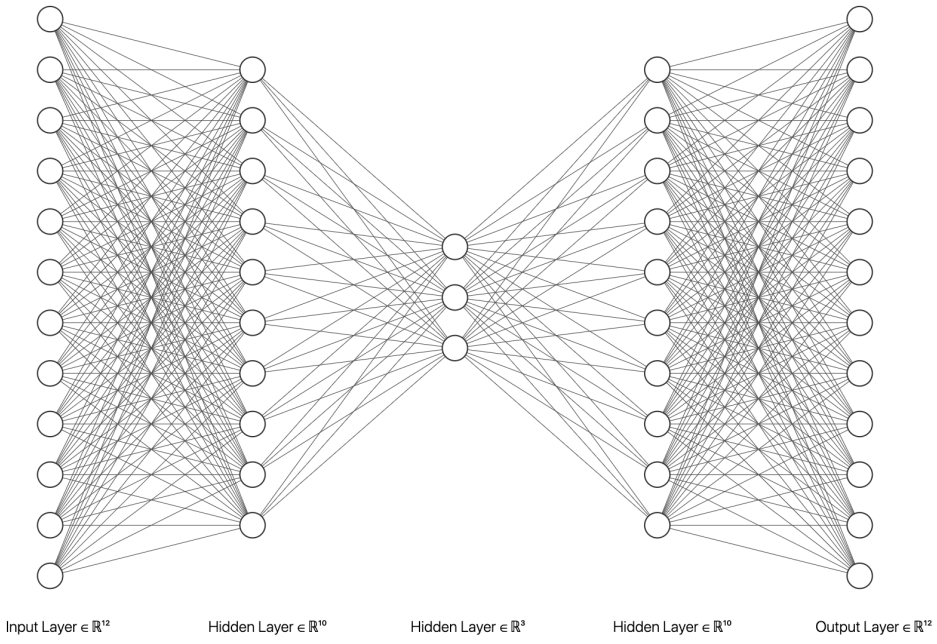

Figure 3. Neural network autoencoder architecture for the anthropogenic and biogenic VBS classes, compressing 12 tracers to 3 superspecies and back again. The central layer of 3 superspecies uses a ReLu activaction function, as well as the output layer, to ensure semi-positivity. The autoencoders for the other VBS classes have analogous architectures, differing only in the size of input and output layers: 18 nodes for POA and 16 nodes for siSOA. This figure was generated using the tool from LeNail (2020).

3.2.2 Conserving Mass and Phase

Section 2.6 proposed a method for conserving total concentration of the VBS tracers in both the superspecies representation and in subsequent reconstruction to decompressed tracers. Applying this method to the cross-sections of VBS class and phase (particle or gas) ensures that the superspecies transformation does not add or remove mass artificially in the gas and particle phases of every class: this necessitates conservation of total gas concentration, total aerosol concentration, and concentration of total organic material (TOM). Phase-specific superspecies are composed of entirely gas or entirely aerosol tracers, conserving information on phase while in the latent space representation.

Phase-specific superspecies require adding another cross-section, halving the number of tracers to be compressed and decompressed by each pair of $\mathbf{B}$ and $\mathbf{W}$, respectively. For this reason, continuing to use 3 superspecies for each phase within each VBS class would reduce the compression factor to slightly over 2.4, not much better than compression factor of around 2 when using the non-lossy partitioning approach. However, using only 1 superspecies per phase per class would fix each corresponding set of tracers to a single shape upon reconstruction, analogously to the discussion in section 3.1. To ensure that this method captures spatiotemporal variability of volatility distributions, while maintaining a decent compression factor, we choose to use 2 superspecies per phase per VBS class. This design choice results in a compression factor of approximately 3.6. Its accuracy is compared to the other strategies in the model selection process in section 3.3.

Figure 4 demonstrates the mass-conserving properties of this approach using representative examples of the primary organic aerosol VBS distribution at two different atmospheric monitoring sites: the Cabauw Experimental Site for Atmospheric Research in the Netherlands, and Mace Head Atmospheric Research Station in Ireland. Mace Head is as a more pristine and remote station (O'Dowd et al., 2014). The legend in Figure 4 shows that POA concentration at Cabauw is two orders of magnitude higher than that at Mace Head, 4.984 $\mu g\ m^{-3}$ compared to 0.032 $\mu g\ m^{-3}$.

Figure 4 compares the primary VBS distribution to the reconstructed primary VBS distributions after mapping to phase-specific superspecies and back again using two sites: Cabauw and Mace Head, as representative examples. Comparing the legends of (a) with (c), it can be seen that total POA concentration, as well as total concentration of tracers in the gas phase, is conserved to machine precision after passing through compression. The same holds for the total concentrations at Mace Head, (b) and (d), at orders of magnitude more dilute. With phase information and concentration conserved, the only source of error caused by compression to superspecies is in the shape of the distribution. This reconstruction error is more apparent at Mace Head in Figure 4 (b) and (d).

3.3 Model Selection

In this section, we compare the four approaches described thus far, and make a judgment about the most promising strategy, evaluated on reconstruction accuracy and physical consistency. The selected approach will be implemented in LOTOS-EUROS v2.2.1 to accelerate the advection operator.

1. NMF/Pseudoinverse linear approach: NMF to find an optimal decompression matrix $\mathbf{W}$, and use its pseudoinverse (with negative elements) $\mathbf{W}^{+}$ as a compression matrix using 3 superspecies per VBS class

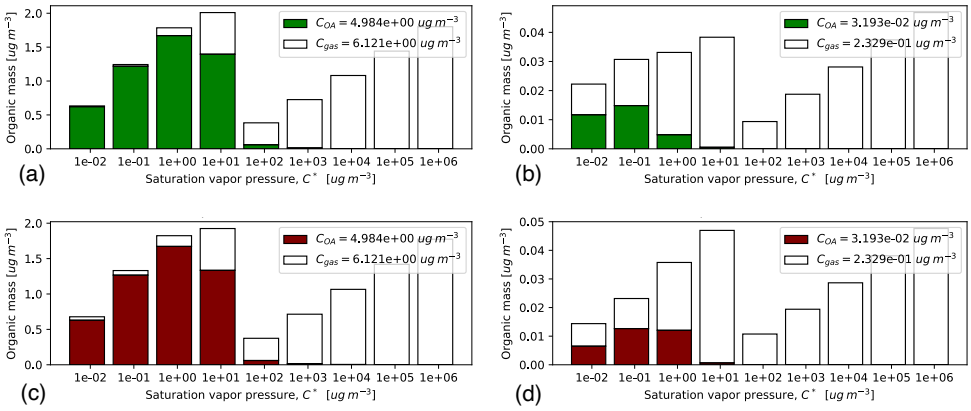

Figure 4. Comparison of the volatility basis set distribution for POA near two sites: Cabauw and Mace Head at a snapshot in time on February 26, 2018. The top row in green shows the distributions as modeled by LOTOS-EUROS at Cabauw (a) and (b) Mace Head. The bottom row in maroon shows the distributions at Cabauw (c) and Mace Head (d) after the non-negative compression/decompression using phase-specific superspecies. Total concentrations are conserved when comparing the legends of the modeled distributions to the reconstructed distributions.

