



Assimilating compact
phase space
retrievals of
atmospheric
composition

A. P. Mizzi et al.

Assimilating compact phase space retrievals of atmospheric composition with WRF-Chem/DART: a regional chemical transport/ensemble Kalman filter data assimilation system

A. P. Mizzi¹, A. F. Arellano², D. P. Edwards¹, J. L. Anderson³, and G. G. Pfister¹

¹National Center for Atmospheric Research, Atmospheric Chemistry Observation and Modeling Laboratory, Boulder, CO, USA

²University of Arizona, Department of Atmospheric Science, Tucson, AZ, USA

³National Center for Atmospheric Research, Institute for Applied Mathematics, Boulder, CO, USA

Received: 3 June 2015 – Accepted: 20 July 2015 – Published: 8 September 2015

Correspondence to: A. P. Mizzi (mizzi@ucar.edu)

Published by Copernicus Publications on behalf of the European Geosciences Union.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures



Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



Abstract

This paper introduces the Weather Research and Forecasting Model with chemistry/Data Assimilation Research Testbed (WRF-Chem/DART) chemical transport forecasting/data assimilation system together with the assimilation of “compact phase space retrievals” of satellite-derived atmospheric composition products. WRF-Chem is a state-of-the-art chemical transport model. DART is a flexible software environment for researching ensemble data assimilation with different assimilation and forecast model options. DART’s primary assimilation tool is the ensemble adjustment Kalman filter. WRF-Chem/DART is applied to the assimilation of Terra/Measurement of Pollution in the Troposphere (MOPITT) carbon monoxide (CO) trace gas retrieval profiles. Those CO observations are first assimilated as quasi-optimal retrievals (QORs). Our results show that assimilation of the CO retrievals: (i) reduced WRF-Chem’s CO bias in retrieval and state space, and (ii) improved the CO forecast skill by reducing the Root Mean Square Error (RMSE) and increasing the Coefficient of Determination (R^2). Those CO forecast improvements were significant at the 95 % level.

Trace gas retrieval data sets contain: (i) large amounts of data with limited information content per observation, (ii) error covariance cross-correlations, and (iii) contributions from the retrieval prior profile that should be removed before assimilation. Those characteristics present challenges to the assimilation of retrievals. This paper addresses those challenges by introducing the assimilation of “compact phase space retrievals” (CPSRs). CPSRs are obtained by preprocessing retrieval datasets with an algorithm that: (i) compresses the retrieval data, (ii) diagonalizes the error covariance, and (iii) removes the retrieval prior profile contribution. Most modern ensemble assimilation algorithms can efficiently assimilate CPSRs. Our results show that assimilation of MOPITT CO CPSRs reduced the number of observations (and assimilation computation costs) by $\sim 35\%$ while providing CO forecast improvements comparable to or better than with the assimilation of MOPITT CO QORs.

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



1 Introduction

There is increased international interest in chemical weather forecasting (Kukkonen et al., 2012; MACC-II Final Report). Such forecasts rely on coupled forecast model/data assimilation systems that ingest a combination of remotely sensed and in situ atmospheric composition observations together with conventional meteorological observations. Generally the remotely sensed observations come in the form of trace gas retrievals. Examples include carbon monoxide (CO) total and partial column or profile retrievals from the Terra/Measurement of Pollution in the Troposphere (MOPITT) and Metop/Infrared Atmospheric Sounding Interferometer (IASI) instruments. The associated data sets are characterized by large numbers of observations with limited information per observation. Such remotely sensed data has been assimilated in various settings (e.g. Bei et al., 2008; Herron-Thorpe et al., 2012; Klonecki et al., 2012; and Gaubert et al., 2014), but there have been only a few papers addressing data compression strategies. Two such papers were Joiner and Da Silva (1998) and Migliorini et al. (2008). This research is inspired by their research and introduces an efficient assimilation strategy that reduces the number of MOPITT CO retrieval observations by ~35%. Greater reductions are possible, for example with IASI CO and ozone (O₃) retrievals, and depend on the number of: (i) levels in the retrieval profile and (ii) linearly independent pieces of information in the retrieval profile.

Joiner and Da Silva (1998) first proposed the idea of using information content to reduce the number of retrieval observations. They suggested projecting retrievals onto the eigenvectors of the observation error covariance matrix and zeroing those coefficients with little or no information. Their approach evolved from one-dimensional retrieval algorithms (e.g. Twomey, 1974; Smith and Woolf, 1976; and Thompson, 1992).

