



1 DCU-accelerated 3DVAR data assimilation with 2 automatic differentiation for WRF-Chem

3 Hancheng Ye ¹, Zengliang Zang¹, Wei You¹, Yiwen Hu¹, Ning Liu¹, Yi Li¹

4 ¹College of Meteorology and Oceanography, National University of Defense Technology, Changsha,
5 China

6 **Correspondence:** Wei You (ywlx_1987@163.com) and Zengliang Zang (zxlqxy@163.com)

7 **Abstract:** This study developed a PyTorch-based three-dimensional variational (3DVAR) data
8 assimilation system (Py3DVAR) for the Weather Research and Forecasting model coupled to
9 Chemistry (WRF-Chem), which integrates automatic differentiation (AD) to replace traditional
10 manual gradient derivation and adopts Deep Computing Unit (DCU) acceleration for high
11 computational efficiency. Py3DVAR enables the simultaneous assimilation of gaseous pollutants
12 (SO₂, NO₂, CO, O₃) and particulate matter (PM_{2.5}, PM₁₀) and supports flexible deployment on both
13 Central Processing Unit (CPU) and DCU computing platforms. To evaluate its performance and
14 efficiency, idealized and real-case assimilation experiments (27 km and 9 km grid resolutions) were
15 conducted, compared against a traditional CPU-parallelized Fortran-based 3DVAR system (Fortran-
16 3DVAR). Idealized results show Py3DVAR effectively propagates observation information,
17 generating increment fields consistent with Fortran-3DVAR. In real-case experiments, Py3DVAR
18 substantially improves the model initial field quality: at 27 km resolution, correlation coefficients
19 (CORR) for SO₂, NO₂, CO, O₃, PM_{2.5}, and PM₁₀ increased by 0.77, 0.51, 0.71, 0.98, 0.60, and 0.69,
20 respectively; corresponding improvements at 9 km resolution are 0.78, 0.98, 0.66, 0.96, 0.63, and
21 0.78. The root mean square error (RMSE) and mean absolute error (MAE) are also significantly
22 reduced, with analysis field accuracy comparable to Fortran-3DVAR. In terms of computational
23 efficiency, Py3DVAR shows remarkable advantages: on the same CPU platform, the total iteration
24 time at 27 km resolution is only 7.1 s, approximately 8.8 times faster than Fortran-3DVAR (62.5 s);
25 on the DCU platform, the speedup reaches 32.7 times at 27 km and 40.3 times at 9 km. A 24-hour
26 forecast test shows that the improved initial fields have sustained positive effects on short-term
27 forecasts: the improvements persist for over 24 hours for SO₂, CO, PM_{2.5}, and PM₁₀, and for over 6
28 hours for NO₂ and O₃. This study confirms that Py3DVAR achieves order-of-magnitude gains in
29 computational efficiency while maintaining accuracy equivalent to traditional assimilation
30 algorithms, providing a flexible new technical pathway for operational atmospheric chemical data
31 assimilation and future intelligent assimilation systems.

32

33

34



35 **1. Introduction**

36 Air quality models (AQMs) are crucial tools for atmospheric pollution simulation and
37 forecasting (Li et al., 2013; Wang et al., 2022; Gao et al., 2024). Errors in their initial conditions can
38 accumulate during model integration, ultimately leading to significant simulation biases. Data
39 assimilation technology, which merges multi-source observations (e.g., from surface stations and
40 satellites) with model background fields using specific weighting, generates a theoretically optimal
41 analysis field. This effectively improves the model's initial conditions and is a core technique for
42 enhancing simulation accuracy, widely used in both research and operational forecasting (Engelen et
43 al., 2009; Ma et al., 2025).

44 Compared to the long-standing development of meteorological data assimilation, atmospheric
45 chemical data assimilation started relatively late, and its theoretical framework and operational
46 applications are still under continuous improvement. Early explorations in atmospheric chemical data
47 assimilation were largely constrained by limitations in observational capabilities and the ability of
48 atmospheric chemistry models to simulate complex mixed aerosol components. Consequently,
49 research initially focused on gaseous precursors like SO₂, NO₂, O₃, and key chemical components
50 like HNO₃, which have relatively well-understood chemical mechanisms. The assimilation of
51 particulate matter (PM) components was less common due to the high complexity of their sources
52 and chemical mechanisms. Elbern et al. (1997) used the RADM2 (Second-Generation Regional Acid
53 Deposition Model) gas-phase chemical mechanism and a variational method to assimilate ozone
54 observations, successfully optimizing the initial concentrations of its precursors NO_x and VOCs, and
55 enabling quantitative analysis of ozone chemical sensitivity. In earlier work, Austin et al. (1992)
56 assimilated LIMS satellite measurements of HNO₃ and O₃ into a 2D chemical model, effectively
57 identifying and correcting systematic biases in the original chemical model, significantly improving
58 the analysis accuracy of stratospheric chemical composition, and demonstrating the key role of four-
59 dimensional assimilation in refining chemical models. Frydendall et al. (2009) developed a data
60 assimilation algorithm based on statistical interpolation and coupled it with the Danish Eulerian
61 Operational Model (DEOM), conducting assimilation studies on surface ozone and effectively
62 improving forecast performance by optimizing assimilation parameters.

63 With advancements in observational technology and the continuous improvement of
64 atmospheric chemistry modeling systems, the focus of atmospheric chemical data assimilation has
65 gradually expanded to include particulate pollutants with more complex compositions and sources.
66 In the early stages, limited by insufficient simulation of aerosol physicochemical processes,
67 assimilation studies often focused on constraining models using total concentration observations like
68 PM_{2.5} and PM₁₀, employing relatively simple methods. For example, Denby et al. (2008) applied
69 both Statistical Interpolation (SI) and Ensemble Kalman Filtering (EnKF) data assimilation methods
70 to an atmospheric chemistry model. They found that both methods significantly corrected the model's
71 systematic underestimation of PM₁₀ concentration, with the SI method reducing the RMSE of daily



72 average concentration from $16.7 \mu\text{g m}^{-3}$ to $9.2 \mu\text{g m}^{-3}$, and noted that the improvement from SI was
73 significantly better than that from EnKF. Tombette et al. (2009) used the Optimal Interpolation
74 method to assimilate surface PM_{10} observations from France into a regional chemical transport
75 model, effectively enhancing the short-term forecast accuracy of total aerosol mass (PM_{10} and $\text{PM}_{2.5}$)
76 across Europe, demonstrating the potential application of this method in operational forecasting.

77 As a core method in data assimilation, the 3DVAR approach has a clear theoretical foundation,
78 mature computational framework, and stable system performance. In terms of implementation, its
79 computational structure is clear, facilitating code development, and the overall computational cost is
80 acceptable with modern computing resources (Jiang et al., 2013; Seo et al., 2023; Chen et al., 2024;
81 Zhou et al., 2025). Given these advantages, it has been widely adopted by major international
82 operational forecasting systems such as ECMWF and NCEP. Pagowski et al. (2010) introduced
83 surface ozone and fine particulate matter ($\text{PM}_{2.5}$) observations into the WRF-Chem model using the
84 GSI (Grid-point Statistical Interpolation) 3DVAR system. Results showed that while assimilation
85 significantly reduced forecast errors, these errors grew rapidly in the early forecast hours.

86 With the advancement of observational techniques and the refinement of model mechanisms,
87 simply correcting total aerosol mass often fails to fundamentally optimize the simulation of its
88 internal chemical composition. The target variables for aerosol assimilation are shifting from total
89 mass concentration towards a multivariate system encompassing key chemical components such as
90 sulfate, nitrate, and organic carbon. Li et al. (2013) integrated a 3DVAR algorithm with the WRF-
91 Chem model employing the MOSAIC aerosol scheme. By assimilating $\text{PM}_{2.5}$ and its chemical
92 component observations, they achieved a comprehensive analysis of the concentrations and size
93 distributions of eight key species, including black carbon, organic carbon, sulfate, and nitrate,
94 significantly improving the simulation and forecast accuracy of these components. Jiang et al. (2013)
95 implemented the assimilation of PM_{10} observations using the 3DVAR method, providing an
96 improved initial field for the WRF-Chem model by optimizing various aerosol components.
97 However, the assimilation effect diminished rapidly due to vertical mixing and advective processes,
98 highlighting the importance of optimizing the aerosol vertical initial state. By designing various
99 assimilation scenarios, Zang et al. (2016) confirmed that introducing aircraft-measured aerosol data
100 into traditional surface-based aerosol data assimilation significantly optimized aerosol mass
101 concentration, vertical distribution, and chemical composition. They also found that the forecast skill
102 improvement from vertical observations lasted longer than that from surface data, underscoring the
103 important influence of vertical resolution on forecast performance.

104 To more directly optimize the three-dimensional distribution characteristics of aerosols, the
105 focus of aerosol data assimilation research has gradually shifted towards more comprehensive
106 aerosol optical parameters (Zhang et al., 2011; Ye et al., 2021). Aerosol extinction coefficient and
107 aerosol optical depth (AOD), as key optical parameters characterizing the vertical distribution and
108 total column burden of aerosols, have become important observational data for the new generation of



109 aerosol data assimilation (Chen et al., 2017; Pang et al., 2020). Aerosol extinction coefficient profiles
110 directly constrain the vertical structure, while AOD provides a key constraint on the total aerosol
111 load in the atmospheric column. Therefore, assimilating aerosol extinction coefficient and AOD data
112 is regarded as an important approach for optimizing the vertical structure of aerosols and enhancing
113 air quality forecasting capability (Pang et al., 2018; Ye et al., 2022). Liang et al. (2020) developed
114 observation operators and their adjoints for aerosol extinction coefficient based on the IMPROVE
115 (Interagency Monitoring of Protected Visual Environments) equation, and constructed a 3D-Var
116 assimilation system around this core. By simultaneously assimilating surface station particulate mass
117 concentrations and lidar aerosol extinction coefficient profiles, they validated the effectiveness of
118 multi-source observation collaborative assimilation in improving PM_{2.5} simulation and forecasting.
119 To simultaneously improve the vertical distribution of aerosols and the near-surface forecast effect,
120 Ye et al. (2020) developed a collaborative assimilation scheme using the WRF-Chem model and the
121 3D-Var method. Building upon the assimilation of surface PM_{2.5} and PM₁₀ observations, they further
122 integrated CALIPSO satellite extinction coefficient data. Comparative analysis of results using
123 ground-based lidar and AOD data showed that incorporating aerosol three-dimensional profile data
124 on top of surface data improved both the aerosol vertical distribution and the forecast of surface
125 pollutants. To investigate the improvement of multi-source data assimilation on the three-
126 dimensional representation of aerosols, Zang et al. (2020) found that after introducing AOD data into
127 the 3D-Var system, the simulated aerosol vertical structure aligned better with CALIPSO satellite
128 observations, revealing the key role of AOD data in constraining the model's aerosol vertical
129 distribution. Wang et al., based on Mie scattering theory and the 3D-Var method, constructed tangent
130 linear and adjoint models for aerosol optical properties, enabling the direct assimilation of satellite
131 AOD. This effectively improved the simulation accuracy of AOD and PM_{2.5}, particularly correcting
132 the model's underestimation during heavy pollution events.

133 However, with the increasing richness of observational data, the ongoing refinement of model
134 grids, and the detailed depiction of complex physicochemical processes and the continuous
135 expansion of model variables in atmospheric chemistry models, the state variables and matrix
136 dimensions in assimilation systems have increased significantly (Liang et al., 2020; Wang et al.,
137 2022). This leads to a sharp expansion in the computational scale of traditional data assimilation,
138 causing traditional methods to face the problem of excessive computational load. This not only
139 prolongs the assimilation time but also makes it difficult to meet the timeliness requirements for real-
140 time rapid air pollution forecasting.

141 In recent years, the rapid advancement of artificial intelligence (AI) technology has provided
142 new ideas for addressing this challenge. GPUs (Graphics Processing Units) and DCUs employ a
143 massively parallel architecture composed of thousands of computing cores. This design characteristic
144 makes them particularly suitable for handling highly parallel operations like high-dimensional matrix
145 computations, providing powerful computing support for large-scale data processing, thereby



146 achieving orders-of-magnitude improvement in computational efficiency at the hardware level
147 (Boudier et al., 2023; De et al., 2021; He et al., 2025). On the other hand, various neural network
148 models possess powerful capabilities for nonlinear fitting and feature learning, enabling them to
149 directly learn complex physical-chemical relationships from vast amounts of data, thus avoiding the
150 tedious process of constructing tangent linear and adjoint models required in traditional methods.
151 Leveraging these advantages, an increasing number of studies in recent years have begun to
152 introduce AI methods into the assimilation process, exploring their potential for enhancing the
153 performance and revolutionizing the process of traditional assimilation algorithms through various
154 neural networks. Intelligent assimilation has thus gradually developed into an emerging research
155 direction aimed at optimizing or replacing traditional assimilation algorithms (Zhang et al., 2021; Li
156 et al., 2024; Xu et al., 2024; Lu et al., 2025).