2. Non-negative linear compression: NMF to find an optimal decompression matrix **W**, as well as a non-negative compression matrix **B** using 3 superspecies per VBS class
3. Non-negative neural network autoencoder: Create a more complicated neural network with ReLU activation functions in the superspecies and output layers, using 3 superspecies per VBS class
4. Mass-conserving, non-negative linear compression with phase specific superspecies: Create **W**, as well as a non-negative compression matrix **B** using 2 superspecies per phase per VBS class

Tables 2 and 3 show RMSE and Bias of the tracers for each VBS class for the 4 approaches, as well as total organic aerosol (TOA) and total organic material (TOM) concentrations.

Table 2. Evaluation RMSE of selected approaches. All values reported in $\mu g\, m^{-3}$.

-	Approach 1	Approach 2	Approach 3	Approach 4
aVOC VBS	4.4×10^{-4}	0.0010	0.0021	0.0011
bVOC VBS	0.0026	0.0078	0.0181	0.0042
POA	0.0109	0.0285	0.0306	0.0142
siSOA	0.0050	0.0086	0.0094	0.0057
TOA	0.0173	0.133	0.101	6.9×10^{-13}
TOM	0.0547	0.240	0.328	1.0×10^{-12}

Using non-negative **B** and **W** to linearly combine tracers into three superspecies shows lower RMSE values than the NN autoencoder in Approach 3,

Table 3. Evaluation bias of selected approaches. All values reported in $\mu g\, m^{-3}$.

	Approach 1	Approach 2	Approach 3	Approach 4
aVOC VBS	2.6×10^{-5}	1.2×10^{-4}	-3.9×10^{-4}	2.8×10^{-20}
bVOC VBS	-1.6×10^{-4}	3.8×10^{-4}	-0.0051	-1.6×10^{-16}
POA	-4.2×10^{-4}	0.0050	-0.0075	-8.8×10^{-18}
siSOA	-9.9×10^{-5}	7.7×10^{-4}	-0.0022	1.2×10^{-19}
TOA	0.0015	0.0657	-0.0346	-1.3×10^{-15}
TOM	-0.00763	0.108	-0.237	-2.1×10^{-15}

with the exception of TOA concentration. This indicates that a linear approach is probably suitable for VBS tracer compression. Using the pseudoinverse $\mathbf{W}^{+}$for compression resulted in much lower RMSE for all the VBS classes, but has the critical weakness of producing a significant amount of negative concentrations for superspecies and subsequently reconstructed tracers as explored in Section 3.2.1. Though the phase-specific superspecies approach does not have as low of RMSE for each VBS class as the pseudoinverse approach, it outperforms the other two non-negative approaches. Moreover, it conserves absolute metrics on compression, ensuring that material will stay in each class and each phase, and no material will be added or removed by compression: for this reason, all biases are negligible to machine precision. Preserving information on phase during compression to superspecies has another advantage. This approach can be used in other processes such as dry deposition, which handles particle and gas tracers separately and is a computationally expensive operator in LOTOS-EUROS. Because the phase-specific superspecies method (Approach 4) is physically consistent while quite accurate in reconstruction error, and is readily extended to other phase-specific processes, it is chosen for implementation in LOTOS-EUROS v2.2.1.

4 Results: Superspecies Implementation in LOTOS-EUROS

The phase-specific, matrix factorization superspecies method (Approach 4) chosen in section 3.3 was implemented in LOTOS-EUROS v2.2.1 as described in section 2.7. This section explores the accuracy and speedup of replacing VBS tracers with ML superspecies in advection, as well as the generalizability of the superspecies to different seasonal conditions and spatial resolutions.

After offline training on data from February 20th through 24th, 2018, the selected superspecies parameterization was implemented into LOTOS-EUROS, and used in the advection operator for a run from February 15th through 28th. This is compared with a control run advecting VBS tracers as a baseline to directly assess the error from advecting superspecies. With the ML parameterization implemented in LOTOS-EUROS, small errors caused by advecting superspecies change subsequent VBS tracer concentrations such that the period of February 20th through 24th differs from the training dataset. In that time period, however, meteorological conditions and other processes independent of the VBS and superspecies parameterization are identical to that of the offline training dataset. For the sake of comparison, the superspecies run and control run are evaluated on February 25th through 28th, even though the superspecies run has the chance to accumulate error and diverge from the control run from the beginning of the simulation on February 15th.

Advecting ML superspecies reproduces the spatial patterns of average TOA across the entire domain. Figure 5 shows average TOA of the control run and the

superspecies run, from February 25th through February 28th. This test time period is well into the model run, 10 days after the begin of the simulation. During this time period and over the entire domain, average bias of TOA of the superspecies run compared to the control run is small and slightly negative, -0.0095 $\mu g\ m^{-3}$. Small average bias is not in itself indicative of low error, as positive and negative bias cancellations throughout the domain and time period are possible. RMSE, an absolute metric, was larger at 0.217 $\mu g\ m^{-3}$. However, the general spatial patterns of total OA across the entire domain are preserved when advecting superspecies.

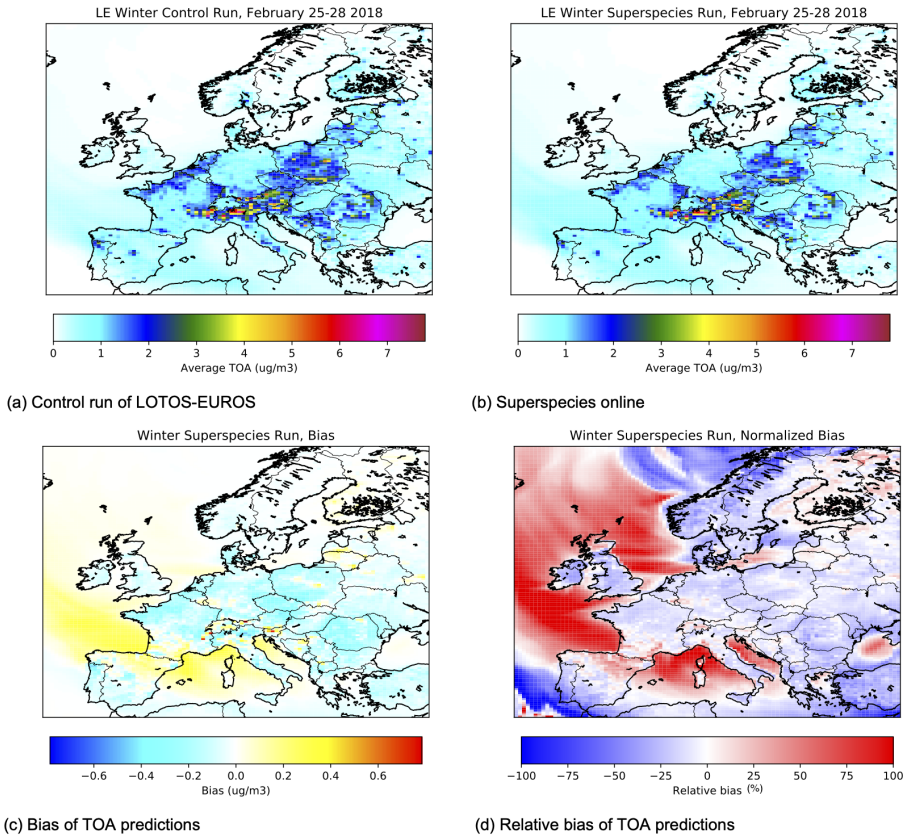

Figure 5. Average TOA for February 25th through 28th 2018, during a 2 week simulation from February 15th through 28th using superspecies matrices optimized offline on winter conditions from February 20th through 24th.