Retrievals are often obtained using the optimal estimation method of Rodgers (2000) to obtain solutions to the “retrieval equation”

$$\mathbf{y}_r = \mathbf{A}\mathbf{y}_t + (\mathbf{I} - \mathbf{A})\mathbf{y}_a + \boldsymbol{\varepsilon} \quad (1)$$

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

⏪

⏩

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

| | |
|--------------------------|--------------|
| Title Page | |
| Abstract | Introduction |
| Conclusions | References |
| Tables | Figures |
| ◀ | ▶ |
| ◀ | ▶ |
| Back | Close |
| Full Screen / Esc | |
| Printer-friendly Version | |
| Interactive Discussion | |

where \mathbf{y}_r is the retrieval profile, \mathbf{A} is the averaging kernel, \mathbf{y}_t is the true atmospheric profile (unknown), \mathbf{I} is the identity matrix, \mathbf{y}_a is the retrieval prior profile, and $\boldsymbol{\varepsilon}$ is the measurement error in retrieval space with error covariance \mathbf{E}_m . Joiner and Da Silva (1998) proposed projecting Eq. (1) onto the trailing left singular vectors from the singular value decomposition (SVD) of: (i) $(\mathbf{I} - \mathbf{A})$ (their Method 1) or (ii) the smoothing error $\mathbf{E}_s = (\mathbf{A} - \mathbf{I})\mathbf{P}_p(\mathbf{A} - \mathbf{I})^T$ where \mathbf{P}_p is the retrieval prior error covariance (their Method 2). Those projections removed the components of the retrieval that the instrument could not measure with sufficient sensitivity. They called that approach “null-space filtering.”

Joiner and da Silva (1998) recognized that when the retrieval was strongly constrained by the retrieval prior profile the assumptions underlying null-space filtering were invalid. For such retrievals, they proposed filtering based on an eigen-decomposition of \mathbf{E}_m (their PED retrievals). Their analysis showed that PED retrievals were well-conditioned and independent of the retrieval prior profile.

Migliorini et al. (2008) noted that the Joiner and DaSilva (1998) filtering depended on their truncation criteria and was therefore somewhat arbitrary. They also showed it was possible to achieve similar filtering results with an alternative approach that used a more well-defined truncation criterion. Migliorini et al. (2008) rearranged Eq. (1) to obtain

$$\mathbf{y}_r - (\mathbf{I} - \mathbf{A})\mathbf{y}_a = \mathbf{A}\mathbf{y}_t + \boldsymbol{\varepsilon}. \quad (2)$$

Following their terminology we call the left side of Eq. (2) a “quasi-optimal” retrieval (QOR). Migliorini et al. (2008) noted that \mathbf{E}_m was unlikely to be diagonal and likely to be poorly conditioned. To address those issues, they applied a SVD transform to Eq. (2) based on the leading left singular vectors of \mathbf{E}_m (similar to that proposed by Anderson, 2003) – this step provided diagonalization of \mathbf{E}_m and used a well-determined truncation cutoff. They also applied a scaling based on the inverse square root of the associated singular values – this step improved numerical conditioning.

Migliorini et al. (2008) continued to reduce the dimension of $\mathbf{A}\mathbf{y}_t$ (i.e., the number of observations) by neglecting those elements whose variability was smaller than the



Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

⏪

⏩

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

measurement error standard deviation (unity in their rotated and scaled system). They proposed identifying those elements with an eigen-decomposition of the covariance of $\mathbf{A}\mathbf{y}_t$. Since that covariance is generally unknown, they replaced it with the forecast error covariance and showed that the resulting dimension was approximately equal to the number of independent linear functions that could be measured to better than noise level.

In this paper, we propose an approach that achieves similar results without needing to approximate the covariance of $\mathbf{A}\mathbf{y}_t$. Our goal is to compress the retrievals and remove those components that are not dependent on the measurements. In so doing we expect to make the assimilation of retrievals more computationally efficient. The rest of this paper is organized as follows: Sect. 2 introduces “compact phase space retrievals,” and Sect. 3 introduces the WRF-Chem/DART regional chemical weather forecast/data assimilation system. Section 4 discusses our experimental design including the study period, model domain, initial/boundary conditions, and relevant WRF-Chem/DART parameter settings. Section 5 discusses the observations that were assimilated in the various experiments described in Sect. 6. The results from those experiments are presented in Sect. 7, and we end with a summary of our thoughts and conclusions in Sect. 8.