157 It should be noted that the early application of AI in assimilation was essentially a local
158 enhancement of the traditional processes, rather than a systematic replacement. Dong et al. (2022)
159 bypassed the explicit development of complex tangent linear and adjoint models required in
160 traditional four-dimensional variational (4DVAR) assimilation by introducing a Bilinear Neural
161 Network (BNN), constructing a purely data-driven ML-4DVAR framework. They validated on the
162 Lorenz-96 model that this framework could significantly improve computational efficiency while
163 ensuring assimilation accuracy. To address the issue of sampling error caused by insufficient
164 ensemble members due to the high computational cost of the classical Ensemble Kalman Filter in
165 high-dimensional systems, Chattopadhyay et al. (2023) proposed a hybrid ensemble Kalman filter.
166 This approach uses deep learning surrogate models to inexpensively generate large ensembles,
167 thereby more accurately estimating the background error covariance and effectively improving the
168 analysis quality of the initial field. Howard et al. (2024) combined machine learning with the
169 Ensemble Kalman Filter (EnKF) and successfully achieved effective assimilation of high-resolution
170 observational data in the Lorenz-96 system. This enhanced assimilation method not only reduced the
171 RMSE of the initial field by about 37% but also significantly improved the accuracy and reliability
172 of forecasts up to 10 days. Addressing the issues of high computational cost and limited forecast
173 improvement in traditional Chemical Transport Model (CTM) assimilation methods, Li et al. (2025)
174 proposed an efficient machine learning self-evolving assimilation system (OIRF-LEnKF v1.0). Tests
175 showed that the assimilation time of this system was only 11.41–16.60% of that of traditional
176 methods. While significantly improving computational efficiency, its self-evolution mechanism
177 notably reduced forecast errors. Compared to the quality of existing reanalysis datasets, the
178 reanalysis dataset generated by this method was of higher quality. Beauchamp et al. (2025)
179 constructed a neural variational assimilation system based on 4DVarNet. This system used a
180 Variational Autoencoder (VAE) with embedded self-attention mechanism as the prior model,
181 replacing the static background error covariance relied upon in traditional assimilation, achieving
182 deep integration of the variational assimilation framework and deep learning. Results indicated that
183 this system significantly outperformed existing operational methods in terms of analysis accuracy,



184 convergence speed, and the ability to resolve small-scale ocean features.

185 In recent years, AI assimilation research is shifting towards an end-to-end paradigm. This
186 paradigm aims to use neural networks to directly establish a nonlinear mapping from the observation
187 field and background field to the model forecast field, fundamentally bypassing the complex
188 processes in traditional assimilation such as background error covariance estimation, observation
189 operator construction, and iterative optimization. Xu et al. (2025) proposed the FuXi-DA end-to-end
190 deep learning assimilation framework, which circumvented the complex processes of traditional
191 assimilation and achieved direct mapping from raw satellite observations to an optimized analysis
192 field. This framework employed multi-branch U-nets and a unified fusion neural network, effectively
193 avoiding complex steps like observation operator construction and observation thinning in traditional
194 assimilation. By assimilating observations from the Advanced Geosynchronous Radiation Imager
195 (AGRI) aboard the Fengyun-4B satellite, the framework was proven to effectively optimize the
196 analysis field, enhance forecast skill, and the results demonstrated good physical plausibility.

197 However, although AI-driven assimilation methods significantly improve computational
198 efficiency, they sacrifice physical interpretability due to the lack of dynamical constraints from
199 physical models and the "black-box" nature of neural networks. This not only limits their in-depth
200 analysis of pollution formation mechanisms but also poses potential risks for their application in
201 operational forecasting requiring high reliability. In contrast, the traditional 3D-Var assimilation
202 method, based on a strict variational optimization theoretical framework, has a clear physical
203 meaning and mathematical traceability for the background field correction process, yielding
204 transparent and highly credible results. However, traditional methods are typically implemented
205 using CPU clusters and traditional languages like Fortran. When facing high-dimensional
206 atmospheric chemical assimilation tasks, they are often constrained by computational efficiency
207 bottlenecks, making it difficult to meet the timeliness requirements of real-time forecasting.

208 To meet the dual demands of physical interpretability and high-performance computational
209 efficiency, this study constructs a three-dimensional variational assimilation system based on the
210 PyTorch deep learning framework, with its core computational workflow can be simultaneously
211 deployed on both CPU and DCU. The core innovation of this system lies in the deep integration of
212 "principle inheritance" and "implementation innovation": At the theoretical level, it strictly adheres
213 to the variational optimization theory of 3DVAR, fully retaining the core logic of cost function
214 construction and gradient descent, ensuring the physical transparency of the assimilation process and
215 the interpretability of the results. At the technical implementation level, it fully leverages the
216 automatic differentiation capability of the PyTorch framework and the native parallel computing
217 advantages of DCU, modernizing the core modules of assimilation. Key steps such as cost function
218 tensor operations, gradient calculation, and matrix multiplication are entirely migrated to the
219 PyTorch framework for execution, utilizing the massive parallel computing power of DCU to break
220 through computational efficiency bottlenecks, ultimately achieving a leap in computational



221 efficiency while maintaining physical consistency. This design strategy preserves the physical
222 reliability of traditional assimilation methods while fully utilizing the extreme computing power of
223 the deep learning hardware ecosystem, establishing a fused assimilation solution for the large-scale,
224 rapid assimilation of multiple pollutants including SO₂, NO₂, CO, O₃, PM_{2.5}, and PM₁₀, which
225 possesses timeliness, transparency, and strict physical interpretability.

226 **2. Methods and data**

227 **2.1 Air quality model**

228 The Weather Research and Forecasting model with Chemistry (WRF-Chem), was jointly
229 developed by the National Center for Atmospheric Research (NCAR), the National Oceanic and
230 Atmospheric Administration (NOAA), and other institutions. Its core strength lies in the fully online,
231 bidirectional coupling between meteorological and chemical fields, enabling a precise description of
232 pollutant transport, transformation, and removal under complex meteorological conditions. The
233 meteorological fields drive the spatiotemporal distribution and transformation of chemical species
234 through physical processes such as dynamic advection, turbulent diffusion, and dry/wet deposition.
235 Conversely, the chemical fields can feedback on meteorological elements through mechanisms like
236 aerosol radiative forcing and cloud condensation nuclei effects. This online coupling ensures
237 spatiotemporal consistency and physical compatibility between atmospheric dynamic processes and
238 chemical transformations.

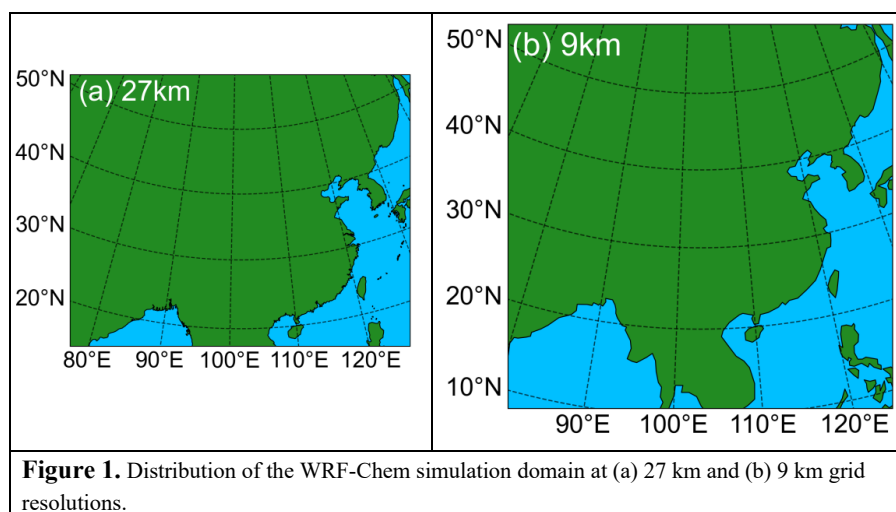
239 The WRF-Chem used in this study integrates the Aerosol Module 5 (AERO5) from the
240 Community Multiscale Air Quality Modeling System (CMAQ). It also features a flexible gas-phase
241 chemistry module that can utilize multiple chemical mechanisms, including the Statewide Air
242 Pollution Research Center (SAPRC), Carbon Bond IV (CB-IV), and Regional Acid Deposition
243 Model version 2 (RADM2) mechanisms (Li et al., 2010, 2011a, 2011b, 2012).

244 Through over two decades of development, this model version incorporates advanced chemical
245 reactions and processes crucial for atmospheric particulate formation. For example, it includes
246 irreversible aqueous-phase processes of glyoxal and methylglyoxal on particle surfaces, which
247 improves the simulation of secondary organic aerosol (SOA) over China (Xing et al., 2018). A
248 parameterization for heterogeneous nitrous acid (HONO) processes has been implemented,
249 enhancing the model's capability to simulate atmospheric oxidizing capacity and SOA (Feng et al.,
250 2005, 2016). The suppression effect of organic coatings on the heterogeneous hydrolysis uptake of
251 dinitrogen pentoxide (N₂O₅) is also parameterized, leading to better simulation of nitrate aerosols.
252 Furthermore, the model incorporates sulfate formation via transition-metal catalysis and Criegee
253 intermediate reactions, improving sulfate simulation. Additionally, a module accounting for aerosol–
254 radiation–photolysis interactions has been developed, enabling assessment of aerosol–cloud
255 interactions, radiative effects of brown carbon aerosols, and direct and indirect effects of various
256 aerosol components.



257 To comprehensively evaluate the computational efficiency and assimilation performance of the
 258 Py3DVAR system, this study designed two single-nested grid simulation schemes based on the
 259 WRF-Chem model. Both schemes share the same vertical configuration of 34 layers, differing only
 260 in horizontal resolution. Scheme 1 employs a coarse-resolution grid with dimensions of 210
 261 (meridional) \times 168 (zonal) grid points and a horizontal resolution of 27 km. This setup is suitable
 262 for rapid verification and benchmark testing. Scheme 2 utilizes a high-resolution grid, sized 600 \times
 263 600 grid points with a 9 km horizontal resolution, approaching the scale of operational assimilation
 264 tasks. It is designed to assess the system's computational performance under high-resolution
 265 conditions. The simulation domains are shown in Fig. 1.

266 To maintain focus and systematically evaluate the core capabilities of the Py3DVAR
 267 framework, this study primarily conducted experimental design and result analysis using the coarse
 268 27 km grid (210 \times 168). The results for the background error covariance, idealized experiments, and
 269 forecast experiments are all presented and discussed at this 27 km resolution. As a supplementary
 270 assessment, a real-case assimilation experiment was performed on the 9 km grid. This serves as a
 271 comparative analysis against the corresponding 27 km results, aiming to examine the scalability of
 272 the Py3DVAR system's computational efficiency and assimilation effectiveness with increasing grid
 273 resolution. Through this "primary-supplementary" experimental design, the study provides an in-
 274 depth analysis of the system's core performance while preliminarily exploring its potential for high-
 275 resolution applications.



276

277 2.2 Data

278 The meteorological fields for the WRF-Chem model were driven by the NCEP FNL (Final
 279 Operational Global Analysis, <https://gdex.ucar.edu/datasets/d083002/>) data, which has a horizontal



280 resolution of $1^\circ \times 1^\circ$ and a temporal resolution of 6 hours. Provided by NCAR and NCEP, the FNL
281 dataset assimilates multi-source global observations from surface stations, radiosondes, satellites, etc.
282 It supplies physically consistent meteorological fields and serves as crucial input for generating
283 reliable initial and boundary conditions for the model. This dataset has been widely used in various
284 mesoscale numerical simulations.

285 Anthropogenic emissions were derived from the 2013 Multi-resolution Emission Inventory for
286 China (MEIC, <http://www.meicmodel.org/index.html>), provided by Tsinghua University. This
287 inventory has a spatial resolution of 6 km, covers most regions of China, and includes emission data
288 for the gaseous pollutants and aerosol precursors required in this study.

289 Chemical boundary conditions for the WRF-Chem model were obtained from the Whole
290 Atmosphere Community Climate Model (WACCM,
291 <https://www.acom.ucar.edu/waccm/download.shtml>) data released by NCAR. This dataset provides
292 the global spatiotemporal distribution of various chemical components, including ozone and aerosol
293 precursors, ensuring reasonable atmospheric chemical boundary conditions for the regional air
294 quality simulations.

295 Observation data were sourced from the China National Environmental Monitoring Centre
296 (CNEMC, <https://air.cnemc.cn:18007/>), primarily comprising hourly surface concentrations of major
297 air pollutants: SO_2 , NO_2 , CO , O_3 , $\text{PM}_{2.5}$, and PM_{10} . The locations of observation sites follow the
298 distribution described in He et al. (2025). To ensure data quality and suitability for 3D-Var
299 assimilation, the following procedures were applied to the raw observational data:

300 (1) Removal of Invalid Values: All concentration records with negative values for any pollutant were
301 identified as invalid and removed.

302 (2) Extreme Value Control: To prevent interference from anomalously high values, maximum
303 concentration thresholds were set for each pollutant (units: $\mu\text{g m}^{-3}$ for all except CO in mg m^{-3}). The
304 specific thresholds were: $500 \mu\text{g m}^{-3}$ for SO_2 , NO_2 , and O_3 ; $600 \mu\text{g m}^{-3}$ for $\text{PM}_{2.5}$; $1200 \mu\text{g m}^{-3}$ for
305 PM_{10} ; and 500mg m^{-3} for CO . Observations exceeding their respective thresholds were discarded.

306 (3) Grid-based Thinning and Averaging: To match the irregularly distributed station observations to
307 the model grid, a grid-averaging method was employed for data thinning. First, all station
308 observations falling within the same WRF-Chem model grid cell were grouped. Then, for the valid
309 observation sequence within each grid cell, an internal consistency check using the 3-standard-
310 deviation criterion was applied to remove outliers that deviated excessively from the mean. Finally,
311 the arithmetic mean was calculated from all observations that passed the quality control in each grid
312 cell, resulting in a single representative concentration value for that grid. This process generated a
313 regular, gridded observational dataset consistent with the model's grid resolution.

314 2.3. Control variable design



315 The WRF-Chem model employed in this study incorporates aerosol variables that encompass a
316 variety of fine and coarse particle components. The fine particle components include: elemental
317 carbon (ECJ), ammonium (NH4AJ), sulfate (SO4AJ), nitrate (NO3AJ_A/B), organic aerosols
318 (ORGAJ, ORGBAJ, SOA_RXT_G/M), nine categories of secondary organic aerosols generated
319 from the semi-volatile gases of primary organic aerosol (POA) (CPO1AER–CPO9AER), chloride
320 (CLAJ), and accumulation-mode fine particles (P25AJ), among others. The coarse particle
321 components include: dust (DUST_1–DUST_5), anthropogenic coarse aerosol (ANTHA), and coarse-
322 mode soil-derived particles (SOILA). Furthermore, the model also includes gaseous pollutants such
323 as SO₂, NO₂, CO, and O₃. The first 11 categories of fine particle components collectively constitute
324 PM_{2.5}. Adding dust along with the anthropogenic and soil-derived coarse-mode aerosols to this forms
325 the complete PM₁₀ composition.