4.1 Seasonal Superspecies

Advecting superspecies was shown to reproduce the spatial patterns of OA in a winter test period from February 25th through 28th. This test period occurred directly after the training test period February 20 through 24th, and has relatively similar conditions to what the superspecies transformation matrices were optimized for.

A run in summer from July 20th through August 1st was chosen to assess the robustness of the winter-optimized superspecies. Summer conditions differ from winter conditions in Europe for several reasons. One, biogenic precursor gases make up a larger contribution to formation of secondary organic aerosol in the summer, partially due to emissions from forests. Two, average temperatures are higher, affecting the partitioning of the VBS by changing the volatility basis set values C^*. The different conditions lead to different modeled compositions of total organic aerosol (TOA). Table 4 compares the modeled average composition of OA for February 25th through 28th to that for July 29th through August 1st.

Table 4. Average TOA composition for the LOTOS-EUROS runs for February and July.

OA Type	February	July
aSOA	0.8%	9.5%
bSOA	4.5%	34.8%
POA	61.2%	12.5%
siSOA	33.5%	43.2%

Though siSOA is on average the largest component of TOA in the run from July 29th through August 1st this is not the full picture, and underscores the importance of bSOA under some conditions. The maximum concentration of surface siSOA over the entire domain over the entire period from July 29th through August 1st was 15.0 $\mu g \quad m^{-3}$, and 99th percentile 1.3 $\mu g \quad m^{-3}$ compared to the maximum bSOA concentration of 100.3 $\mu g\ m^{-3}$ and 99th percentile 9.4 $\mu g\ m^{-3}$. This indicates that although siSOA may dominate in background conditions and when TOA is low, bSOA is the dominant component of TOA in other conditions.

4.1.1 Domain-wide assessment

Figure 6 shows average surface TOA, as predicted by the control run (a), the run with superspecies advected (b), and the bias and relative bias of the superspecies run with regards to the control, (c) and (d) respectively. The spatial patterns of TOA are visually different from the winter conditions in Figure 5. Primary organic emissions corresponding to POA are often the largest contributor to winter TOA, and for this time period TOA is most concentrated in the Po Valley, Czechia, and Poland. The winter superspecies run is able to recreate these large regions of high TOA, as well as other smaller but distinct pockets of TOA, such as Madrid (the most populous city in Spain) and northwest Portugal, a region with a lot of industrial activity. In contrast, summer TOA is concentrated around southern Germany, Switzerland, Austria, and Slovenia. Many places in this region are forested, and contribute to TOA via emission of biogenic precursors of bSOA. The superspecies run shown in (b) is able to capture these spatial patterns, but with a strong bias. For this reason, other regions with high biogenic emissions become visually apparent in (b), such as southern Sweden, Finland Proper, and northwestern Russia, which are all heavily forested and which is taken into account in LOTOS-EUROS via land use maps and tree-species emissions. (Manders et al., 2017).

The superspecies, optimized on winter conditions and tested on a 2 week run in July, show a large positive bias over the areas with high average TOA, especially heavily forested regions. RMSE for TOA over the whole domain and time period is 2.12 $\mu g\ m^{-3}$, with an average bias of 0.321 $\mu g\ m^{-3}$. RMSE of the tracers from the biogenic VBS for all times and gridcells is 0.66 $\mu g\ m^{-3}$, an order of magnitude

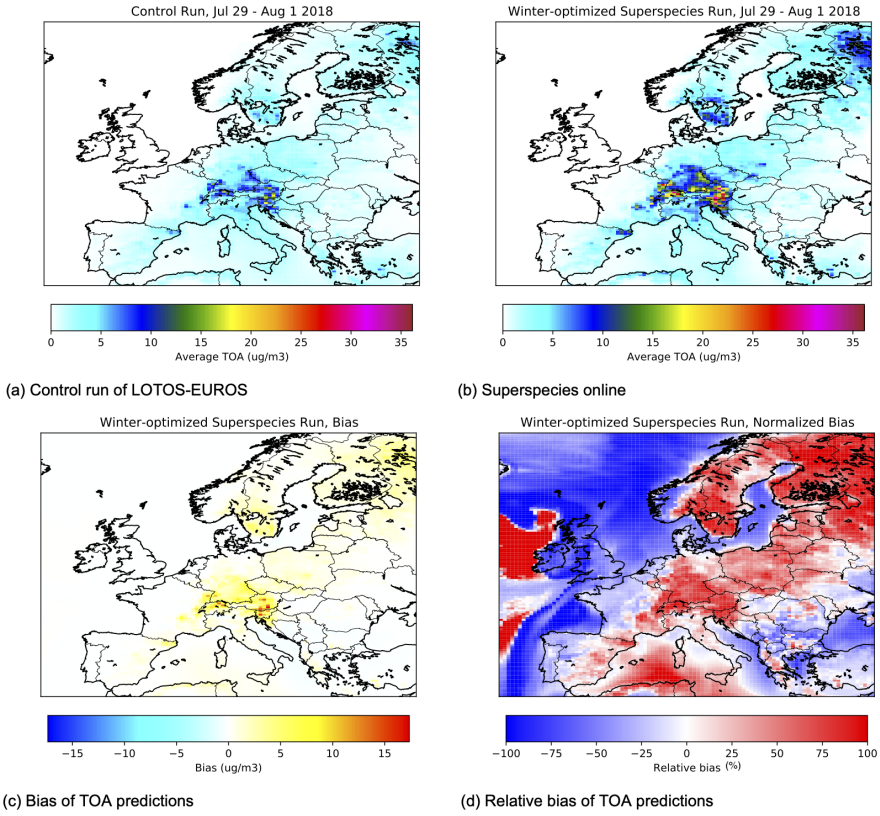

Figure 6. Average TOA for July 29th through August 1st 2018, during a 2 week simulation from July 19th through August 1st using superspecies matrices optimized offline on winter conditions from February 20th through 24th.

higher than tracers from the other VBS classes: the class of tracers with the next highest RMSE value is the siSOA VBS class, at 0.062 $\mu g\ m^{-3}$. The average bSOA bias (bias of total biogenic aerosol neglecting gaseous tracers) is 0.068 $\mu g\ m^{-3}$, three orders of magnitude smaller than the maximum bSOA bias of 82.9 $\mu g\ m^{-3}$. Overestimation of bSOA in the superspecies run under some conditions is likely due to errors in decompression, artificially shifting mass to lower volatility bins. However, the large positive bias in parts of the domain indicate that this tendency to overestimate bSOA only happens in certain conditions: namely, forested regions. The following section analyzes one gridcell in a forested region, and finds additional temporal patterns where bSOA is significantly overestimated, leading to overestimation of TOA.