2 Assimilation of compact phase space retrievals

We can rewrite Eq. (2) as

$$\mathbf{y}_r - (\mathbf{I} - \mathbf{A})\mathbf{y}_a - \boldsymbol{\varepsilon} = \mathbf{A}\mathbf{y}_t. \quad (3)$$

In Eq. (3) the averaging kernel \mathbf{A} is singular and has low rank. Therefore the information content for each component of $\mathbf{A}\mathbf{y}_t$ is relatively small, and the assimilation is inefficient because one must assimilate the entire profile to get the same information as can be compressed into a number of processed observations equal to the rank of \mathbf{A} . To compress Eq. (3) we propose transforming it with the leading left singular vectors from

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

⏪

⏩

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

a SVD of the averaging kernel i.e., $\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^T$. We denote the truncated system with the subscript zero so that $\mathbf{A}_0 = \mathbf{U}_0\mathbf{S}_0\mathbf{V}_0^T$, the truncated averaging kernel is obtained by setting the trailing singular values and vectors to zero. The transformed system has the form

$$5 \quad \mathbf{U}_0^T(\mathbf{y}_r - (\mathbf{I} - \mathbf{A})\mathbf{y}_a - \boldsymbol{\varepsilon}) = \mathbf{S}_0\mathbf{V}_0^T\mathbf{y}_t. \quad (4)$$

Migliorini et al. (2008) showed that subtracting $(\mathbf{I} - \mathbf{A})\mathbf{y}_a$ from Eq. (3) removes all contribution from the retrieval prior profile. Equations (3) and (4) confirm their result because the leading left singular vectors of \mathbf{A} span its range so the left side of Eq. (3) should project completely onto \mathbf{U}_0^T . Following that transform \mathbf{E}_m becomes $\mathbf{U}_0^T\mathbf{E}_m\mathbf{U}_0$ which may still be non-diagonal and poorly conditioned. Therefore, we apply an SVD transform and inverse scaling similar to that used by Migliorini et al. (2008). If the SVD of $\mathbf{U}_0^T\mathbf{E}_m\mathbf{U}_0$ has the form $\mathbf{U}_0^T\mathbf{E}_m\mathbf{U}_0 = \Phi\Sigma\Psi$, the transformed and conditioned form of Eq. (4) is

$$10 \quad \Sigma^{-1/2}\Phi^T\mathbf{U}_0^T(\mathbf{y}_r - (\mathbf{I} - \mathbf{A})\mathbf{y}_a - \boldsymbol{\varepsilon}) = \Sigma^{-1/2}\Phi^T\mathbf{S}_0\mathbf{V}_0^T\mathbf{y}_t. \quad (5)$$

This approach compresses Eq. (3) so that the dimension of the “compact phase space retrieval” (CPSR) profile on the left side of Eq. (5) is identical to number of independent linear functions of the atmospheric profile to which the instrument is sensitive. It differs from Migliorini et al. (2008) in that it does not require an eigen-decomposition of the forecast error covariance and does not rely on comparison of the eigenvalues and the rotated/scaled observation error variance to compress the system. As with Migliorini et al. (2008) the final form of the observation error covariance is the truncated identity matrix.

The assimilation of CPSRs should produce results similar to the assimilation of QORs except for: (i) the effect of assimilation sub-processes like horizontal localization and inflation, and (ii) differences in the observation error due to the CPSR compression transform. The QOR and CPSR observation errors are different because the compression transform projects the errors onto the leading left singular vectors of the averaging kernel and retains only those components that lie in the range of the averaging kernel.