326 If all the aforementioned aerosol components were directly used as control variables for
327 assimilation, it would provide detailed compositional resolution but would also lead to a drastic
328 increase in computational scale, significantly raising storage demands and reducing computational
329 efficiency. To balance assimilation accuracy with computational cost, this study rationally merges
330 specific components based on their physicochemical properties and particle size characteristics: the
331 dual nitrate components are merged into a single variable; the four organic aerosol components are
332 combined into one group; the nine secondary organic aerosols originating from POA semi-volatile
333 gases are uniformly merged; the dust components are optimally merged according to particle size
334 differences—DUST_1 is retained as an independent control variable, while DUST_2 to DUST_5 are
335 merged into a single variable. Through this merging process, the number of aerosol control variables
336 is reduced from the original numerous components to 12. Gaseous pollutants (SO₂, NO₂, CO, O₃) are
337 directly used as control variables for assimilation. Consequently, the assimilation system constructed
338 in this study utilizes a total of 16 control variables. Table 1 provides the detailed mapping between
339 the WRF-Chem model variables and the 3DVAR control variables.

340

341

342

343

344

345

346

347



Model variables	3DVAR control variables
ECJ	EC
NH4AJ	NH4
SO4AJ	SO4
NO3AJ_A NO3AJ_B	NO3
ORGAJ ORGBAJ SOA_RXT_G SOA_RXT_M	ORG
CPO1AER CPO2AER ... CPO9AER	CPO
CLAJ	CL
P25AJ	P25
DUST_1	DUST1
DUST_2 DUST_3 DUST_4 DUST_5	DUST2
ANTHA	ANTHA
SOILA	SOILA
SO2	SO2
NO2	NO2
CO	CO
O3	O3

348 **Table 1.** Mapping between WRF-Chem model variables and 3DVAR control variables.

349

350 **2.4 Three-dimensional variational assimilation theory**

351 The three-dimensional variational assimilation method is grounded in probability and statistical
 352 theory, treating both the simulated background field and observations as useful information
 353 containing errors. This method constructs a cost function that incorporates the weighted errors of the
 354 background field and the observations, and seeks the state that minimizes this function, known as the
 355 analysis field. Statistically, this analysis field is considered the optimal estimate of the true
 356 atmospheric state. The 3D-Var objective function is:

357
$$J(x) = \frac{1}{2}(x - x_b)^T B^{-1}(x - x_b) + \frac{1}{2}(Hx - y)^T R^{-1}(Hx - y) \quad (1)$$

358 where x is the analysis field, representing the set of control variables at all grid points in the
 359 model. Its dimensionality is jointly determined by the number of control variables, vertical layers,
 360 and horizontal grid points. x_b is the background field, provided by the WRF-Chem model
 361 simulation, and has the same dimensions as x . B is the background error covariance matrix,
 362 representing the statistical characteristics of the background field errors. H is the observation
 363 operator that maps from model space to observation space. y is the observation vector, containing the
 364 quality-controlled observational data, with dimensions of $\text{Obs_Station} \times 6$, where Obs_Station
 365 represents the number of valid observation sites, and 6 corresponds to the six types of pollutants. R is
 366 the observation error covariance matrix, representing the uncertainty of the observational data, and
 367 has dimensions consistent with y . Following Li et al. (2013), the observational error standard



368 deviation in this study is set to half of the background error standard deviation.

369 The background error covariance matrix B has a very large dimension. Storing and performing
370 iterative computations with it directly would pose serious challenges to DCU memory and
371 computational efficiency. To optimize the solution of the cost function, this study refers to the
372 method of Li et al., factorizing and structurally simplifying the B matrix.

373 The specific processing is as follows: First, the B matrix is decomposed into the form $B =$
374 DCD^T , where D is a diagonal matrix composed of the background error standard deviations, and C is
375 the background error correlation matrix. In practical assimilation, the analysis field x is typically not
376 solved for directly. Instead, the incremental field $\delta x = x - x_b$ is introduced to solve for the analysis
377 field indirectly, and the control variable is transformed to δz , where the control variables satisfy
378 $\delta x = DC^{1/2}\delta z$. Defining the transformation operator $F = DC^{1/2}$, the cost function can be rewritten in
379 terms of δz as:

$$380 \quad J(\delta z) = \frac{1}{2}\delta z^T \delta z + \frac{1}{2}(HF\delta z - d)^T R^{-1}(HF\delta z - d) \quad (2)$$

381 where $d = y - Hx_b$ is the innovation vector (observation increment). After minimizing $J(\delta z)$ to
382 obtain the optimal δz , the final analysis field x_a is obtained through the following transformation:

$$383 \quad x_a = x_b + F\delta z \quad (3)$$

384 2.5 Gradient calculation method

385 During the minimization of the cost function, both the Py3DVAR system under the Python
386 framework and the Fortran-3DVAR system under the Fortran framework employ the Limited-
387 memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm to achieve minimization.
388 Specifically, the Py3DVAR system directly utilizes the built-in torch.optim.LBFGS optimizer from
389 the PyTorch framework, which is optimized for accelerated tensor operations. In contrast, the
390 Fortran-3DVAR system implements the L-BFGS algorithm manually in Fortran, employing
391 multi-CPU-thread parallel computing. While the optimization objective is identical for both
392 assimilation systems, they differ fundamentally in their underlying code architecture and
393 parallelization strategies.

394 The Py3DVAR system, built upon the PyTorch deep learning framework, fully leverages its
395 automatic differentiation capability for gradient computation. When minimizing the cost function
396 $J(\delta z)$, the system automatically calculates the gradient $\partial J/\partial \delta z$ via the backpropagation algorithm.
397 This method constructs a computational graph, recording all tensor operations during the forward
398 pass, and then automatically computes the gradients of the objective function with respect to each
399 parameter during the backward pass. This automatic differentiation mechanism not only avoids
400 complex analytical derivation, simplifying the code development process, but also ensures both
401 numerical accuracy and computational efficiency in gradient calculations.



402 In comparison, the traditional Fortran-based Fortran-3DVAR system employs manually derived
 403 analytical gradient formulas for computation. This approach requires pre-deriving the analytical
 404 expression for the gradient of the cost function with respect to the control variable δz through
 405 mathematical analysis:

$$406 \quad \nabla J(\delta z) = \delta z + F^T H^T R^{-1} (H F \delta z - d) \quad (4)$$

407 This derived gradient formula is then explicitly implemented in the program. Although this
 408 method has been widely used in traditional scientific computing, the derivation process is complex
 409 and prone to error, and it struggles to fully exploit the advantages of parallel computing.

410 While the two gradient calculation methods are mathematically equivalent, they exhibit
 411 significant differences in their implementation approach and computational efficiency.

412 2.5 Background error covariance modeling based on tensor product decomposition and the NMC 413 Method

414 Furthermore, the background error correlation matrix C is decomposed into the tensor product
 415 of correlation matrices along the longitude, latitude, and vertical directions, i.e., $C = C_x \otimes C_y \otimes C_z$.
 416 For the horizontal correlation matrices C_x and C_y , horizontal isotropy is assumed. Consequently,
 417 they can be represented by the same decay function, differing only in their dimensions. Their
 418 physical meaning is to describe the horizontal correlation between different grid points in the
 419 longitudinal or latitudinal direction, specifically fitted using a Gaussian kernel function:

$$420 \quad C(x_1, x_2) = \exp[-(x_1 - x_2)^2 / (2L^2)] \quad (5)$$

421 Here, L is the horizontal correlation scale parameter for the background error. The method for
 422 obtaining the specific value of L will be introduced later. x_1 and x_2 represent horizontal grid points.

423 The key components of the background error covariance matrix B —namely, the background
 424 error standard deviation matrix D , the vertical background error correlation matrix C_z , and the
 425 horizontal correlation scale parameter L —are statistically estimated using the NMC method (Parrish
 426 et al., 1992). The specific procedure is as follows: First, 24-hour and 48-hour forecast experiments are
 427 initiated daily at 00:00 UTC for 30 consecutive days. The forecast difference field ε is then
 428 constructed from forecasts valid at the same time, calculated as $\varepsilon = X_{48h} - X_{24h}$.

429 The background error standard deviation matrix D is a diagonal matrix. The values for each
 430 control variable in D are computed from the standard deviation of the 30-day forecast difference
 431 field:

$$432 \quad D = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (\varepsilon_i - \bar{\varepsilon})^2} \quad (6)$$

433 where N is the number of days ($N=30$ in this study). The vertical profiles of the background error
 434 standard deviation for each control variable are shown in Figure 2.



435 Vertical Correlation Matrix C_z : This is obtained by calculating the correlation coefficients of
436 the same variable between different model layers. Its matrix elements can be expressed as:

$$437 \quad [C_z]_{mn} = \frac{\text{cov}(m,n)}{\sigma_{\epsilon_m} \cdot \sigma_{\epsilon_n}} \quad (7)$$

438 Here, ϵ_m and ϵ_n represent the forecast differences at the m-th and n-th layers, respectively, cov
439 denotes the covariance, and σ represents the standard deviation. The distribution of the vertical
440 autocorrelation matrix for the background error of each control variable is shown in Figure 3.

441 Horizontal Correlation Scale Parameter L: This is estimated by analyzing the decay of
442 correlation for the same variable across different horizontal grid distances. The correlation
443 coefficient $\rho(d)$ between all pairs of horizontal grid points is computed as a function of their
444 distance d. L is defined as the maximum grid distance d at which the horizontal correlation
445 coefficient $\rho(d)$ decays to $e^{-1/2}$:

$$446 \quad L = \max\{d \mid \rho(d) \geq e^{-1/2}\} \quad (8)$$

447 The horizontal correlation scale parameter L quantitatively describes the decay characteristics of
448 background error correlation between different grid points with increasing distance. Its physical
449 meaning is that when the distance between two horizontal grid points exceeds L, the correlation
450 coefficient between their background errors falls below $e^{-1/2}$ (approximately 0.6065), indicating
451 that their statistical correlation has significantly weakened. This parameter quantifies the decay rate
452 of horizontal correlation and is a key parameter for describing the horizontal structure of the
453 background error covariance. The computed values of L are used to construct the horizontal
454 correlation matrices C_x and C_y , which characterize the spatial horizontal correlation structure and
455 directly determine the shape of the Gaussian correlation kernel function. The decay of the horizontal
456 background error correlation coefficient with distance and the corresponding L value for each control
457 variable are shown in Fig. 4.

458 To achieve efficient computation, the three correlation matrices C_x , C_y and C_z are further
459 decomposed via Cholesky decomposition. Each is expressed as the product of a lower triangular
460 matrix and an upper triangular matrix (e.g., $C_x = C_{x_{low}} \cdot C_{x_{up}}$), where these two triangular matrices
461 are transposes of each other (i.e., $C_{x_{up}} = C_{x_{low}}^T$, where T denotes transpose). During the actual
462 assimilation process, only the upper triangular matrices for these three directions are retained and
463 used, thereby avoiding direct manipulation of the full B matrix.




Figure 2. Vertical profiles of background error standard deviation for each control variable.

464

465 Fig. 2 shows the vertical profiles of the background error standard deviation, which quantifies
 466 the uncertainty of the model background field across different vertical layers and for different
 467 variables. In the 3D-Var assimilation process, the D matrix determines the relative distribution of
 468 analysis increments, induced by observational information, among the control variables. A larger
 469 background error standard deviation for a variable indicates higher uncertainty in the background
 470 field simulation, allowing for greater adjustment by observational information during assimilation.
 471 Conversely, a smaller standard deviation implies higher confidence in the background field, resulting
 472 in more moderate increments. Fig. 2 reveals significant differences in the near-surface values among
 473 the variables. The background error standard deviations for DUST1 (exceeding $10 \mu\text{g m}^{-3}$) and
 474 DUST2 (exceeding $30 \mu\text{g m}^{-3}$) are the largest, substantially higher than those of other variables (all
 475 below $5 \mu\text{g m}^{-3}$). With increasing height, the background error standard deviations for all variables
 476 generally exhibit a decaying trend, and the differences between them gradually diminish. In terms of
 477 profile shape and decay rate, DUST1 and DUST2 decay most rapidly, while the decay of the other
 478 variables is relatively gradual.




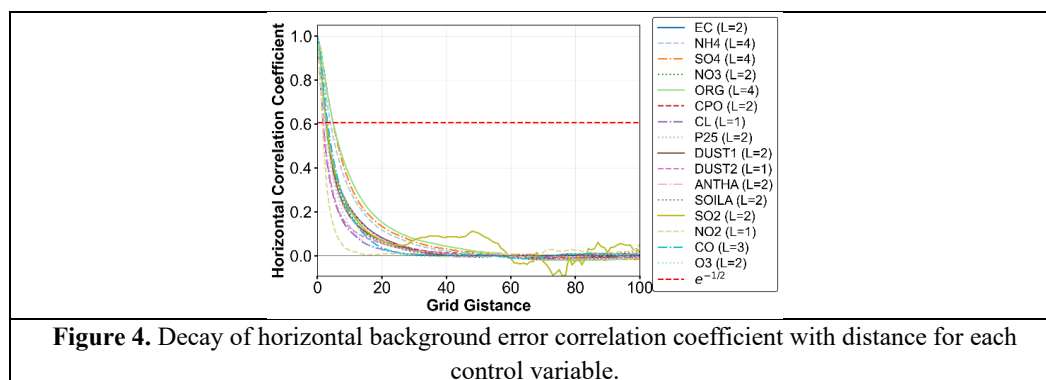
Figure 3. Distribution of the vertical autocorrelation matrix of background error for each control variable.

479

480 Fig. 3 displays the vertical autocorrelation matrices of the background error for each control



481 variable. These matrices play a key role in the assimilation system by constraining the propagation
 482 and distribution of observational information across the model's vertical layers. Analyzing their
 483 spatial structure shows that, although the detailed distributions differ among variables, each exhibits
 484 strong near-surface correlations, and these correlations decrease with increasing height. This
 485 indicates that observational information from the surface layer plays a more dominant role in
 486 adjusting and correcting the lower-level state of the model background field, while adjustments to
 487 higher levels are more strongly constrained.