4.1.2 Case Study: Summer Night in a Forest

We choose a single gridcell over a forested area to take a closer look at the superspecies tendency to lead to high bSOA. We study the LOTOS-EUROS gridcell containing the Schönbuch Natural Reserve in southwest Germany, which is 156 square kilometers and 85% forested. Figure 7a shows the temporal variation of TOA

649 in the Schönbuch from July 29th through August 1st. This overestimation
650 systematically occurs at night, with the night of July 30th to July 31st a
651 particularly high TOA event showing the highest bias.

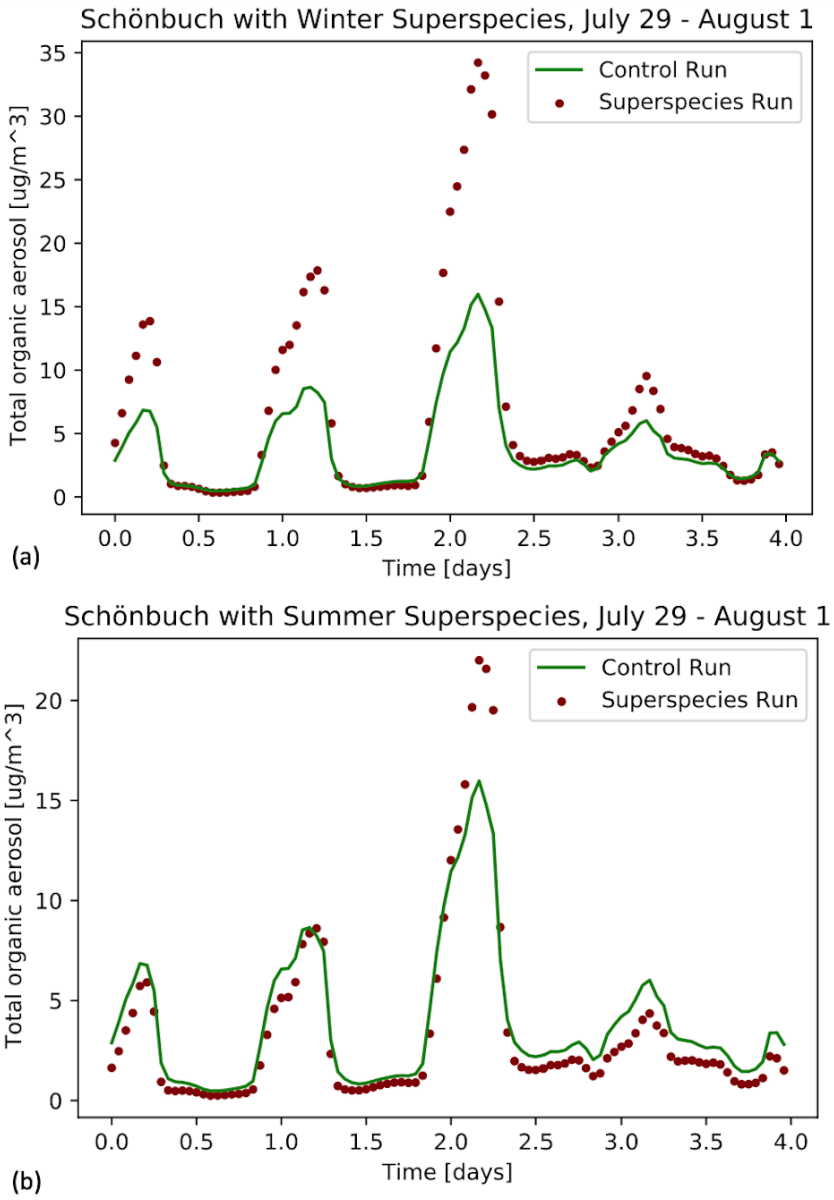

Figure 7. Temporal variation of TOA over Schönbuch from July 29th through August 1st using (a) winter-optimized superspecies and (b) summer-optimized species. The maroon points of TOA as predicted with when advecting superspecies are compared to the green line of TOA as modeled by the LE control run used as a baseline.

Examining Figure 7a, the peak overestimation occurs on 05:00 July 31st and overestimates total bSOA with a factor between 2 and 2.5 times that of the control run. The superspecies run has a bSOA concentration of 32.9 $\mu g\ m^{-3}$, which comprises 99% of total OA concentration for that gridcell and time. The control run concentration of bSOA is 14.1 $\mu g\ m^{-3}$, about 95% of TOA for that simulation. By 09:00 July 31st, both runs return to a total bSOA concentration of less than 3.5 $\mu g\ m^{-3}$. This night episode of high bSOA contains the largest overpredictions for that particular gridcell in the whole time period. However, it is illustrative of a failure mode of the winter-optimized superspecies to capture the total concentration of bSOA, and ultimately TOA due to the importance of bSOA contributions in this example. The spatial patterns and temporal patterns of the superspecies run compared to the control run show that the superspecies are limited in their ability to model conditions over forested areas on summer nights.

Given that winter-optimized superspecies showed limitations in capturing high bSOA events over forested areas at night, we investigate whether superspecies optimized on summer conditions and implemented online reproduce high bSOA conditions with more accuracy. The phase and mass conserving method selected in Section 3 is applied to model output from July 23rd through 28th, 2018, to obtain a superspecies parameterization optimized on summer conditions. The loss of accuracy is evaluated for the subsequent 4 days, from July 29th to August 1st, 2018.

The superspecies approach optimized on summer conditions shows a much lower bias than the winter-optimized superspecies. The temporal behavior of summer-optimized superspecies from July 29th through August 1st is shown in Figure 7b. Comparing Figure 7a to 7b, it can be seen that the spatiotemporal pattern of bSOA bias is fixed by using summer-optimized superspecies, which do not show the same nightly overestimation pattern of winter-optimized superspecies. Total bSOA is even slightly underestimated in the day when using summer-optimized superspecies.

Averaged over the entire domain and time period of July 29th through August 1st, the summer-optimized superspecies display a slightly negative average bias for bSOA of -0.023 $\mu g\ m^{-3}$. Small pockets of TOA overestimation (within 10 $\mu g\ m^{-3}$) still occur in the same regions as the winter-optimized superspecies: over highly forested areas. Quantitatively, the superspecies optimized on summer conditions result in more accurate runs in the summer. The RMSE over the whole domain of of time-averaged TOA was 0.98 $\mu g\ m^{-3}$ when using summer-optimized superspecies, less than half of the RMSE of 2.12 $\mu g\ m^{-3}$ when using winter-optimized superspecies. RMSE of the tracers from the biogenic VBS (both gas and particle phases) for all times and gridcells is reduced by a factor of 2, at 0.32 $\mu g\ m^{-3}$ compared to 0.66 $\mu g\ m^{-3}$.

The limitation of winter-optimized superspecies and the subsequent improvement in accuracy when using summer-optimized superspecies indicates that this method might be best applied to different seasons: creating seasonal-specific superspecies might result in higher accuracy. However, in superspecies trained on either conditions, the biogenic VBS tracers in the summer show significantly higher error than the tracers of the other VBS classes, with the siSOA VBS class having the next highest RMSE value at 0.050 $\mu g\ m^{-3}$.