In summary the steps for obtaining CPSRs from trace-gas retrievals are as follows (assuming the retrieval equation has the same form as Eq. 2):

- Obtain the retrieval and retrieval prior profiles, the averaging kernel, and the observation error covariance for a particular horizontal location.
- Subtract the retrieval prior term $(\mathbf{I} - \mathbf{A})\mathbf{y}_a$ from the retrieval profile \mathbf{y}_r . This yields the QOR as defined by Eq. (3).
- Perform a SVD of the averaging kernel. Form a transform matrix from the left singular vectors associated with the non-zero singular values. Left multiply the QOR by the transpose of the transform matrix. This yields the truncated QOR profile as defined by Eq. (4).
- Left multiply the observation error covariance by the transpose of the transform matrix. Right multiply that matrix product by the transform matrix. This yields the truncated observation error covariance.
- Perform a SVD of the truncated observation error covariance. Scale the left singular vectors with the inverse square root of their respective singular values. Left multiply the truncated QOR profile by the transpose of the scaled left singular vector matrix. This yields the CPSR profile as defined by Eq. (5).
- As a check left multiply the truncated observation error covariance by the transpose of the scaled left singular vector matrix. Right multiply that matrix product by the scaled left singular vector matrix. The result should be an identity matrix with rank equal to the number of non-zero CPSRs from the previous step.
- Assimilate the non-zero CPSRs with unitary error variance.

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

⏪

⏩

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



3 WRF-Chem/DART – a regional chemical transport/data assimilation system

The Weather Research Forecasting Model with chemistry/Data Assimilation Research Testbed (WRF-Chem/DART) system is the WRF-Chem chemical transport model (www2.acd.ucar.edu/wrf-chem) coupled with the DART (www.image.ucar.edu/DARes/DART) ensemble adjustment Kalman filter (Anderson, 2001, 2003) data assimilation system. WRF-Chem/DART is an extension of WRF/DART (www.cawcr.gov.au and references therein). WRF-Chem is the National Center for Atmospheric Research (NCAR) regional Weather Research and Forecasting (WRF) model (www.wrf-model.org) with chemistry. WRF-Chem is a regional model that predicts conventional weather together with the emission, transport, mixing, and chemical transformation of atmospheric trace gasses and aerosols. WRF-Chem is collaboratively developed and maintained by the National Oceanic and Atmospheric Administration/Earth System Research Laboratory (NOAA/ESRL), Pacific Northwest National Laboratory (PNNL), and NCAR/Atmospheric Chemistry Observation and Modeling Laboratory (ACOM). WRF-Chem is documented in Grell et al. (2004), discussed in Kukkonen et al. (2012), and has been applied in various research settings (e.g., Pfister et al., 2011, 2013).

DART (Anderson et al., 2009) is a community resource for ensemble data assimilation (DA) research developed and maintained by the NCAR/Data Assimilation Research Section (DARes). DART is a flexible software environment for studying the interaction between different assimilation methods, observation platforms, and forecast models. WRF-Chem and DART are state-of-the-art tools for studying the impact of assimilating trace gas retrievals on conventional and chemical weather analyses and forecasts.

GMDD

8, 7693–7725, 2015

Assimilating compact
phase space
retrievals of
atmospheric
composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

⏪

⏩

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

4 Study period, domain, initial conditions, boundary conditions, emissions, and initial ensemble generation

We conducted continuous cycling experiments with WRF-Chem/DART for the period of 00:00 UTC, 1 June 2008 to 00:00 UTC, 1 July 2008 with 6 h cycling (00:00 UTC, 06:00 UTC, 12:00 UTC, and 18:00 UTC). To facilitate a large number of experiments, we used a reduced ensemble of 20 members and a horizontal resolution of 100 km (101 × 41 grid points). We used 34 vertical grid points with a model top at 10 hPa and ~ 15 grid points below 500 hPa. WRF-Chem ran with the Model for Ozone and Related Chemical Tracers (MOZART-4) chemistry and Goddard Chemistry Aerosol Radiation and Transport (GOCART) model aerosol options (Colarco et al., 2009; Emmons et al., 2010). Ideally for chemical transport forecast experiments we would like an ensemble size of at least 40 members, a horizontal resolution of no larger than 20 km, and a vertical grid with at least 50 grid points. We expect our small ensemble/coarse resolution cycling results, as they pertain to the assimilation of QORs and CPSRs, will apply to larger ensembles with higher resolutions. However, as the vertical resolution increases, the sensitivity to vertical localization may increase so that tuning of the vertical localization length would be necessary.