488 Fig. 4 presents the decay curves of the horizontal background error correlation coefficient with
 489 grid distance for each control variable, along with their corresponding horizontal correlation scale
 490 parameter L. This correlation coefficient determines the horizontal spread and weight allocation of
 491 observational information. The overall pattern shows that the horizontal correlation coefficients for
 492 all variables decay with increasing grid distance, characterized by an initially fast then slower decay
 493 rate. This is consistent with the physical mechanism represented by the Gaussian kernel function:
 494 "strong correlation at short distances and weak correlation at long distances." The decay rates differ
 495 significantly among the variables, leading to calculated L values distributed between 1 and 4. This
 496 reflects the heterogeneity in the horizontal spatial correlation of background errors for different
 497 pollutants. Notably, the decay curve for SO₂ shows noticeable fluctuations beyond a grid distance of
 498 20, exhibiting less smoothness compared to other variables. This phenomenon may stem from the
 499 following reason: SO₂, primarily emitted from point sources (e.g., power plants, industrial zones),
 500 has a background error field that exhibits stronger spatial heterogeneity and discontinuity. When the
 501 distance exceeds the effective propagation range of the observational information, residual local error
 502 signals cannot be smoothly fitted, ultimately manifesting as curve fluctuations.

503 3. Experimental design and evaluation metrics

504 3.1 Experimental design

505 To systematically evaluate the performance and effectiveness of the Py3DVAR assimilation
 506 system, this study designed three sets of numerical experiments: an idealized experiment, real case
 507 assimilation experiments, and a short-term forecast experiment.



508 First, an idealized twin-observation experiment was designed based on the 27 km grid. Its
509 purpose is to test the system's capability to generate physically reasonable increment fields under
510 observations of varying magnitudes. Two observation sites were positioned far apart to avoid spatial
511 correlation interference, ensuring the independence of constraints from different observation
512 magnitudes. The low-concentration observation site (37.2104°N, 116.0108°E) had prescribed values
513 of 0.001 $\mu\text{g m}^{-3}$ for SO₂, NO₂, O₃, PM_{2.5}, and 0.0015 $\mu\text{g m}^{-3}$ for PM₁₀, and 0.001 mg m⁻³ for CO. The
514 high-concentration site (37.8666°N, 92.3785°E) had values of 300 $\mu\text{g m}^{-3}$ for SO₂, NO₂, O₃, 500 μg
515 m⁻³ for PM_{2.5}, 700 $\mu\text{g m}^{-3}$ for PM₁₀, and 300 mg m⁻³ for CO. By comparing the analysis field with
516 the background field, the ability of the assimilation system to produce physically sound increment
517 fields under different observation intensities can be verified.

518 In addition to the twin-point idealized experiment, a real case assimilation experiment was
519 conducted for the background field at 00:00 UTC on October 27, 2023. This experiment serves as the
520 core comparative test of this study and is the only one covering both grid resolutions. Unlike the
521 idealized experiment, it utilized all quality-controlled real observation data and was executed
522 simultaneously on both the 27 km and 9 km grids. For each grid, three parallel assimilation runs
523 were performed: (1) using the Fortran-3DVAR system on the CPU platform (labeled Fortran-
524 3DVAR); (2) using the Py3DVAR system on the CPU platform (labeled Py3DVAR_CPU); and (3)
525 using the Py3DVAR system on the DCU platform (labeled Py3DVAR_DCU). This resulted in a total
526 of six assimilation experiments across the two grids. The Fortran-3DVAR system runs exclusively on
527 the CPU platform, whereas the Py3DVAR system is deployed on both CPU and DCU platforms,
528 enabling a multi-dimensional comparison of assimilation performance and computational efficiency.

529 Finally, based on the results of the real case assimilation, a 24-hour short-term forecast
530 experiment was conducted on the 27 km grid. Two types of initial fields were used: the original, non-
531 assimilated background field (serving as the control experiment, labeled Control) and the assimilated
532 analysis fields (from Fortran-3DVAR and Py3DVAR_DCU, respectively). By comparing the forecast
533 results from the control and assimilation experiments, the improvement effect of the assimilation
534 analysis on short-term air quality forecasting and its persistence can be assessed.

535 To ensure rigorous comparison, all aforementioned experiments used identical background error
536 covariance, background fields, and observation data. All computational processes employed 32-bit
537 floating-point precision to avoid introducing extraneous numerical errors due to precision
538 differences. This approach aims to accurately test the result fidelity of the Py3DVAR system and
539 quantitatively compare the computational efficiency of the two systems and the acceleration effect of
540 the DCU processor.

541 The specifications of the computing platforms involved in this study are as follows: The CPU
542 computing platform is a single node equipped with two Intel Xeon 7495 processors (128 physical
543 cores @ 3.0 GHz) and 768 GB of memory. All 128 physical cores were utilized for parallel



544 computing during the experiments. The DCU acceleration platform centers on a K500SM_AI DCU
545 with 64 GB of device memory, supporting the parallel accelerated computations for the
546 Py3DVAR_DCU experiments.

547 **3.2 Evaluation methods**

548 To quantitatively evaluate the assimilation and forecasting performance of Py3DVAR, this study
549 employs three statistical metrics: the correlation coefficient (CORR), root mean square error
550 (RMSE), and mean absolute error (MAE). CORR measures the linear correlation between simulated
551 and observed values; RMSE is more sensitive to extreme errors and reflects the overall deviation of
552 the simulation results; MAE measures the average absolute error between simulated and observed
553 values and is less sensitive to outliers. The formulas for each metric are as follows:

$$554 \quad \text{CORR} = \frac{\sum_{i=1}^N (M_i - \bar{M})(O_i - \bar{O})}{\sqrt{\sum_{i=1}^N (M_i - \bar{M})^2 \sum_{i=1}^N (O_i - \bar{O})^2}} \quad (9)$$

$$555 \quad \text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (M_i - O_i)^2} \quad (10)$$

$$556 \quad \text{MAE} = \frac{1}{N} \sum_{i=1}^N |M_i - O_i| \quad (11)$$

557 Where:

558 M_i and O_i represent the simulated value and observed value for the i -th sample, respectively.

559 \bar{M} and \bar{O} are the average values of the simulated and observed values, respectively.

560 N is the total number of samples.

561 **4. Results and evaluation**

562 **4.1 Evaluation of idealized experiment results**

563 Fig. 5 and 6 present the surface increment fields generated by the Py3DVAR and Fortran-
564 3DVAR assimilation systems, respectively, in the idealized experiment, where red indicates positive
565 increments and blue indicates negative increments. It should be noted that the $\text{PM}_{2.5}$ and PM_{10}
566 increment fields shown in the figures are not obtained directly through assimilation. Instead, they are
567 synthesized from the increments of multiple aerosol component control variables: the sum of the
568 increments of the first 8 control variables constitutes the $\text{PM}_{2.5}$ increment field, and the sum of the
569 increments of the first 12 control variables constitutes the PM_{10} increment field. In contrast, the
570 increment fields for the four gaseous pollutants (SO_2 , NO_2 , CO , O_3) are directly derived from their
571 corresponding assimilation control variables.

572 For all variables in both systems, the increments exhibit an approximately circular distribution



573 centered on the observation point, with the maximum value at the center decaying gradually with
 574 increasing distance. This spatial structure aligns with the theoretical characteristic in 3DVAR where
 575 observational information spreads horizontally according to the background error horizontal
 576 correlation. Further comparison of the increment fields for different pollutants reveals significant
 577 differences in their horizontal distribution ranges. This is primarily governed by the horizontal
 578 correlation scale parameter L of the background error for each control variable. A larger L value
 579 indicates a wider horizontal spread of observational information, resulting in a correspondingly
 580 larger horizontal distribution range for the increment field.

581 From the comparison of the extremum values (retaining two decimal places) of the increment
 582 fields presented in Fig. 5 and 6, it is clearly evident that the results from the Py3DVAR and Fortran-
 583 3DVAR systems exhibit exceptionally high consistency. Within the bounds of numerical precision,
 584 the difference in extremum values for gaseous pollutant CO is 0 mg m^{-3} , for gaseous pollutants SO₂,
 585 NO₂, and O₃ it does not exceed $0.03 \text{ } \mu\text{g m}^{-3}$, and for particulate pollutants (PM_{2.5}, PM₁₀) it does not
 586 exceed $0.04 \text{ } \mu\text{g m}^{-3}$. These extremely small differences confirm that although Py3DVAR has
 587 undergone fundamental reconstruction in terms of gradient calculation (automatic differentiation)
 588 and optimization algorithm (PyTorch L-BFGS), it has not compromised the scientific accuracy of the
 589 assimilation. In terms of core assimilation capability, it possesses assimilation performance and
 590 numerical reliability comparable to the traditional manually coded Fortran-3DVAR system.

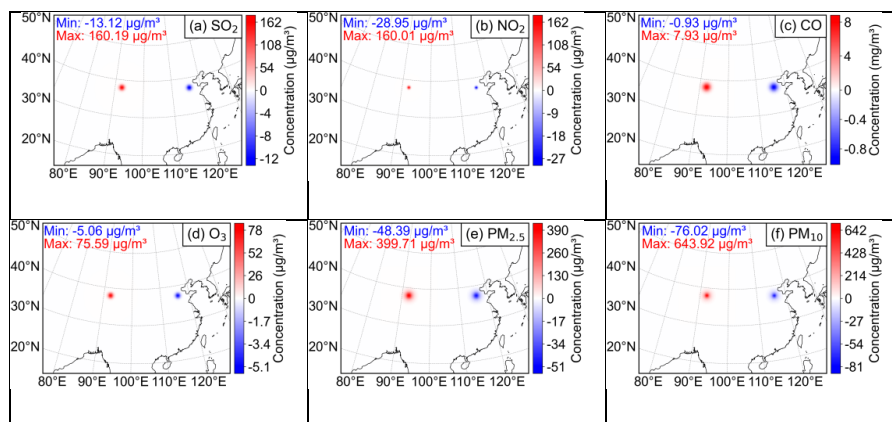
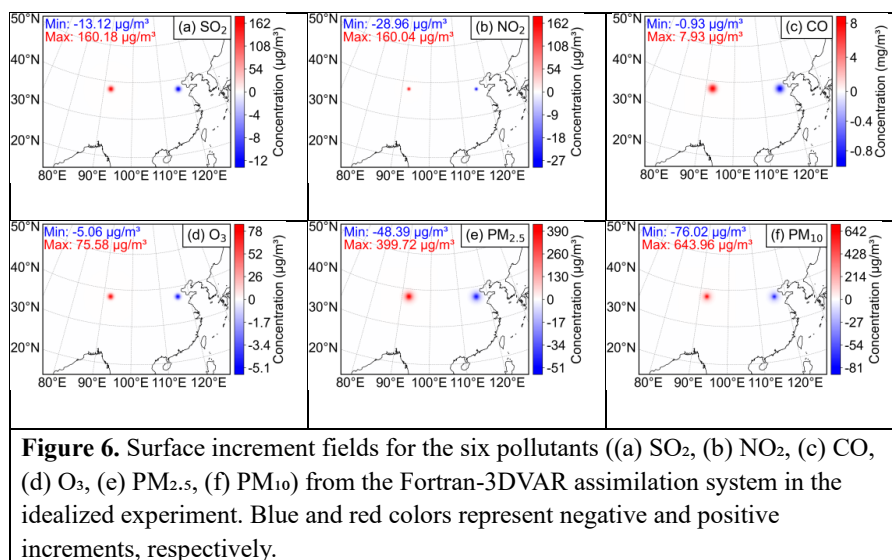


Figure 5. Surface increment fields for the six pollutants ((a) SO₂, (b) NO₂, (c) CO, (d) O₃, (e) PM_{2.5}, (f) PM₁₀) from the Py3DVAR assimilation system in the idealized experiment. Blue and red colors represent negative and positive increments, respectively.

591

592



593 Fig. 7 and 8 present the vertical cross-sections of the increment fields generated by the
 594 Py3DVAR and Fortran-3DVAR assimilation systems, respectively, in the idealized experiment, with
 595 red and blue denoting positive and negative increments, respectively. The distribution of vertical
 596 increments for various pollutants from both systems exhibits a distinct columnar structure. The
 597 deepest color and thickest column in the near-surface layer indicate the largest increment values
 598 there. As height increases, the color fades and the column narrows, corresponding to a weakening of
 599 the increment magnitude. This vertical distribution pattern is fully consistent with the theoretical
 600 expectation encoded in the vertical correlation matrix C_z of the background error covariance,
 601 namely that error correlation peaks in the near-surface layer and decays with altitude. Specifically,
 602 increments above approximately level 10 are near zero, indicating that the effective vertical
 603 propagation range of surface observation information is about 10 layers. At higher levels, constrained
 604 by the decreasing correlation, the magnitude of increment adjustments becomes minimal.

605 Comparing the vertical increment fields of different pollutants reveals differences in the
 606 thickness of their columnar structures. This thickness essentially reflects the horizontal spread of the
 607 observational information; a thicker column represents a larger horizontal influence area. Comparing
 608 these with the horizontal increment fields in Fig. 5 and 6 shows a clear correspondence between the
 609 column thickness in the vertical cross-sections and the distribution range in the horizontal fields.
 610 Both are determined by the horizontal correlation scale parameter L of the background error for each
 611 control variable – a larger L value leads to a wider horizontal spread and, consequently, a thicker
 612 column in the vertical profile. This further verifies the global constraining effect of the background
 613 error's horizontal correlation characteristics on the spatial propagation of observational information.

614 Comparing the Py3DVAR and Fortran-3DVAR systems shows that the distribution patterns,
 615 column thickness, and decay trends of the vertical increment fields for all pollutants are highly



616 consistent between them. This finding corroborates the comparison results of the horizontal
 617 increment fields mentioned previously. The consistency in both vertical and horizontal dimensions
 618 fully validates the ability of the Py3DVAR assimilation system to maintain physical rationality and
 619 numerical accuracy within the three-dimensional space, demonstrating that its reliability and
 620 effectiveness are comparable to the mature Fortran-3DVAR assimilation system.