4.2 Towards Operational Forecasting on Higher-Resolution Domains

LOTOS-EUROS is one model in the ensemble used in the Copernicus Atmospheric Modeling Service (CAMS) operational forecasts, which requires all models to include secondary organic aerosol (SOA) representation by 2022. The domain used in CAMS operational forecasts has a higher resolution and wider

domain than the domain used by MACC: 0.1 °by 0.1° for 420 by 700 gridcells compared to the 0.50° by 0.25° used in the MACC domain, and extending past Moscow, Russia. The change of resolution and domain increases the number of gridcells by a factor of 20. One result of this is many more gridcells and many more computations. Another result is that the operator splitting timestep Δt needs to decrease in order to satisfy the Courant-Friedrichs-Lewy criterion as the gridcell distance is smaller. With a smaller operator splitting timestep, the call to the advection operator occurs more times per simulated hour. With every call to the advection operator, compression of tracers into superspecies and decompression of superspecies back into tracers also occurs.

We investigate how the superspecies approach, optimized on model output from February 20th through 24th on the coarse-resolution MACC domain, generalizes to a 2 week run on the extended high-resolution CAMS domain. Figure 8 shows the time-averaged TOA concentration across the entire CAMS domain for the test period of February 25th-28th, 2018, chosen for ease of comparison with the winter run on the MACC domain.

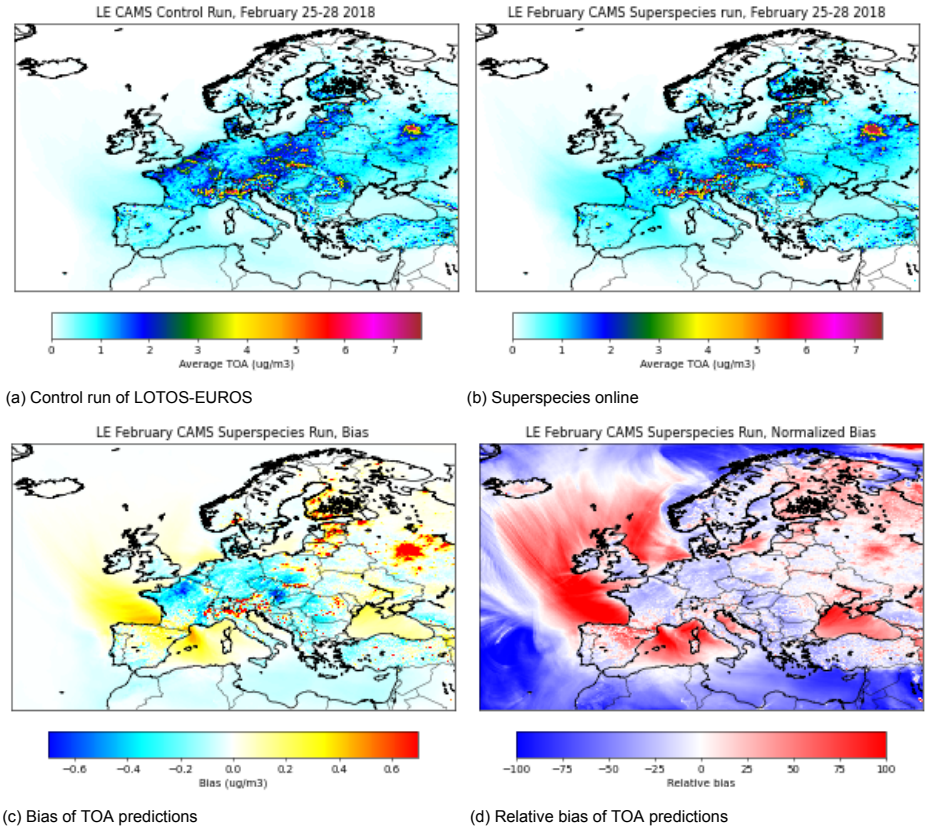

Figure 8. Time averaged TOA for the period of February 25th through 28th on the high-resolution domain used in CAMS operational forecasting, from control and superspecies runs, as well as bias and relative bias. The superspecies were optimized on model output from a simulation using the coarse-resolution MACC domain.

The superspecies run has a positive bias for TOA of 0.019 $\mu g\ m^{-3}$, with visible overestimation in the area near Moscow, Russia, which is not in the MACC grid used to optimize the compression/decompression matrices.

The colorbar limits of Figure 8 (a), (b), and (c) were adjusted for visual comparison with Figure 5. For this reason, colors at the upper or lower limits should be interpreted as greater or equal to the limit. Though the maximum gridcell concentration of time-averaged TOA from both the superspecies run and the control run was 28.2 $\mu g\ m^{-3}$, 99.85% of the gridcells had a time-averaged TOA under 7.6 $\mu g\ m^{-3}$, which was chosen as the upper limit of the colorbar. This means that only 0.15% of the gridcells in Figures 8a and 8b exceed the limit shown in the colorbar. Neglecting the highest 0.15% of average TOA, the spatial patterns of the CAMS control run in Figure 8a are visually very similar to those of the that the CAMS superspecies run in Figure 8b. Both show spatial patterns similar to the simulations performed on the MACC grid for the same time period. The same approach is done for the bias shown in Figure 8c, with very few gridcells in the CAMS simulation exceeding the maximum error of time-averaged TOA on the MACC grid. The maximum absolute error of time-averaged TOA between the superspecies run and the control run was 8.9 $\mu g\ m^{-3}$, but 99.2% of all gridcells had an absolute error of less than 0.70 $\mu g\ m^{-3}$. Less than 1% of the gridcells in Figure 8c exceed the colorbar limit.

The largest bias for TOA over all cells and the entire test time period (not time averaged) was 89 $\mu g\ m^{-3}$, corresponding to a gridcell in northwestern Spain, near Ponferrada. This gridcell also showed the highest time-averaged TOA concentration of 32.0 $\mu g\ m^{-3}$ for the superspecies run, compared to 19.4 $\mu g\ m^{-3}$ for the control run.

At the highest positive bias of 89 $\mu g\ m^{-3}$, TOA concentration as modeled by the superspecies run is 206.4 $\mu g\ m^{-3}$ while the control run predicts a TOA concentration of 117.4 $\mu g\ m^{-3}$. TOA during this event is composed almost wholly of primary material: the superspecies run models a POA concentration of 205.9 $\mu g\ m^{-3}$ (99.78% of TOA concentration) while the control run POA concentration is 117.1 $\mu g\ m^{-3}$ (99.75 %). Rather than error compounding and leading to divergence from the control run, the superspecies run restabilizes without error accumulation for the rest of the simulation: TOA concentration in the superspecies run converges to that of to the control run. This indicates that other processes in LOTOS-EUROS can correct temporary overestimation arising from the superspecies parameterization as the simulation progresses.

4.3 Speed Improvement

The advection operator has an outer for-loop over all tracers that are transported. Using superspecies instead of VBS tracers reduces the overall amount of variables in the outer for-loop. With the superspecies selected in Section 3, 16 superspecies (two gas and two particle superspecies for each of the four VBS classes) are advected rather than the 58 VBS tracers.

The MACC run on the small domain was run sequentially on one node. Wall time for the advection operator when advecting superspecies rather than VBS tracers was 6790 seconds, 56% of the time of (1.8 times faster than) the 12073 seconds to advect all tracers in the control run.