We used NCEP Global Forecast Model (GFS) 0.5° six-hour forecasts for the WRF-Chem initial/boundary conditions. Our model domain extends from ~ 176 to ~ 50° W and from ~ 7 to ~ 54° N. We used the WRF Preprocessing System (WPS) to interpolate the GFS forecasts to our domain and generate the deterministic boundary conditions. We used the WRF Data Assimilation System (WRFDA) (http://www2.mmm.ucar.edu/wrf/users/wrfda/Docs/user_guide_V3.7/WRFDA_Users_Guide.pdf) to generate the initial meteorology ensemble.

For the chemistry initial and lateral boundary conditions, we used global simulations from the NCAR MOZART-4 model. The fire emissions came from the Fire Inventory from NCAR (FINNv1; Wiedinmyer et al., 2011), and the Model of Emissions of Gases and Aerosols from Nature (MEGAN; Guenther et al., 2012) calculated the

GMDD

8, 7693–7725, 2015

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



biogenic emissions as part of the WRF-Chem forecast. The anthropogenic emissions were based on the US Environmental Protection Agency's (EPA's) 2005 National Emissions Inventory (NEI-2005). We used or adapted existing ACOM/WRF-Chem utilities (<https://www2.acd.ucar.edu/wrf-chem/wrf-chem-tools-community>) to generate the initial chemistry ensembles with a Gaussian distribution from a specified mean and standard deviation. That distribution was truncated at the tails to include 95 % of the distribution. Similar utilities were used to generate the emission ensembles. We excluded the distribution tails to avoid the potential for the extreme values to cause numerical problems in the chemistry algorithms. Although we recognize that the assimilation cycling results may be sensitive to the emission perturbation horizontal correlation lengths (e.g. Pagowski et al., 2012), that issue was not particularly relevant to our study so we set the horizontal correlation lengths to zero.

5 Meteorology observations and satellite trace gas retrievals

At each cycle time, depending on the experiment we assimilated meteorology and/or chemistry observations with DART and then advanced the analysis ensemble to the next cycle time with WRF-Chem. The 6 h forecast ensemble was then used as the first guess for the next ensemble DA step.

We assimilated conventional meteorological observations and CO trace gas retrievals from MOPITT. The meteorological observations were NCEP Automated Data Processing (ADP) upper air and surface observations (PREPBUFR observations). They included air temperature, sea level pressure, surface winds, dew point temperature, sea surface temperature, and upper level winds from various observing platforms. We refer to those observations as the "MET OBS."

We also assimilated MOPITT partial column/profile CO retrievals. MOPITT is an instrument flying on NASA's Earth Observing System Terra spacecraft. MOPITT's spatial resolution is 22 km at nadir, and it sees the earth in 640 km wide swaths. MOPITT uses gas correlation spectroscopy to measure CO in a thermal-infrared (TIR) band

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

⏪

⏩

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

near 4.7 μm and a near-infrared (NIR) band near 2.3 μm . TIR radiances are sensitive to CO in the mid- and upper-troposphere while NIR measures the CO total column. Worden et al. (2010) and Deeter et al. (2011, 2012, and 2013) have shown that the sensitivity to CO in the lower troposphere is significantly greater for retrievals exploiting simultaneous TIR and NIR than for retrievals based on TIR alone. MOPITT started data collection in March 2000. We used the MOPITT v5 TIR/NIR products described in Deeter et al. (2013). We refer to the MOPITT observations as the “CHEM OBS.”

The retrieval error covariance \mathbf{C}_r associated with each MOPITT CO retrieval profile is provided as part of the data product. That error variance is derived as part of the retrieval process based on a specified a priori error covariance \mathbf{C}_a . Under the optimal estimation theory of Rodgers (2000) \mathbf{C}_r is related to \mathbf{C}_a through the averaging kernel \mathbf{A} by $\mathbf{C}_r = (\mathbf{I} - \mathbf{A})\mathbf{C}_a$. The measurement error in retrieval space \mathbf{C}_m is also related to \mathbf{C}_a to \mathbf{A} by $\mathbf{C}_m = (\mathbf{I} - \mathbf{A})\mathbf{C}_a(\mathbf{I} + (\mathbf{A} - \mathbf{I})^T)$. Generally for retrieval data sets \mathbf{C}_a , \mathbf{C}_r , and \mathbf{C}_m are non-diagonal.