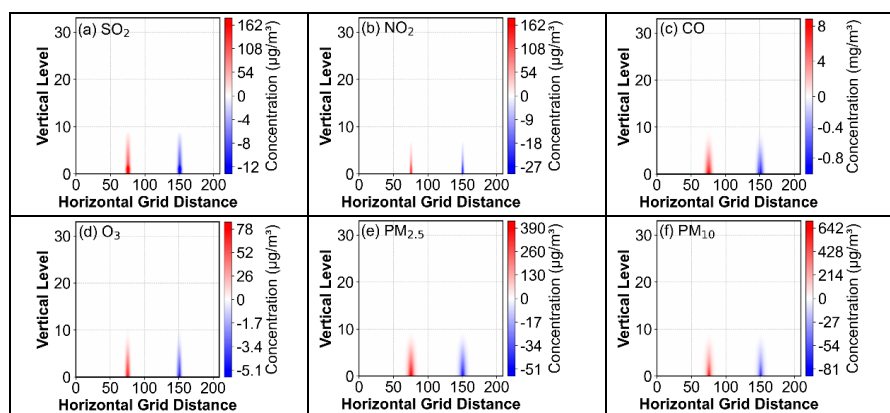


Figure 7. Vertical cross-sections of the increment fields for the six pollutants ((a) SO₂, (b) NO₂, (c) CO, (d) O₃, (e) PM_{2.5}, (f) PM₁₀) from the Py3DVAR assimilation system in the idealized experiment. Blue and red colors represent negative and positive increments, respectively.

621

622

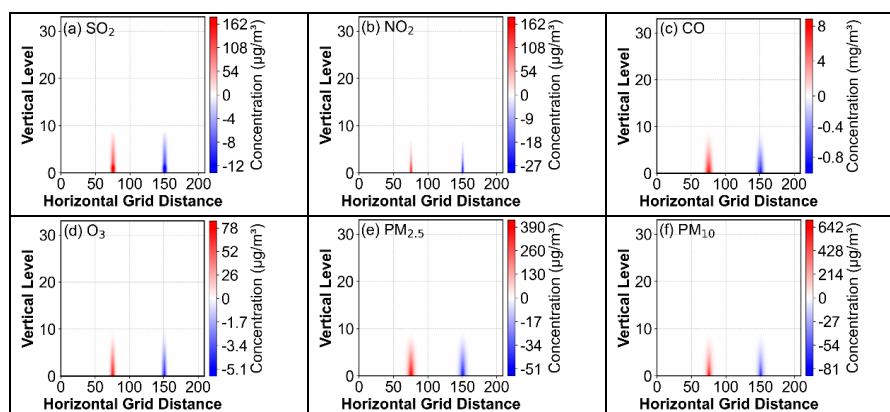


Figure 8. Vertical cross-sections of the increment fields for the six pollutants ((a) SO₂, (b) NO₂, (c) CO, (d) O₃, (e) PM_{2.5}, (f) PM₁₀) from the Fortran-3DVAR assimilation system in the idealized experiment. Blue and red colors represent negative and positive increments, respectively.

623 To provide a more intuitive assessment of the assimilation performance of Py3DVAR versus
 624 Fortran-3DVAR, Fig. 9 compares the background field, the observation field, and the analysis fields



625 from the two assimilation systems under conditions of extremely high (a) and extremely low (b)
 626 observed concentrations. The bar chart clearly shows that for both the Py3DVAR and Fortran-
 627 3DVAR systems, their respective analysis fields are closer to the observed values than the
 628 background field is. This indicates that the assimilation process significantly improves the initial
 629 background field. This improvement effect remains stable under both extremely high and extremely
 630 low concentration conditions, demonstrating the robustness of the assimilation systems.

631 By comparing the analysis fields of the ideal experiments of the two assimilation systems
 632 (results retained to two decimal places), it can be observed that under extremely high observation
 633 concentrations, the differences in SO₂, NO₂, CO, O₃, PM_{2.5}, and PM₁₀ in the analysis fields of the
 634 Py3DVAR and Fortran-3DVAR assimilation systems are 0.02 μg m⁻³, 0.01 μg m⁻³, 0 μg m⁻³, 0 μg
 635 m⁻³, 0 μg m⁻³, and 0.02 μg m⁻³ respectively. Under extremely low observation concentrations, the
 636 differences in SO₂, NO₂, CO, O₃, PM_{2.5}, and PM₁₀ in the analysis fields of the two systems are 0 μg
 637 m⁻³, 0.01 μg m⁻³, 0 μg m⁻³, 0 μg m⁻³, 0.05 μg m⁻³, and 0.08 μg m⁻³ respectively. It is evident that
 638 regardless of extremely high or low observation concentrations, the differences in the analysis fields
 639 of various pollutants between the two systems are very small, and this small difference range is
 640 highly consistent with the previous comparison results of the increment fields. In general, the
 641 Py3DVAR system can effectively use observation information to correct the background field, and
 642 its analysis field results are highly consistent with the mature Fortran-3DVAR system, exhibiting
 643 reliable assimilation performance.






Figure 9. Comparison of the background field (gray), observation field (black), Fortran-3DVAR analysis field (red), and Py3DVAR analysis field (blue) under extremely high (a) and extremely low (b) observation concentration conditions.

644

645 **4.2 Analysis of initial field assimilation effects: comparing control and Py3DVAR/Fortran-**
646 **3DVAR assimilation experiments at the 27 km grid resolution**

647 Fig. 10 and 11 present scatter comparisons between simulated results and observations for the
648 control experiment, the Py3DVAR_DCU assimilation experiment, and the Fortran-3DVAR
649 assimilation experiment at the 27 km grid resolution, respectively. In these figures, blue dots
650 represent the control experiment, while red dots represent the assimilation experiments. Although the
651 Py3DVAR assimilation system was tested on both CPU and DCU platforms, the results differ only
652 due to the computational platform. Therefore, the Py3DVAR_DCU assimilation experiment is
653 selected for analysis here. Regarding the distribution pattern, the blue dots from the Control
654 experiment are relatively scattered and deviate considerably from the 1:1 line. In contrast, the red
655 dots from the assimilation experiments show a clear convergence towards the 1:1 line, indicating that
656 both assimilation systems effectively improve the quality of the initial field, with a very pronounced
657 improvement effect.

658 Examining the specific evaluation metrics, the degree of improvement in the initial fields varies
659 among different pollutants. After assimilation, the CORR values for all pollutants increased
660 significantly. The correlation coefficients for SO₂ and CO both exceeded 0.8, while those for NO₂,
661 O₃, PM_{2.5}, and PM₁₀ reached above 0.9. Concurrently, the RMSE and MAE for all pollutants
662 decreased substantially compared to the Control experiment, further confirming the corrective effect
663 of assimilation on initial field biases.

664 By comparing the assimilation performance of the two systems, Py3DVAR_DCU and Fortran-
665 3DVAR, it can be observed that the red circles are highly close in distribution in the scatter plot,
666 indicating good consistency in the assimilation performance of the two systems. From the
667 perspective of quantitative indicators, for the six pollutants (SO₂, NO₂, CO, O₃, PM_{2.5}, and PM₁₀), the
668 degree of improvement in CORR, RMSE, and MAE between the two systems is basically equivalent,
669 with only a minor difference of 0.01 μg m⁻³ in both the MAE of O₃ and the RMSE of PM_{2.5}. This
670 quantitative result is consistent with the difference characteristics of the increment fields and analysis
671 fields in the previous ideal experiments, further confirming the consistency of the two systems in
672 core assimilation performance.

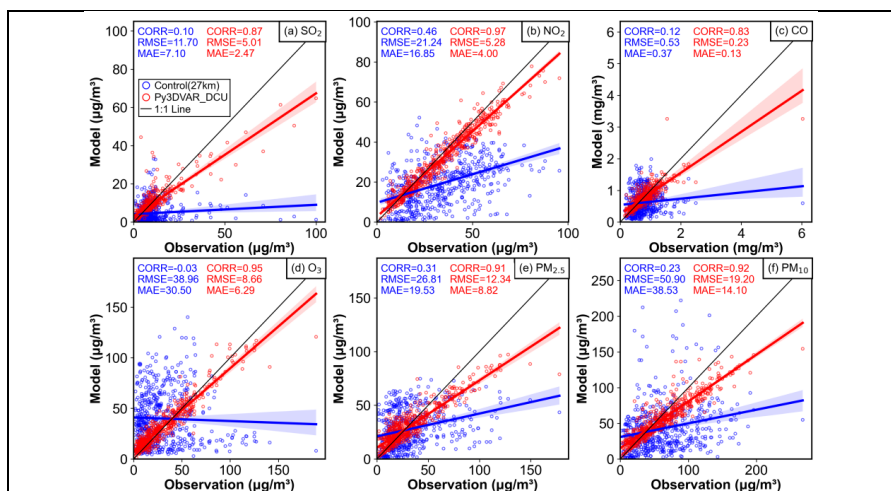


Figure 10. Scatter plots of the initial field for the Control and Py3DVAR_DCU assimilation experiments at the 27 km grid resolution. Red circles represent the Py3DVAR_DCU experiment, blue circles represent the Control experiment, and the black line is the 1:1 line.

673

674

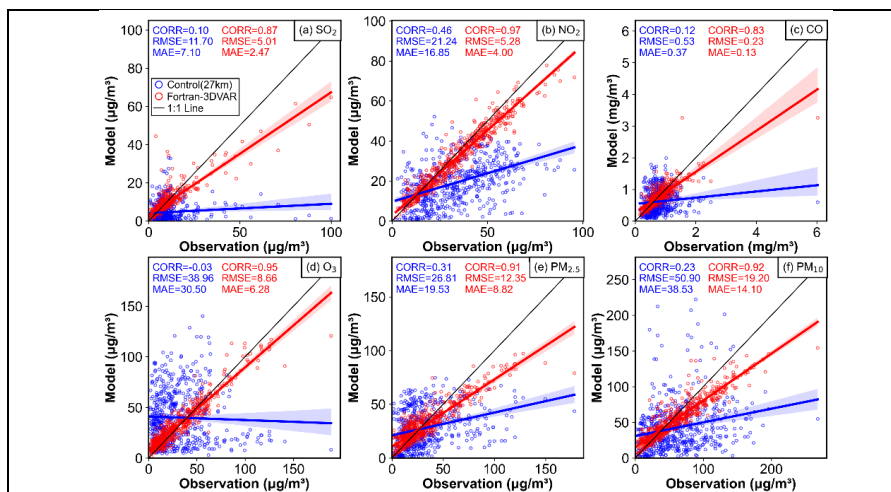


Figure 11. Scatter plots of the initial field for the Control and Fortran-3DVAR assimilation experiments at the 27 km grid resolution. Red circles represent the Fortran-3DVAR experiment, blue circles represent the Control experiment, and the black line is the 1:1 line.

675

676

4.3 Analysis of initial field assimilation effects: comparing control and Py3DVAR/Fortran-3DVAR assimilation experiments at the 9 km grid resolution

677

678

Fig. 12 and 13 display scatter comparisons between simulated results and observations for the control experiment, the Py3DVAR_DCU assimilation experiment, and the Fortran-3DVAR



679 assimilation experiment, respectively, at the 9 km grid resolution. In these plots, blue dots represent
 680 the control experiment, while red dots represent the assimilation experiments. It should be noted that
 681 the 9 km grid experiments serve as a supplement to the 27 km grid experiments, aiming to
 682 preliminarily investigate the system's performance in a higher-resolution scenario.

683 From the control experiment (Control), it can be observed that the model simulations for all six
 684 pollutants are generally lower than the observations. Compared to the control experiment at the 27
 685 km grid, certain differences exist, primarily due to variations in model errors at different grid
 686 resolutions and inconsistencies in the scale of the emission inventory. The assimilation results show
 687 that both Py3DVAR_DCU and Fortran-3DVAR effectively correct the systematic underestimation in
 688 the background field. After assimilation, the scatter points of the analysis fields align more closely
 689 with the 1:1 line, indicating that the assimilation systems maintain good performance at the 9 km
 690 grid. A further comparison between the two assimilation experiments reveals that their scatter
 691 distributions are highly overlapping, and the values of various evaluation metrics are essentially
 692 consistent. This aligns with the conclusions drawn from Figures 10 and 11 for the 27 km grid,
 693 demonstrating that Py3DVAR achieves assimilation effectiveness comparable to the traditional
 694 Fortran-3DVAR across different resolutions.