The high resolution required for CAMS operational forecasts increases the computational intensity of the simulations. The CAMS runs for both the control and superspecies runs were performed using domain decomposition over 24 computing nodes with each node computing a subdomain of 175 by 70 gridcells.

Using the VBS on the CAMS domain, advection wall time more than doubled from 34959 seconds to 74762 seconds. With superspecies advected instead of VBS tracers, wall time for the advection operator was then reduced to 49473 seconds. Advecting superspecies on the CAMS domain took about 66% of the time that advecting all the VBS tracers took, a speedup of approximately 1.5.

The timing results suggest that advection wall time depends linearly on number of tracers, which is expected behavior given the structure of the advection operator: an outer for-loop over all tracers. VBS inclusion approximately doubles the tracers and therefore should be expected to double the computation time of the advection operator. The set of 16 superspecies is slightly more than a quarter the size of the 58 VBS tracers. The VBS tracers are slightly more than half of the 104 advected tracers (some of the other 64 tracers, like radicals, are not advected), so the theoretical estimate of timing is 59% of the original timing when using 16 superspecies instead of 58 VBS tracers.

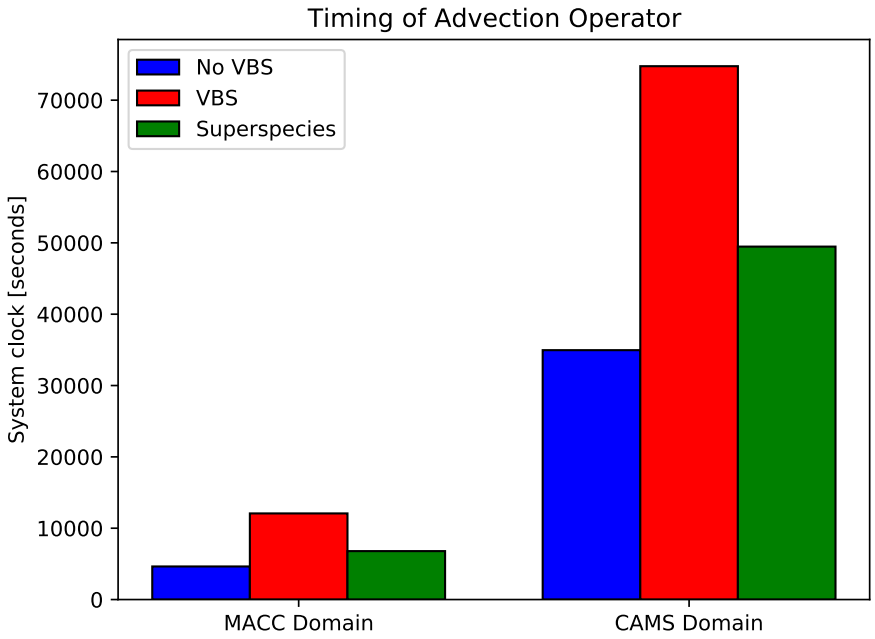

Figure 9. Use of the 58 VBS tracers approximately doubles the wall time spent on advection calculations. Advecting superspecies takes 56% and 66% of the time compared to advecting VBS tracers on the MACC and CAMS domains, respectively.

5 Conclusions

The high-dimensional modeling of organic aerosol processes via four VBS classes is computationally expensive in LOTOS-EUROS v2.2.1, slowing the advection operator down by a factor of 2. We developed unsupervised machine learning (ML) methods to reduce the dimension of VBS tracers to a set of superspecies and reduce the computational burden on the advection operator. These methods were refined to ensure physical consistency, including semi-positive constraints, mass conservation, and information on phase. Multiple approaches were compared in Section 3 and a linear approach based on non-negative matrix factorization was judged to be the most appropriate after being evaluated on reconstruction accuracy and physical consistency. This approach creates 16 phase-specific, class-specific superspecies (a compression factor of 3.6) while preserving phase and conserving total concentration to machine precision. Section 4 explores the results of implementing the ML superspecies in LOTOS-EUROS v2.2.1. The ML parameterization ran stably without runaway error for a model simulation of 2 weeks. Higher bias of total OA concentration was shown when the superspecies, optimized to reconstruct winter OA patterns, were used in a 2 week run in the summer. During the summer run, the bias showed a clear spatiotemporal pattern, with biogenic SOA overestimated over forests at night. The ML approach was retrained on model output from summer conditions and implemented in LOTOS-EUROS v2.2.1 to reduce high bias. The resuts of this case study indicate that the superspecies might work best when optimized for season-specific conditions.

We found that the ML superspecies trained on the coarse-resolution MACC domain performed well when used on the fine-resolution domain used in CAMS operational forecasts for a period of 2 weeks. In an analysis period of 4 days performed at the end of the 2 week CAMS run, over 99% of all gridcells showed an absolute bias of time-averaged TOA within the maximum error of the MACC grid. Evaluating a gridcell that exceeded the maximum average error, we found that high overestimation of total OA concentration occured at a high OA event, and converged back to the baseline simulation as time progressed rather than displaying continued error growth.

Advecting superspecies reduced the wall time spent on the advection operator: advecting superspecies took 56% to 66% of the time that it took to advect VBS tracers. Timing experiments indicate a linear dependence of wall time on number of tracers to advect, an expected relation from the structure of the advection operator, which uses a for-loop over all advected tracers. With linear dependence demonstrated, the design choice of compression factor (number of superspecies) can already give an estimate of theoretical speedup.

The use of physically consistent machine learning to find superspecies allows for inclusion of organic aerosol processes without doubling the computational burden on the advection operator. Preserving information on phase of the superspecies allows for their future use in phase-specific processes such as dry deposition, which can be computationally intensive in LOTOS-EUROS. Though demonstrated on organic aerosol species in a regional CTM as a case study, this approach readily generalizes to other tracers, processes, and models. As physical consistency and computational efficiency are widely desired aspects of numerical modeling in the physical sciences, this approach could be adapted for use in comprehensive earth system models with the purpose of providing forecasts of global atmospheric composition, for example GEOS-CF (Keller et al., 2021). More generally, this method contributes additional physical consistency to a widely used linear dimensionality reduction technique (non-negative matrix factorization) that can be used to reversibly map between high and low detail in earth system models.

Acknowledgments

This work was supported in part by the UC Davis CeDAR (Center for Data Science and Artificial Intelligence Research) Innovative Data Science Seed Funding Program. We thank Christoph Keller (NASA Goddard Space Flight Center) for valuable discussion throughout the project.

The exact version of LOTOS-EUROS used in this work including the superspecies extension, as well as all Python code used for analysis of model output and figure generation, is available at https://doi.org/10.5281/zenodo.6601166.

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., ... Zheng, X. (2016). Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)* (pp. 265–283).

Bergström, R., Denier van der Gon, H. A. C., Prévôt, A. S. H., Yttri, K. E., & Simpson, D. (2012). Modelling of organic aerosols over europe (2002-2007) using a volatility basis set (vbs) framework: application of different assumptions regarding the formation of secondary organic aerosol. *Atmospheric Chemistry and Physics*, *12*(18), 8499–8527. Retrieved from `https://acp.copernicus.org/articles/12/8499/2012/` doi: 10.5194/acp-12-8499-2012

Beucler, T., Pritchard, M., Rasp, S., Ott, J., Baldi, P., & Gentine, P. (2021). Enforcing analytic constraints in neural networks emulating physical systems. *Physical Review Letters*, *126*(9), 098302.