6 Experimental design

We conducted two basic experiments: (i) a control experiment where we assimilated only MET OBS (MET DA); and (ii) a chemical data assimilation experiment where we assimilated MET OBS and MOPITT CO partial column retrievals in the form of QORs (MOP QOR). In addition we conducted an experiment where we converted the CHEM OBS to CPSRs and assimilated the CPSRs (MOP CPSR). We also conducted sensitivity experiments where we: (i) zeroed the observation error covariance cross-correlations (as opposed to using a SVD transformation for diagonalization), and (ii) applied vertical localization. Those experiments are denoted MOP NROT and MOP LOC respectively. The suite of experiments is summarized in Table 1.

For all experiments we used: (i) DART horizontal localization, (ii) DART prior adaptive inflation, (iii) no posterior inflation, (iv) full interaction between all observations and all state variables – i.e., MET OBS update chemistry state variables and CHEM OBS up-

date meteorology state variables (joint assimilation of MET and CHEM OBS), (v) DART clamping on chemistry state variables to constrain the posterior ensemble members to be positive, and (vi) the reported MOPITT retrieval error covariance as the observation error covariance.

5 For the MOP QOR experiment, the MOPITT CO retrievals were converted to QORs using an algorithm similar to that described by Migliorini et al. (2008) except we did not perform their second forecast error covariance based filtering.

7 Results

7.1 The control and chemical data assimilation experiments

10 The MET DA and MOP QOR experiments are intended to identify the impact of assimilating chemistry observations. Figure 1a shows shaded contours of CO in parts per billion (ppb) at ~ 1000 hPa from the 6 h forecast valid at 00:00 UTC, 29 June 2008 and compares the MET DA and MOP QOR experiments. It shows that over the course of MOP QOR the assimilation of CO retrievals reduced the: (i) positive CO bias found in polluted areas of MET DA (i.e., metropolitan areas with high CO emissions – San Francisco, Los Angeles, Chicago, and the northeast US), and (ii) negative CO bias found in nonpolluted areas in MET DA (Hawaii, east Pacific, southeast US, and Baja). The MET DA biases could result from model errors such as: (i) emission errors – CO emissions too high in polluted areas and too low in nonpolluted areas, (ii) transport errors – insufficient CO transport away from polluted areas and insufficient transport toward nonpolluted areas, and/or (iii) chemistry errors – CO destruction is too weak in polluted areas and too strong in nonpolluted areas. The MET DA biases could also result from initial/boundary conditions errors that were corrected by the assimilation of MOPITT CO in MOP QOR.

25 Figure 1b shows the assimilated CO retrievals for the 18:00 UTC, 28 June 2008 update cycle in the upper panel and the corresponding increments in the lower panel.

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



Comparison of those panels shows that the assimilation step adjusted the CO concentrations primarily along the satellite observation paths, which is a consequence of assimilation sparse observations. The DA adjustments in Fig. 1b are generally consistent with the differences between MOP QOR and MET DA in Fig. 1a (CO increases in nonpolluted areas – east of San Francisco, the southeast US, and Baja).

Figure 2 shows time series of the domain average CO from the MET DA and MOP QOR experiments in retrieval and state space. The dots represent the retrieval space results where the cool colors (blue and black) show the forecasts, and the warm colors (red and magenta) show the analyses. The green dots represent the MOPITT retrievals. The solid lines show state space results. Figure 2 has several interesting results. First, MET DA had a negative bias of ~ 10 ppb in retrieval space. Second, assimilation of MOPITT CO reduced that bias by ~ 5 ppb. Finally, in state space MOP QOR increased the mean CO by ~ 5 ppb. Those results are consistent with Fig. 1 which shows a large number of nonpolluted areas in MET DA with a negative bias and a small number of polluted areas with a positive bias.

Figure 3 shows vertical profiles of the time (00:00 UTC, 25 June 2008 to 00:00 UTC, 29 June 2008) and horizontal domain average CO in retrieval space. It shows that the MOPITT profile had greater vertical variability (moderate CO near the surface, low CO in the mid-troposphere (500–400 hPa), high CO in the upper troposphere (300–200 hPa), and low CO near the tropopause (200–100 hPa)) than the MET DA and MOP QOR profiles. It also shows that the assimilation of MOPITT CO had positive impacts throughout the troposphere with greatest improvement in the upper troposphere. Figure 3 shows that there were differences in the MOP QOR/MET DA bias reduction between: (i) the upper and lower troposphere (greater magnitude negative bias reduction in the upper troposphere and lesser magnitude positive bias reduction in the lower troposphere), and (ii) the forecast and the analysis (greater bias reduction in the analysis than in the forecast). Those results expand our understanding of the bias in Figs. 1 and 2. In Fig. 3 the forecast and analysis show greater bias reduction in the upper troposphere. That suggests that the domain averages in Fig. 2 were dominated by bias