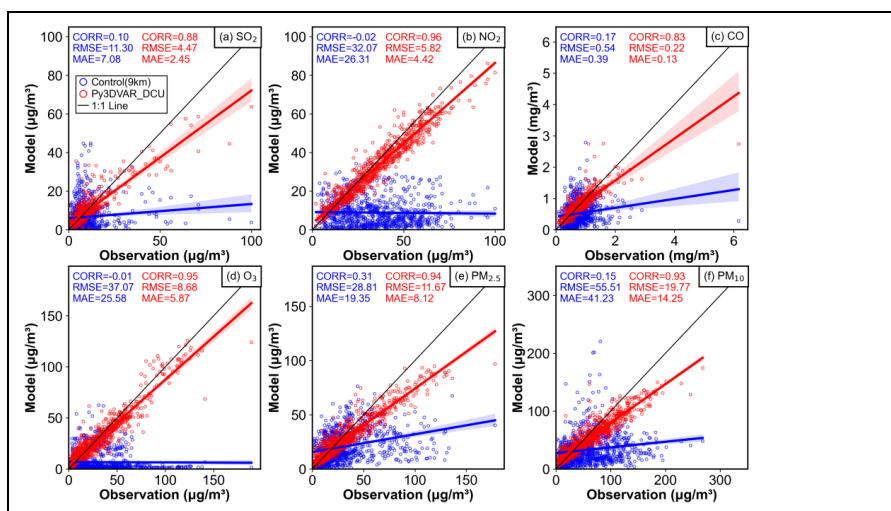
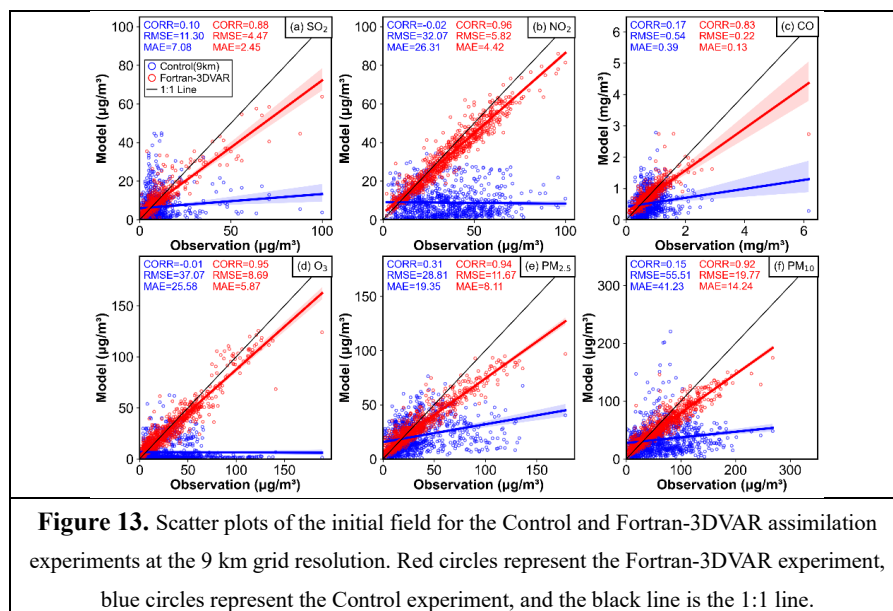


Figure 12. Scatter plots of the initial field for the Control and Py3DVAR_DCU assimilation experiments at the 9 km grid resolution. Red circles represent the Py3DVAR_DCU experiment, blue circles represent the Control experiment, and the black line is the 1:1 line.

695



696

697 4.4 Analysis of the Assimilation System Iteration Process and Cost Function Descent 698 Characteristics

699 Fig. 14 (a) and (b) illustrate the descent trends of the cost function values with iteration counts
700 for different assimilation systems at the 27 km and 9 km grid resolutions, respectively. The figures
701 show that the cost function values for both types of assimilation systems on different computing
702 platforms decrease rapidly as iterations proceed, exhibiting typical nonlinear convergence
703 characteristics: the descent rate is fastest during the initial iterations, gradually narrows as iterations
704 increase, and the curves flatten, indicating that the assimilation systems gradually approach the
705 optimal solution and the convergence rate slows down.

706 At the 27 km grid resolution, the numerical characteristics of the cost functions across the
707 assimilation systems show a high degree of consistency. For the Py3DVAR system, the initial cost
708 function values on the CPU and DCU platforms are 127,484.29 and 127,484.30, respectively, with
709 final convergence costs of 24,906.77 and 24,906.77. For the Fortran-3DVAR system, the initial cost
710 is 127,483.55 and the final convergence cost is 24,906.90. The comparison reveals that the cost value
711 differences for the Py3DVAR system between computing platforms are minimal, attributable to
712 minor numerical deviations resulting from floating-point operation rule differences between the CPU
713 and DCU platforms. The differences in initial and final costs between Py3DVAR and Fortran-
714 3DVAR are slightly larger than the former but remain within an acceptable range overall. These
715 discrepancies are inherent to numerical computations under different algorithmic frameworks and
716 computing architectures, are reasonable, and align with the pattern of numerical deviations in the



717 assimilation results discussed earlier. In summary, the final convergence costs for the three
718 experiment sets are very close, indicating that despite differences in iteration paths, the systems
719 ultimately achieve equivalent convergence outcomes. This conclusion corroborates the earlier
720 assertion of "highly consistent assimilation effectiveness."

721 Analyzing the dynamic characteristics of the iteration process reveals that the initial iteration
722 stage is crucial for cost function reduction, and the Py3DVAR system exhibits a significantly higher
723 cost reduction rate than Fortran-3DVAR. Specifically, Py3DVAR achieves a substantial decrease in
724 the cost function value in the very first iteration, whereas the reduction in the first iteration for
725 Fortran-3DVAR is relatively limited. This difference stems from the distinct implementation
726 approaches of the optimization algorithms in the two systems: Fortran-3DVAR employs a manually
727 coded L-BFGS optimization algorithm, while Py3DVAR relies on the built-in gradient-descent-class
728 optimizer from the PyTorch framework. The differences in engineering implementation directly lead
729 to the divergence in convergence rates, further highlighting the optimization efficiency advantage of
730 the Py3DVAR system. Additionally, comparing the Py3DVAR_CPU and Py3DVAR_DCU
731 experiments shows that, although they use identical code and input data, subtle differences exist in
732 the iteration process due to hardware architecture variations between the computing platforms.
733 However, the overall iteration trends are highly consistent, demonstrating good cross-platform
734 stability.

735 In terms of the number of iterations required for convergence, the Py3DVAR system requires
736 only 15 and 14 iterations to converge on the CPU and DCU platforms, respectively, while Fortran-
737 3DVAR requires 25 iterations, a significantly higher count. The core reason for this disparity is that,
738 although both systems are based on the core principles of the L-BFGS algorithm, there are
739 differences in the engineering details of their implementations—specifically in convergence
740 threshold settings, step-size adaptation strategies, and Hessian matrix approximation update methods.
741 Furthermore, the slower initial cost reduction rate of Fortran-3DVAR ultimately necessitates more
742 iterations to meet the convergence criteria. These results fully demonstrate the more efficient
743 convergence characteristics of the Py3DVAR system.

744 As a comparative test, the pattern of cost function descent at the 9 km grid resolution is highly
745 consistent with that at the 27 km grid: all assimilation systems exhibit the "rapid initial descent
746 followed by gradual flattening" convergence feature, and the differences in iteration processes
747 between Py3DVAR and Fortran-3DVAR also follow the same pattern observed at the 27 km grid.
748 This indicates that the Py3DVAR assimilation system maintains good numerical stability and
749 convergence consistency across different horizontal resolutions.




Figure 14. Descent of the cost function with iteration number for different assimilation systems at (a) 27 km and (b) 9 km grid resolutions. Red, blue, and green lines represent the Py3DVAR_DCU, Py3DVAR_CPU, and Fortran-3DVAR experiments, respectively.

750

751 4.5 Analysis of Iteration Time Consumption in the Assimilation System

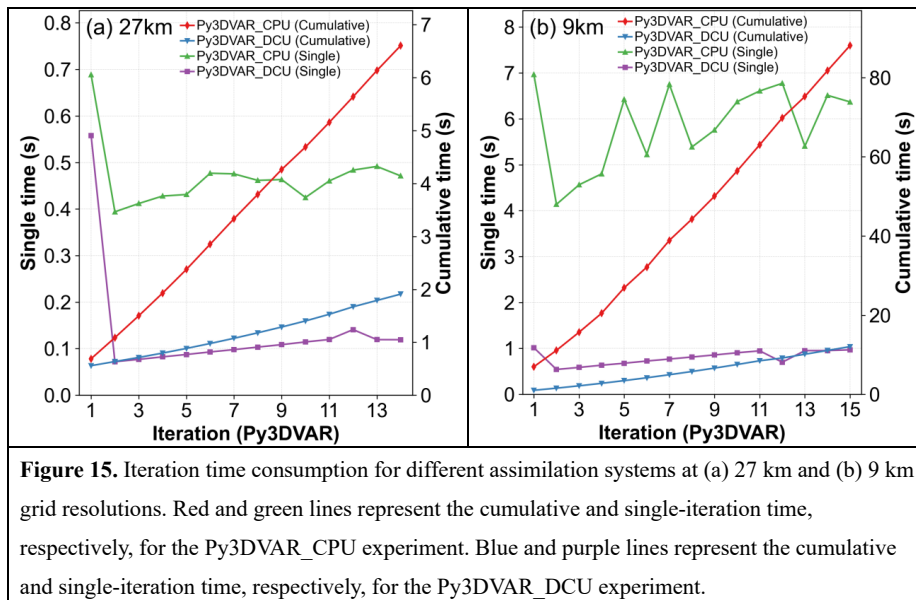
752 Fig. 15 (a) and (b) present the characteristics of single-iteration time consumption and
 753 cumulative time consumption with iteration counts for the Py3DVAR_CPU and Py3DVAR_DCU
 754 experiments at the 27 km and 9 km grid resolutions, respectively. Overall, the iteration time
 755 consumption of the Py3DVAR assimilation system on both computing platforms follows a consistent
 756 pattern: the first iteration requires significantly more time than subsequent iterations. Specifically, the
 757 first iteration time exceeds 0.6 s for the Py3DVAR_CPU experiment and 0.5 s for the
 758 Py3DVAR_DCU experiment. In contrast, the time for subsequent iterations decreases markedly and
 759 stabilizes. For Py3DVAR_CPU, the single-iteration time remains in the range of 0.4 - 0.5 s, while
 760 for Py3DVAR_DCU, it stabilizes around 0.1 s. This phenomenon stems from the computational
 761 characteristics of the PyTorch framework: during the first execution, the automatic differentiation
 762 mechanism needs to construct and initialize the computational graph, incurring additional overhead.
 763 Subsequent iterations can directly reuse the already-built computational graph without repeated
 764 initialization, leading to significantly reduced and stable time consumption. It is noteworthy that the
 765 acceleration advantage of the DCU computing platform is not fully manifested in the first iteration;
 766 its core benefit primarily emerges in subsequent iterations. Compared to the CPU platform, the DCU
 767 substantially reduces the single-iteration time for later iterations due to its parallel computing
 768 architecture, further highlighting the efficiency enhancement of heterogeneous computing for
 769 assimilation.

770 Analyzing the trend of cumulative time consumption, the cumulative time curves for both
 771 experiments exhibit an almost linear increase, indicating that the total computation time is
 772 approximately linearly correlated with the number of iterations. This characteristic reflects the good
 773 computational stability and predictability of the system during the iteration process.

774 In the comparative experiments at the 9 km grid, due to the significantly larger grid size and
 775 expanded computational scale, both single-iteration time and cumulative time increase substantially.



776 However, the overall pattern of time consumption remains consistent with that observed at the 27 km
 777 grid: higher overhead for the first iteration, stable time for subsequent iterations, and an
 778 approximately linear growth in cumulative time. This further demonstrates that Py3DVAR maintains
 779 stable computational characteristics across different resolutions.



780

781 4.6 Comparison of Computational Efficiency for Assimilation Systems at Different Grid 782 Resolutions

783 Fig. 16 (a) and (b) present the comparison of total iteration times for the Fortran-3DVAR,
 784 Py3DVAR_CPU, and Py3DVAR_DCU experiments at the 27 km and 9 km grid resolutions,
 785 respectively.

786 At the 27 km grid, the total iteration times for Fortran-3DVAR, Py3DVAR_CPU, and
 787 Py3DVAR_DCU are 62.5 s, 7.1 s, and 1.9 s, respectively. Although both Fortran-3DVAR and
 788 Py3DVAR_CPU run on the same CPU platform, the latter achieves a speedup of 8.8 times. This
 789 significant performance improvement stems primarily from fundamental differences in their
 790 computational paradigms: Fortran-3DVAR employs loop-based serial computing with manually
 791 partitioned parallel regions, making it difficult to fully exploit the parallel potential of multi-core
 792 CPUs. In contrast, Py3DVAR leverages the vectorized operations and automatic parallel
 793 optimization of the PyTorch framework, enabling efficient utilization of multi-core CPU resources
 794 and thereby substantially enhancing computational efficiency. This demonstrates that Py3DVAR
 795 represents not only a code-level refactoring but also an innovation in computational paradigm.

796 Furthermore, Py3DVAR_DCU, running on the DCU platform, achieves a 3.7 times speedup



797 compared to Py3DVAR_CPU, reducing the total time to under 1.9 s and demonstrating preliminary
 798 potential for near-real-time assimilation. Compared to the traditional Fortran-3DVAR,
 799 Py3DVAR_DCU achieves an even greater speedup of 32.7 times, highlighting its significant
 800 performance advantages within the existing hardware architecture.

801 At the 9 km grid, as the computational scale increases, the time consumption for all systems
 802 rises substantially, but the acceleration pattern exhibits new characteristics: the speedup of
 803 Py3DVAR_CPU relative to Fortran-3DVAR decreases slightly to 5.1 times, while the speedup of
 804 Py3DVAR_DCU relative to Fortran-3DVAR increases further to 40.3 times. The speedup of
 805 Py3DVAR_DCU relative to Py3DVAR_CPU also increases to 7.8 times. This indicates that as the
 806 problem scale expands, reliance solely on multi-core CPU parallelism gradually reaches a bottleneck,
 807 whereas the advantages of massive parallel computing on the DCU platform become more fully
 808 realized. With smaller grid scales, the parallel computing units of the DCU may not be fully utilized,
 809 resulting in a limited performance gap compared to the CPU. However, when the computational
 810 scale increases significantly, the DCU can more efficiently schedule its vast number of computing
 811 cores, thereby exhibiting more substantial acceleration effects.

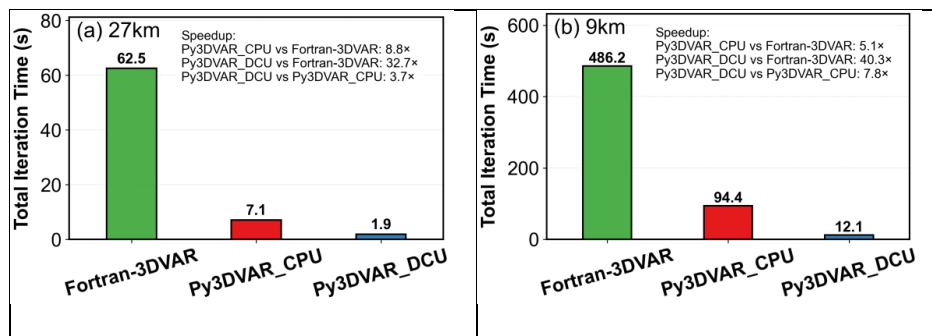


Figure 16. Total L-BFGS gradient descent iteration time and speedup ratios of different experiments for assimilation systems at (a) 27 km and (b) 9 km grid resolutions. Green, red, and blue bars represent the Fortran-3DVAR, Py3DVAR_DCU, and Py3DVAR_CPU experiments, respectively.