Brasseur, G. P., & Jacob, D. J. (2017). *Modeling of atmospheric chemistry.* Cambridge University Press. doi: 10.1017/9781316544754

Brenowitz, N. D., & Bretherton, C. S. (2019). Spatially extended tests of a neural network parametrization trained by coarse-graining. *Journal of Advances in Modeling Earth Systems*, *11*(8), 2728-2744. Retrieved from `https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019MS001711` doi: https://doi.org/10.1029/2019MS001711

Chollet, F., et al. (2015). *keras.*

de Gouw, J. A., Middlebrook, A. M., Warneke, C., Goldan, P. D., Kuster, W. C., Roberts, J. M., ... Bates, T. S. (2005). Budget of organic carbon in a polluted atmosphere: Results from the new england air quality study in 2002. *Journal of Geophysical Research: Atmospheres*, *110*(D16). Retrieved from `https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2004JD005623` doi: https://doi.org/10.1029/2004JD005623

Denier van der Gon, H. A. C., Bergström, R., Fountoukis, C., Johansson, C., Pandis, S. N., Simpson, D., & Visschedijk, A. J. H. (2015). Particulate emissions from residential wood combustion in europe – revised estimates and an evaluation. *Atmospheric Chemistry and Physics*, *15*(11), 6503–6519. Retrieved from `https://acp.copernicus.org/articles/15/6503/2015/` doi: 10.5194/acp-15-6503-2015

Donahue, N. M., Epstein, S. A., Pandis, S. N., & Robinson, A. L. (2011). A two-dimensional volatility basis set: 1. organic-aerosol mixing thermodynamics. *Atmospheric Chemistry and Physics*, *11*(7), 3303–3318. Retrieved from `https://acp.copernicus.org/articles/11/3303/2011/` doi: 10.5194/acp-11-3303-2011

Donahue, N. M., Henry, K. M., Mentel, T. F., Kiendler-Scharr, A., Spindler, C., Bohn, B., ... Baltensperger, U. (2012). Aging of biogenic secondary organic aerosol via gas-phase oh radical reactions. *Proceedings of the National Academy of Sciences*, *109*(34), 13503-13508. Retrieved from `https://www.pnas.org/doi/abs/10.1073/pnas.1115186109` doi:

10.1073/pnas.1115186109

Donahue, N. M., Robinson, A., Stanier, C., & Pandis, S. (2006). Coupled partitioning, dilution, and chemical aging of semivolatile organics. *Environmental science & technology*, *40*(8), 2635–2643.

EEA. (2005). *Image2000 and clc2000. products and methods. corine land cover updating for the year 2000.* Ispra Italy.

Gery, M. W., Whitten, G. Z., Killus, J. P., & Dodge, M. C. (1989). A photochemical kinetics mechanism for urban and regional scale computer modeling. *Journal of Geophysical Research*, *94*(D10), 12925. Retrieved from http://doi.wiley.com/10.1029/JD094iD10p12925 doi: 10.1029/JD094iD10p12925

Hayes, P. L., Carlton, A. G., Baker, K. R., Ahmadov, R., Washenfelder, R. A., Alvarez, S., ... Jimenez, J. L. (2015). Modeling the formation and aging of secondary organic aerosols in los angeles during calnex 2010. *Atmospheric Chemistry and Physics*, *15*(10), 5773–5801. Retrieved from https://acp.copernicus.org/articles/15/5773/2015/ doi: 10.5194/acp-15-5773-2015

Heald, C. L., Jacob, D. J., Park, R. J., Russell, L. M., Huebert, B. J., Seinfeld, J. H., ... Weber, R. J. (2005). A large organic aerosol source in the free troposphere missing from current models. *Geophysical Research Letters*, *32*(18). Retrieved from https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2005GL023831 doi: https://doi.org/10.1029/2005GL023831

Janssen, R. H. H., Tsimpidi, A. P., Karydis, V. A., Pozzer, A., Lelieveld, J., Crippa, M., ... Locoge, N. (2017). Influence of local production and vertical transport on the organic aerosol budget over paris. *Journal of Geophysical Research: Atmospheres*, *122*(15), 8276-8296. Retrieved from https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2016JD026402 doi: https://doi.org/10.1002/2016JD026402

Jiang, J., Aksoyoglu, S., El-Haddad, I., Ciarelli, G., Denier van der Gon, H. A. C., Canonaco, F., ... Prévôt, A. S. H. (2019). Sources of organic aerosols in europe: a modeling study using camx with modified volatility basis set scheme. *Atmospheric Chemistry and Physics*, *19*(24), 15247–15270. Retrieved from https://acp.copernicus.org/articles/19/15247/2019/ doi: 10.5194/acp-19-15247-2019

Jimenez, J. L., Canagaratna, M. R., Donahue, N. M., Prevot, A. S. H., Zhang, Q., Kroll, J. H., ... Worsnop, D. R. (2009). Evolution of organic aerosols in the atmosphere. *Science*, *326*(5959), 1525-1529. Retrieved from https://www.science.org/doi/abs/10.1126/science.1180353 doi: 10.1126/science.1180353

Kaiser, J. W., Heil, A., Andreae, M. O., Benedetti, A., Chubarova, N., Jones, L., ... van der Werf, G. R. (2012). Biomass burning emissions estimated with a global fire assimilation system based on observed fire radiative power. *Biogeosciences*, *9*(1), 527–554. Retrieved from http://www.biogeosciences.net/9/527/2012/bg-9-527-2012-discussion.html doi: 10.5194/bg-9-527-2012

Keller, C. A., & Evans, M. J. (2019). Application of random forest regression to the calculation of gas-phase chemistry within the geos-chem chemistry model v10. *Geoscientific Model Development*, *12*(3), 1209–1225. Retrieved from https://gmd.copernicus.org/articles/12/1209/2019/ doi: 10.5194/gmd-12-1209-2019

Keller, C. A., Knowland, K. E., Duncan, B. N., Liu, J., Anderson, D. C., Das, S., ... Pawson, S. (2021). Description of the nasa geos composition forecast modeling system geos-cf v1.0. *Journal of Advances in Modeling Earth Systems*, *13*(4), e2020MS002413. Retrieved from https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2020MS002413 (e2020MS002413 2020MS002413) doi: https://doi.org/10.1029/2020MS002413

Kelp, M. M., Jacob, D. J., Kutz, J. N., Marshall, J. D., & Tessum, C. W. (2020). Toward stable, general machine-learned models of the atmospheric chemical system. *Journal of Geophysical Research: Atmospheres*, *125*(23), e2020JD032759. Retrieved from `https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2020JD032759` (e2020JD032759 2020JD032759) doi: https://doi.org/10.1029/2020JD032759

Kelp, M. M., Jacob, D. J., Lin, H., & Sulprizio, M. P. (2022). An online-learned neural network chemical solver for stable long-term global simulations of atmospheric chemistry. *Journal of Advances in Modeling Earth Systems*, *n/a*(n/a), e2021MS002926. Retrieved from `https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2021MS002926` (e2021MS002926 2021MS002926) doi: https://doi.org/10.1029/2021MS002926

Kelp, M. M., Tessum, C. W., & Marshall, J. D. (2018). Orders-of-magnitude speedup in atmospheric chemistry modeling through neural network-based emulation. *arXiv preprint arXiv:1808.03874*.