812

813 4.7 Comparative analysis of 24-hour forecast performance between Py3DVAR and Fortran- 814 3DVAR at the 27 km grid resolution

815 Fig. 17 and 18 display the temporal evolution of the correlation coefficient (CORR), root mean
 816 square error (RMSE), and mean absolute error (MAE) for various pollutants during the 24-hour
 817 forecast period, initialized from the analysis fields of Py3DVAR and Fortran-3DVAR, respectively.
 818 Comparing the forecast performance of the two systems reveals that the curves of the forecast
 819 metrics for all pollutants are fundamentally consistent in terms of morphological characteristics,
 820 temporal trends, and numerical values. This finding corroborates the previous results showing minor



821 differences in idealized experiment increment fields, similar analysis field effectiveness in the real
822 case, and comparable cost function convergence characteristics during iteration. It fully demonstrates
823 that the Py3DVAR system achieves performance consistent with the mature Fortran-3DVAR system
824 not only in initial field assimilation but also in subsequent forecast applications, thoroughly verifying
825 the reliability and practical utility of its assimilation results.

826 Regarding the evolution of forecast performance, although the initial field is significantly
827 optimized by assimilation, this positive effect decays rapidly as the forecast lead time increases. This
828 decay primarily stems from inherent model uncertainties, spatiotemporal errors in emission
829 inventories, and inaccuracies in the simulated meteorological fields affecting pollutant dispersion and
830 transport. These errors accumulate during the forecast process and gradually become dominant,
831 thereby progressively diminishing the positive effect of the assimilation improvements.

832 It is noteworthy that the duration of the positive assimilation effect varies significantly among
833 different pollutants. The improvement for SO₂, CO, PM_{2.5}, and PM₁₀ can be maintained for over 24
834 hours, whereas the positive effects for NO₂ and O₃ are shorter, lasting approximately 8 hours and 9
835 hours, respectively. This difference is mainly determined by the respective physicochemical
836 properties and dominant life-cycle processes of the pollutants. SO₂, CO, and particulate matter
837 typically have longer atmospheric lifetimes, and their concentration distributions are more influenced
838 by large-scale transport processes. Consequently, improvements in the initial field yield more lasting
839 positive effects. In contrast, NO₂ and O₃ have strong photochemical reactivity. Their concentrations
840 during daytime are more significantly governed by local photochemical processes and boundary
841 layer variations, causing the initial field improvements from assimilation to dissipate quickly.

842 In summary, Py3DVAR not only possesses assimilation capabilities comparable to Fortran-
843 3DVAR but also produces positive effects with good forecast persistence. Although this effect decays
844 with increasing forecast lead time, its improvement for short-term forecasts of various pollutants,
845 especially long-lived species, is clear and effective.




Figure 17. Temporal evolution of the 24-hour forecast performance for the Py3DVAR_DCU assimilation experiment and the Control experiment for the six pollutants (SO₂, NO₂, CO, O₃, PM_{2.5}, PM₁₀). Each column represents one of the three evaluation metrics: correlation coefficient (CORR), root mean square error (RMSE), and mean absolute error (MAE), in sequence. The red lines represent the results of the Py3DVAR_DCU experiment, and the blue lines represent the results of the Control experiment.




Figure 18. Temporal evolution of the 24-hour forecast performance for the Fortran-3DVAR assimilation experiment and the Control experiment for the six pollutants (SO₂, NO₂, CO, O₃, PM_{2.5}, PM₁₀). Each column represents one of the three evaluation metrics: correlation coefficient (CORR), root mean square error (RMSE), and mean absolute error (MAE), in sequence. The red lines represent the results of the Fortran-3DVAR experiment, and the blue lines represent the results of the Control experiment.

847

848 **5. Summary and outlook**

849 Based on the three-dimensional variational (3D-VAR) assimilation principle and integrated with
 850 the PyTorch deep learning framework, this study successfully developed a DCU-accelerated



851 atmospheric chemical three-dimensional variational data assimilation system (Py3DVAR). The
852 system is built upon the CMAQ-AERO5 aerosol scheme coupled with the WRF-Chem model. It
853 consolidates 27 aerosol species into 12 aerosol control variables and directly designates four gaseous
854 pollutants as gaseous assimilation control variables. Simultaneously, the background error covariance
855 matrix is constructed using the NMC method, enabling the simultaneous assimilation of ground-
856 observed concentrations of SO₂, NO₂, CO, O₃, PM_{2.5}, and PM₁₀. At the algorithm implementation
857 level, the core computations of the system replace the traditional Fortran-based loop structures with
858 Torch tensors. Leveraging PyTorch's automatic differentiation mechanism and its built-in L-BFGS
859 optimizer, the system achieves efficient minimization on DCU hardware.

860 To comprehensively evaluate the system's performance and computational efficiency, this study
861 conducted an idealized experiment, multi-platform and multi-resolution real-case assimilation
862 experiments, and a short-term forecast experiment, all compared and validated against the traditional
863 Fortran-3DVAR system. The idealized experiment results demonstrate that Py3DVAR can propagate
864 observation information both horizontally and vertically through the background error covariance,
865 generating Gaussian-distributed increment fields. These fields show high consistency with those
866 from Fortran-3DVAR, verifying the physical soundness of the system.

867 Real-case assimilation experiments further confirm that Py3DVAR significantly optimizes the
868 model's background field, substantially improves the correlation (CORR) between the analysis field
869 and observations, and reduces the root mean square error (RMSE) and mean absolute error (MAE).
870 At the 27 km grid resolution, the Py3DVAR_DCU experiment increased the CORR for SO₂, NO₂,
871 CO, O₃, PM_{2.5}, and PM₁₀ by 0.77, 0.51, 0.71, 0.98, 0.60, and 0.69, respectively. The RMSE was
872 reduced by 6.69 μg·m⁻³ (57.2%), 15.96 μg·m⁻³ (75.1%), 0.30 mg·m⁻³ (56.6%), 30.30 μg·m⁻³
873 (77.8%), 14.47 μg·m⁻³ (54.0%), and 31.70 μg·m⁻³ (62.3%), respectively. The MAE was reduced by
874 4.63 μg·m⁻³ (65.2%), 12.85 μg·m⁻³ (76.3%), 0.24 mg·m⁻³ (64.9%), 24.21 μg·m⁻³ (79.3%), 10.71
875 μg·m⁻³ (54.8%), and 24.43 μg·m⁻³ (63.4%), respectively.

876 At the 9 km grid resolution, the Py3DVAR_DCU experiment increased the CORR for SO₂,
877 NO₂, CO, O₃, PM_{2.5}, and PM₁₀ by 0.78, 0.98, 0.66, 0.96, 0.63, and 0.78, respectively. The RMSE was
878 reduced by 6.83 μg·m⁻³ (60.4%), 26.25 μg·m⁻³ (81.8%), 0.32 mg·m⁻³ (59.3%), 28.39 μg·m⁻³
879 (76.6%), 15.14 μg·m⁻³ (56.5%), and 35.74 μg·m⁻³ (64.4%), respectively. The MAE was reduced by
880 4.63 μg·m⁻³ (65.4%), 21.89 μg·m⁻³ (83.2%), 0.26 mg·m⁻³ (66.7%), 19.71 μg·m⁻³ (77.1%), 11.23
881 μg·m⁻³ (58.0%), and 26.98 μg·m⁻³ (65.4%), respectively.

882 The assimilation accuracy at both grid resolutions is comparable to that of Fortran-3DVAR,
883 while Py3DVAR exhibits exceptional computational efficiency. Even on identical CPU hardware,
884 Py3DVAR achieves an 8.8-fold speedup compared to the original loop-based Fortran-3DVAR due to
885 its use of vectorized PyTorch tensor computations, revealing the inherent efficiency advantages of
886 modern computational frameworks. When deployed on the DCU platform, the acceleration benefits



887 are even more substantial. At the 27 km grid, the total iteration time is reduced from 62.5 s for the
888 Fortran-CPU version to 1.9 s, representing a speedup of 32.7-fold. At the larger-scale 9 km grid, the
889 speedup further increases to 40.3-fold, highlighting the immense potential of DCU's massively
890 parallel computing for handling high-dimensional assimilation problems.

891 The forecast experiment results indicate that the improvements from assimilation have a
892 sustained positive effect on short-term air quality forecasts. The positive effects for SO₂, CO, PM_{2.5},
893 and PM₁₀ last over 24 hours, while those for NO₂ and O₃ last over 6 hours. Furthermore, the decay
894 trends of the forecast improvement effects are consistent between Py3DVAR_DCU and
895 Py3DVAR_CPU, further validating the stability and consistency of the system's assimilation
896 effectiveness.

897 In summary, the Py3DVAR system developed in this study successfully achieves orders-of-
898 magnitude improvements in computational efficiency while preserving the physical interpretability
899 and numerical accuracy of traditional 3D-VAR assimilation. This is accomplished through a
900 synergistic innovation encompassing both the algorithm implementation paradigm (shifting from
901 manual derivation to automatic differentiation) and the computational hardware architecture
902 (transitioning from CPU to DCU). This work not only provides a new solution that combines
903 timeliness and reliability for the operational application of atmospheric chemical data assimilation
904 but also, by being built upon the PyTorch framework, lays a crucial technical foundation for the
905 future integration of deep learning models and the advancement towards intelligent assimilation.

906

907 *Code and data availability.* The WRF-Chem model source code can be downloaded at the WRF
908 model download page (<https://www.mmm.ucar.edu/models/wrf>). The Py3DVAR system was
909 developed by the authors. The exact version of the aerosol Py3DVAR system code and input data for
910 supporting this paper are available at <https://doi.org/10.5281/zenodo.18401481> (You, 2026).

911

912 *Author contributions.* Hancheng Ye conceived the study, developed the Py3DVAR assimilation
913 system, performed all numerical simulations and assimilation experiments, conducted the formal
914 analysis, and wrote the original manuscript. Zengliang Zang and Wei You contributed to the
915 methodology design and assisted in the validation and evaluation of the assimilation system. Yiwen
916 Hu, Ning Liu and Yi Li discussed the assimilation effects of Py3DVAR, provided supervision for the
917 research, and reviewed and edited the manuscript. All authors have read and agreed to the published
918 version of the manuscript.

919

920 *Competing interests.* The contact author has declared that neither they nor their co-authors have any
921 competing interests.



922

923 *Acknowledgements.* We would like to express our sincere gratitude to the National Centers for
924 Environmental Prediction (NCEP) for providing the FNL Operational Global Analysis dataset, and to
925 the China National Environmental Monitoring Center (CNEMC) for making available the
926 observational data of atmospheric pollutants.

927

928 *Financial support.* This research is supported by the National Natural Science Foundation of China
929 (Grant Nos. 42430612, 42475166 and 42405149).

930

931 **References**

- 932 Austin, J.: Toward the four dimensional assimilation of stratospheric chemical constituents, *J.*
933 *Geophys. Res.-Atmos.*, 97, 2569-2588, <https://doi.org/10.1029/91JD02603>, 1992.
- 934 Beauchamp, M., Karagali, I., Gacitúa, G., Høyer, J. L., Ballarotta, M., and Fablet, R.: Multiscale
935 neural assimilation scheme for high-resolution sea surface temperature reconstruction from satellite
936 observations, *Sci Rep*, 15, 39928, <https://doi.org/10.1038/s41598-025-23682-9>, 2025.
- 937 Boudier, P., Fillion, A., Gratton, S., Gürol, S., and Zhang, S.: Data Assimilation Networks, *J Adv*
938 *Model Earth Syst*, 15, e2022MS003353, <https://doi.org/10.1029/2022MS003353>, 2023.
- 939 Chattopadhyay, A., Nabizadeh, E., Bach, E., and Hassanzadeh, P.: Deep learning-enhanced
940 ensemble-based data assimilation for high-dimensional nonlinear dynamical systems, *J. Comput.*
941 *Phys.*, 477, 111918, <https://doi.org/10.1016/j.jcp.2023.111918>, 2023.
- 942 Chen, D., Liu, Z., Davis, C., and Gu, Y.: Dust radiative effects on atmospheric thermodynamics and
943 tropical cyclogenesis over the Atlantic Ocean using WRF-Chem coupled with an AOD data
944 assimilation system, *Atmos. Chem. Phys.*, 17, 7917-7939, <https://doi.org/10.5194/acp-17-7917-2017>,
945 2017.
- 946 Chen, H., Gao, J., Sun, T., Chen, Y., Wang, Y., and Carlin, J. T.: Assimilation of Water Vapor
947 Retrievals from ZDR Columns Using the 3DVar Method for Improving the Short-Term Prediction of
948 Convective Storms, *Mon. Wea. Rev.*, 152, 1077-1095, <https://doi.org/10.1175/MWR-D-23-0196.1>,
949 2024.
- 950 De Luca, P., Galletti, A., Giunta, G., and Marcellino, L.: Recursive filter based GPU algorithms in a
951 Data Assimilation scenario, *J. Comput. Sci.*, 53, 101339, <https://doi.org/10.1016/j.jocs.2021.101339>,
952 2021.
- 953 Denby, B., Schaap, M., Segers, A., Builtjes, P., and Horálek, J.: Comparison of two data assimilation



- 954 methods for assessing PM10 exceedances on the European scale, *Atmos. Environ.*, 42, 7122 – 7134,
955 <https://doi.org/10.1016/j.atmosenv.2008.05.058>, 2008.
- 956 Parrish, D. F. and Derber, J. C.: The National Meteorological Center' s Spectral Statistical-
957 Interpolation Analysis System, *Mon. Wea. Rev.*, 120, 1747-1763, [https://doi.org/10.1175/1520-0493\(1992\)120%3C1747:TNMCSS%3E2.0.CO;2](https://doi.org/10.1175/1520-0493(1992)120%3C1747:TNMCSS%3E2.0.CO;2), 1992.
- 959 Dong, R., Leng, H., Zhao, J., Song, J., and Liang, S.: A Framework for Four-Dimensional Variational
960 Data Assimilation Based on Machine Learning, *Entropy*, 24, 264, <https://doi.org/10.3390/e24020264>,
961 2022.
- 962 Elbern, H., Schmidt, H., and Ebel, A.: Variational data assimilation for tropospheric chemistry
963 modeling, *J. Geophys. Res.-Atmos.*, 102, 15967-15985, <https://doi.org/10.1029/97JD01213>, 1997.
- 964 Engelen, R. J., Serrar, S., and Chevallier, F.: Four-dimensional data assimilation of atmospheric CO2
965 using AIRS observations, *J. Geophys. Res.-Atmos.*, 114, 2008JD010739,
966 <https://doi.org/10.1029/2008JD010739>, 2009.
- 967 Feng, T., G. Li, J. Cao, N. Bei, Z. Shen, W. Zhou, S. Liu, T. Zhang, Y. Li, G., Zhang, R., Fan, J., and
968 Tie, X.: Impacts of black carbon aerosol on photolysis and ozone, *J. Geophys. Res.*, 110,
969 10.1029/2005jd005898, 2005.
- 970 Feng, T., N. Bei, R. Huang, J. Cao, Q. Zhang, W. Zhou, X. Tie, S. Liu, T. Zhang, X. Su, W. Lei, L. T.
971 Molina, and G. Li.: Summertime ozone formation in Xi'an and surrounding areas, China, *Atmos.*
972 *Chem. Phys.*, 16, 4323-4342, 2016.
- 973 Frydendall, J., Brandt, J., and Christensen, J. H.: Implementation and testing of a simple data
974 assimilation algorithm in the regional air pollution forecast model, DEOM, *Atmos. Chem. Phys.*, 9,
975 5475-5488, <https://doi.org/10.5194/acp-9-5475-2009>, 2009.
- 976 Gao, C., Zhang, X., Xiu, A., Tong, Q., Zhao, H., Zhang, S., Yang, G., Zhang, M., and Xie, S.:
977 Intercomparison of multiple two-way coupled meteorology and air quality models (WRF v4.1.1-
978 CMAQ v5.3.1, WRF-Chem v4.1.1, and WRF v3.7.1-CHIMERE v2020r1) in eastern China, *Geosci.*
979 *Model Dev.*, 17, 2471-2492, <https://doi.org/10.5194/gmd-17-2471-2024>, 2024.
- 980 He, X., Zang, Z., Li, Y., Liu, N., Hu, Y., Liu, L., and You, W.: Inversion of a Near-Real-Time China
981 Gridded Hourly SO2 Emission Inventory Using Deep Learning Combined With 3D-Variational
982 Assimilation, *Geophys. Res. Lett.*, 52, e2024GL114375, <https://doi.org/10.1029/2024GL114375>,
983 2025.
- 984 Howard, L. J., Subramanian, A., and Hoteit, I.: A Machine Learning Augmented Data Assimilation
985 Method for High-Resolution Observations, *J Adv Model Earth Syst*, 16, e2023MS003774,
986 <https://doi.org/10.1029/2023MS003774>, 2024.



- 987 Jiang, Z., Liu, Z., Wang, T., Schwartz, C. S., Lin, H., and Jiang, F.: Probing into the impact of
988 3DVAR assimilation of surface PM10 observations over China using process analysis, *JGR*
989 *Atmospheres*, 118, 6738-6749, <https://doi.org/10.1002/jgrd.50495>, 2013.
- 990 Li, C., Fang, H., Cao, X., Duan, D., Xiao, C., Huang, H., Ren, G., Lin, Y., and Cai, Y.: A New
991 Method for Reconstruction of Regional Three-Dimensional Electron Density Distributions Using AI-
992 Based Data Assimilation Method and Incoherent Scatter Radar Measurements, *Geophys. Res. Lett.*,
993 51, e2024GL112352, <https://doi.org/10.1029/2024GL112352>, 2024.
- 994 Li, G., Lei, W., Zavala, M., Volkamer, R., Dusanter, S., Stevens, P., and Molina, L. T.: Impacts of
995 HONO sources on the photochemistry in Mexico City during the MCMA-2006/MILAGO Campaign,
996 *Atmos. Chem. Phys.*, 10, 6551-6567, 10.5194/acp-10-6551-2010, 2010.
- 997 Li, G., Bei, N., Tie, X., and Molina, L. T.: Aerosol effects on the photochemistry in Mexico City
998 during MCMA-2006/MILAGRO campaign, *Atmos. Chem. Phys.*, 11, 5169-5182, 10.5194/acp-11-
999 5169-2011, 2011a.
- 1000 Li, G., Zavala, M., Lei, W., Tsimpidi, A. P., Karydis, V. A., Pandis, S. N., agaratna, M. R., and
1001 Molina, L. T.: Simulations of organic aerosol concentrations in Mexico City using the WRF-CHEM
1002 model during the MCMA-2006/MILAGRO campaign, *Atmos. Chem. Phys.*, 11, 3789-3809,
1003 10.5194/acp-11-3789-2011, 2011b.
- 1004 Li, G., Lei, W., Bei, N., and Molina, L. T.: Contribution of garbage burning to chloride and PM2.5 in
1005 Mexico City, *Atmos. Chem. Phys.*, 12, 8751-8761, 10.5194/acp-12-8751-2012, 2012.
- 1006 Li, H., Yang, T., Kong, L., Zhang, D., Tang, G., and Wang, Z.: OIRF-LEnKF v1.0: A Self-evolving
1007 Data Assimilation System by Integrating Incremental Machine Learning with a Localized EnKF for
1008 Enhanced PM2.5 Chemical Component Forecasting and Analysis, *EGUsphere* [preprint],
1009 <https://doi.org/10.5194/egusphere-2025-3960>, 2025.
- 1010 Li, Z., Zang, Z., Li, Q. B., Chao, Y., Chen, D., Ye, Z., Liu, Y., and Liou, K. N.: A three-dimensional
1011 variational data assimilation system for multiple aerosol species with WRF/Chem and an application
1012 to PM2.5 prediction, *Atmos. Chem. Phys.*, 13, 4265-4278, [https://doi.org/10.5194/acp-13-4265-](https://doi.org/10.5194/acp-13-4265-2013)
1013 2013, 2013.
- 1014 Liang, Y., Zang, Z., Liu, D., Yan, P., Hu, Y., Zhou, Y., and You, W.: Development of a three-
1015 dimensional variational assimilation system for lidar profile data based on a size-resolved aerosol
1016 model in WRF-Chem model v3.9.1 and its application in PM2.5 forecasts across China, *Geosci.*
1017 *Model Dev.*, 13, 6285-6301, <https://doi.org/10.5194/gmd-13-6285-2020>, 2020.
- 1018 Lu, F.: U-Net Kalman Filter (UNetKF): An Example of Machine Learning-Assisted Data
1019 Assimilation, *J Adv Model Earth Syst*, 17, e2023MS003979,
1020 <https://doi.org/10.1029/2023MS003979>, 2025.



- 1021 Ma, X., Liu, H., and Peng, Z.: Improving WRF-Chem PM_{2.5} predictions by combining data
1022 assimilation and deep-learning-based bias correction, *Environ. Int.*, 195, 109199,
1023 <https://doi.org/10.1016/j.envint.2024.109199>, 2025.
- 1024 Pagowski, M., Grell, G. A., McKeen, S. A., Peckham, S. E., and Devenyi, D.: Three-dimensional
1025 variational data assimilation of ozone and fine particulate matter observations: some results using the
1026 Weather Research and Forecasting-Chemistry model and Grid-point Statistical Interpolation, *Quart.*
1027 *J. Royal. Meteor. Soc.*, 136, 2013-2024, <https://doi.org/10.1002/qj.700>, 2010.
- 1028 Pang, J., Liu, Z., Wang, X., Bresch, J., Ban, J., Chen, D., and Kim, J.: Assimilating AOD retrievals
1029 from GOCI and VIIRS to forecast surface PM_{2.5} episodes over Eastern China, *Atmos. Environ.*,
1030 179, 288-304, <https://doi.org/10.1016/j.atmosenv.2018.02.011>, 2018.
- 1031 Pang, J., Wang, X., Shao, M., Chen, W., and Chang, M.: Aerosol optical depth assimilation for a
1032 modal aerosol model: Implementation and application in AOD forecasts over East Asia, *Sci. Tot.*
1033 *Environ.*, 719, 137430, <https://doi.org/10.1016/j.scitotenv.2020.137430>, 2020.
- 1034 Seo, M.-G. and Kim, H. M.: Effect of meteorological data assimilation using 3DVAR on high-
1035 resolution simulations of atmospheric CO₂ concentrations in East Asia, *Atmos. Pollut. Res.*, 14,
1036 101759, <https://doi.org/10.1016/j.apr.2023.101759>, 2023.
- 1037 Tombette, M., Mallet, V., and Sportisse, B.: PM₁₀ data assimilation over Europe with the optimal
1038 interpolation method, *Atmos. Chem. Phys.*, 9, 57-70, <https://doi.org/10.5194/acp-9-57-2009>, 2009.
- 1039 Wang, D., You, W., Zang, Z., Pan, X., Hu, Y., and Liang, Y.: A three-dimensional variational data
1040 assimilation system for aerosol optical properties based on WRF-Chem v4.0: design, development,
1041 and application of assimilating Himawari-8 aerosol observations, *Geosci. Model Dev.*, 15, 1821-
1042 1840, <https://doi.org/10.5194/gmd-15-1821-2022>, 2022.
- 1043 Wei You.: A PyTorch-based three-dimensional variational data assimilation system with automatic
1044 differentiation and DCU acceleration for atmospheric chemicals: development and application to
1045 WRF-Chem, Zenodo [data set], <https://doi.org/10.5281/zenodo.18401481>, 2026.
- 1046 Xing, L., Wu, J., Elser, M., Tong, S., Liu, S., Li, X., Liu, L., Cao, J., Zhou, J., El-Haddad, I., Huang,
1047 R., Ge, M., Tie, X., Prévôt, A. S. H., and Li, G.: Wintertime secondary organic aerosol formation in
1048 Beijing-Tianjin-Hebei (BTH): Contributions of HONO sources and heterogeneous reactions, *Atmos.*
1049 *Chem. Phys. Discuss.*, 2018, 1-25, 10.5194/acp-2018-770, 2018.
- 1050 Xu, H., Duan, Y., and Xu, X.: Exploring the integration of a global AI model with traditional data
1051 assimilation in weather forecasting, *Environ. Res. Lett.*, 19, 124079, <https://doi.org/10.1088/1748-9326/ad93e8>, 2024.
- 1053 Xu, X., Sun, X., Han, W., Zhong, X., Chen, L., Gao, Z., and Li, H.: FuXi-DA: a generalized deep
1054 learning data assimilation framework for assimilating satellite observations, *npj Clim Atmos Sci*, 8,



- 1055 156, <https://doi.org/10.1038/s41612-025-01039-3>, 2025.
- 1056 Ye, H., Pan, X., You, W., Zhu, X., Zang, Z., Wang, D., Zhang, X., Hu, Y., and Jin, S.: Impact of
1057 CALIPSO profile data assimilation on 3-D aerosol improvement in a size-resolved aerosol model,
1058 *Atmos. Res.*, 264, 105877, <https://doi.org/10.1016/j.atmosres.2021.105877>, 2021.
- 1059 Ye, H., You, W., Zang, Z., Pan, X., Wang, D., Zhou, N., Hu, Y., Liang, Y., and Yan, P.: Observing
1060 system simulation experiment (OSSE)-quantitative evaluation of lidar observation networks to
1061 improve 3D aerosol forecasting in China, *Atmos. Res.*, 270, 106069,
1062 <https://doi.org/10.1016/j.atmosres.2022.106069>, 2022.
- 1063 Zang, Z., Li, Z., Pan, X., Hao, Z., and You, W.: Aerosol data assimilation and forecasting
1064 experiments using aircraft and surface observations during CalNex, *Tellus. B.: Chem. Phys. Meteo.*,
1065 68, <https://doi.org/10.3402/tellusb.v68.29812>, 2016.
- 1066 Zang, Z., You, W., Ye, H., Liang, Y., Li, Y., Wang, D., Hu, Y., and Yan, P.: 3DVAR Aerosol Data
1067 Assimilation and Evaluation Using Surface PM_{2.5}, Himawari-8 AOD and CALIPSO Profile
1068 Observations in the North China, *Remote. Sens.*, 14, 4009, <https://doi.org/10.3390/rs14164009>,
1069 2022.
- 1070 Zhang, J., Campbell, J. R., Reid, J. S., Westphal, D. L., Baker, N. L., Campbell, W. F., and Hyer, E.
1071 J.: Evaluating the impact of assimilating CALIOP-derived aerosol extinction profiles on a global
1072 mass transport model: CALIOP DATA ASSIMILATION, *Geophys. Res. Lett.*, 38, n/a-n/a,
1073 <https://doi.org/10.1029/2011GL047737>, 2011.
- 1074 Zhang, J., Spurr, R. J. D., Reid, J. S., Xian, P., Colarco, P. R., Campbell, J. R., Hyer, E. J., and Baker,
1075 N. L.: Development of an Ozone Monitoring Instrument (OMI) aerosol index (AI) data assimilation
1076 scheme for aerosol modeling over bright surfaces – a step toward direct radiance assimilation in
1077 the UV spectrum, *Geosci. Model Dev.*, 14, 27 – 42, <https://doi.org/10.5194/gmd-14-27-2021>, 2021.
- 1078 Zhou, Y., Sun, W., Liu, Z., Gao, L., Chen, D., Feng, J., Zhang, T., and Zhou, Z.: Further development
1079 and application of the WRFDA-Chem three-dimensional variational (3DVAR) system: Joint
1080 assimilation of satellite AOD retrievals and surface observations, *Atmos. Res.*, 316, 107942,
1081 <https://doi.org/10.1016/j.atmosres.2025.107942>, 2025.