Köble, R., & Seufert, G. (2001). Novel maps for forest tree species in europe. In *Proceedings of the 8th european symposium on the physico-chemical behaviour of air pollutants:"a changing atmosphere* (pp. 17–20).

Lane, T. E., Donahue, N. M., & Pandis, S. N. (2008). Effect of no x on secondary organic aerosol concentrations. *Environmental science & technology*, *42*(16), 6022–6027.

Lee, D. D., & Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, *401*(6755), 788–791.

LeNail, A. (2020). Publication-ready nn-architecture schematics. *NN SVG*.

Liao, H., Henze, D. K., Seinfeld, J. H., Wu, S., & Mickley, L. J. (2007). Biogenic secondary organic aerosol over the united states: Comparison of climatological simulations with observations. *Journal of Geophysical Research: Atmospheres*, *112*(D6).

Manders, A. M. M., Builtjes, P. J. H., Curier, L., Denier van der Gon, H. A. C., Hendriks, C., Jonkers, S., ... Schaap, M. (2017). Curriculum vitae of the lotos–euros (v2.0) chemistry transport model. *Geoscientific Model Development*, *10*(11), 4145–4173. Retrieved from `https://gmd.copernicus.org/articles/10/4145/2017/` doi: 10.5194/gmd-10-4145-2017

Manders-Groot, A. M. M., Segers, A. J., & Jonkers, S. (2021). Lotos-euros v2.2.000 reference guide. *TNO Reports*.

Marsland, S. (2014). *Machine learning: An algorithmic perspective, second edition* (2nd ed.). Chapman & Hall/CRC.

Matsui, H. (2017). Development of a global aerosol model using a two-dimensional sectional method: 1. model design. *Journal of Advances in Modeling Earth Systems*, *9*(4), 1921–1947.

Murphy, B. N., Donahue, N. M., Fountoukis, C., Dall'Osto, M., O'Dowd, C., Kiendler-Scharr, A., & Pandis, S. N. (2012). Functionalization and fragmentation during ambient organic aerosol aging: application of the 2-d volatility basis set to field studies. *Atmospheric Chemistry and Physics*, *12*(22), 10797–10816.

Murphy, B. N., & Pandis, S. N. (2009). Simulating the formation of semivolatile primary and secondary organic aerosol in a regional chemical transport model. *Environmental science & technology*, *43*(13), 4722–4728.

Ng, N. L., Kroll, J. H., Keywood, M. D., Bahreini, R., Varutbangkul, V., Flagan, R. C., ... Goldstein, A. H. (2006). Contribution of first-versus second-generation products to secondary organic aerosols formed in the oxidation of biogenic hydrocarbons. *Environmental science & technology*, *40*(7), 2283–2297.

O'Dowd, C., Ceburnis, D., Ovadnevaite, J., Vaishya, A., Rinaldi, M., & Facchini, M. (2014). Do anthropogenic, continental or coastal aerosol sources impact on a

marine aerosol signature at mace head? *Atmospheric Chemistry and Physics*, *14*(19), 10687–10704.

Paatero, P., & Tapper, U. (1994). Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, *5*(2), 111–126.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830.

Rasp, S. (2020). Coupled online learning as a way to tackle instabilities and biases in neural network parameterizations: general algorithms and lorenz 96 case study (v1. 0). *Geoscientific Model Development*, *13*(5), 2185–2196.

Rasp, S., Pritchard, M. S., & Gentine, P. (2018). Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, *115*(39), 9684–9689.

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

Seinfeld, J. H., & Pandis, S. N. (2006). *Atmospheric chemistry and physics : from air pollution to climate change*. Hoboken, N.J. J. Wiley.

Shrivastava, M. K., Lane, T. E., Donahue, N. M., Pandis, S. N., & Robinson, A. L. (2008). Effects of gas particle partitioning and aging of primary emissions on urban and regional organic aerosol concentrations. *Journal of Geophysical Research: Atmospheres*, *113*(D18).

Sturm, P. O. (2021). Advecting superspecies: Reduced order modeling of organic aerosols in lotos-euros using machine learning. *TU Delft Education Repository*. Retrieved from `http://resolver.tudelft.nl/uuid:2c3be50e-5340-4495-a0b7-1670db9be329`

Sturm, P. O., & Wexler, A. S. (2020). A mass-and energy-conserving framework for using machine learning to speed computations: a photochemistry example. *Geoscientific Model Development*, *13*(9), 4435–4442.

Sturm, P. O., & Wexler, A. S. (2022). Conservation laws in a neural network architecture: enforcing the atom balance of a julia-based photochemical model (v0.2.0). *Geoscientific Model Development*, *15*(8), 3417–3431. Retrieved from `https://gmd.copernicus.org/articles/15/3417/2022/` doi: 10.5194/gmd-15-3417-2022

Tsimpidi, A. P., Karydis, V. A., Pozzer, A., Pandis, S. N., & Lelieveld, J. (2014). Oracle (v1. 0): module to simulate the organic aerosol composition and evolution in the atmosphere. *Geoscientific Model Development*, *7*(6), 3153–3172.

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

Walcek, C. J. (2000). Minor flux adjustment near mixing ratio extremes for simplified yet highly accurate monotonic calculation of tracer advection. *Journal of Geophysical Research: Atmospheres*, *105*(D7), 9335–9348.

Yuval, J., O'Gorman, P. A., & Hill, C. N. (2021). Use of neural networks for stable, accurate and physically consistent parameterization of subgrid atmospheric processes with good performance at reduced precision. *Geophysical Research Letters*, *48*(6), e2020GL091363.

Yuval, J., Pritchard, M., Gentine, P., Zanna, L., & Fan, J. (2021). Call for papers on machine learning and earth system modeling. *Eos*, *102*.

Zanten, M. v., Sauter, F., Wichink Kruit, R., Jaarsveld, J. v., Pul, W. v., & Wichink Kruit, R. (2010). *Description of the DEPAC module. Dry deposition modelling with DEPAC_GCN2010* (Tech. Rep.). Bilthoven, The Netherlands: Rijksinstituut voor volksgezondheid en Milieu, RIVM report 680180001. Retrieved from `http://www.rivm.nl/Documenten_en_publicaties/Wetenschappelijk/Rapporten/2010/oktober/Description_of_the_DEPAC_module_Dry_deposition_modelling_with_DEPAC_GCN2010`

Zhang, L. (2001). A size-segregated particle dry deposition scheme for an atmospheric aerosol module. *Atmospheric Environment*, *35*(3), 549–560. Retrieved from `http://www.sciencedirect.com/science/article/pii/S1352231000003265` doi: 10.1016/S1352-2310(00)00326-5

Zhao, B., Shrivastava, M., Donahue, N. M., Gordon, H., Schervish, M., Shilling, J. E., ... Fast, J. D. (2020). High concentration of ultrafine particles in the amazon free troposphere produced by organic new particle formation. *Proceedings of the National Academy of Sciences*, *117*(41), 25344–25351.