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



and upper troposphere and a minimum in the middle troposphere. The green curves represent the truncated profiles which are obtained by: (i) projecting the full retrieval profile or the QOR profile onto the leading left singular vectors of the associated averaging kernel to get the projection coefficients (e.g., $\mathbf{c}_r = \mathbf{U}_0^T \mathbf{y}_r$ where \mathbf{c}_r is the projection coefficient vector for the full retrieval and $\mathbf{c}_{\text{qor}} = \mathbf{U}_0^T (\mathbf{y}_r - (\mathbf{I} - \mathbf{A}) \mathbf{y}_a - \boldsymbol{\varepsilon}) = \mathbf{U}_0^T \mathbf{y}_{\text{qor}}$, where \mathbf{c}_{qor} is the coefficient vector for the QOR profile – see Eq. (4) in Sect. 2), and (ii) performing the inverse projection by multiplying the leading singular vectors by their respective projection coefficients and summing those dot-products (e.g., $\hat{\mathbf{y}}_r = \mathbf{c}_r \mathbf{U}_0$ is the truncated retrieval profile – denoted MOP-Trc and $\hat{\mathbf{y}}_{\text{qor}} = \mathbf{c}_{\text{qor}} \mathbf{U}_0$ is the truncated QOR profile – denoted QOR-Trc). The forward transform in (i) is analogous to the first part of the CPSR transform in Eq. (4). The inverse transform in (ii) brings the result of forward transform in (i) back to state space. The inverse transform is not part of the CPSR algorithm.

In Fig. 4a the residuals are defined as the difference between the full and truncated profiles (e.g., $\mathbf{y}_r - \hat{\mathbf{y}}_r$ is the full retrieval residual – denoted MOP-Res and $\mathbf{y}_{\text{qor}} - \hat{\mathbf{y}}_{\text{qor}}$ is the QOR residual – denoted QOR-Res). If the full profiles project completely onto the leading singular vectors, the residuals are zero. The upper panel of Fig. 4a shows that the transform in (i) has the greatest impact near the surface and the upper troposphere and the least impact in the middle troposphere. When the truncation residuals are nonzero, the original profiles contain components that are not in the range of the averaging kernel. That can also indicate a contribution from the retrieval prior term $((\mathbf{A} - \mathbf{I}) \mathbf{y}_a)$ but not always. When components of the retrieval prior term lie in the range of the averaging kernel, they cannot be removed by the transform in (i) and are therefore not included in the residual. For example in the upper panel of Fig. 4a the similarity between MOP-Trc and QOR-Ret shows that the MOPITT retrieval was strongly influenced by the retrieval prior and that most of the prior contribution was removed by the transform in (i). That also shows that most of the prior contribution was not in the range of the averaging kernel. However not all was outside the range, and the difference between MOP-Trc and QOR-Ret shows that most was inside the range. This analysis shows that the influence

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

of the retrieval prior term cannot be completely removed by projecting the retrieval onto the range of the averaging kernel. The results show that it is necessary to use the Migliorini et al. (2008) quasi-optimal subtraction in Eq. (2) to remove the retrieval prior contribution. Comparison of QOR-Ret and QOR-Trc in the lower panel of Fig. 4a shows that QOR-Ret lies completely within the range of the averaging kernel. That result was expected from the discussion of Eqs. (3) and (4).

In summary Fig. 4a shows the state space impacts from applying the Migliorini et al. (2008) quasi-optimal subtraction and the CPSR transform in (i). It also shows that the quasi-optimal subtraction was necessary to remove the influence of the retrieval prior. Thus in CPSRs the quasi-optimal subtraction removes the influence of the retrieval prior and projection onto the leading singular vectors of the averaging kernel provides the data compression.

In Fig. 4a the average number of leading singular vectors was ~ 2.3 . CPSRs reduced the number of observations by ~ 7.7 per MOPITT profile. After thinning there were $\sim 30\,000$ MOPITT profiles per assimilation cycle. That implies a CPSR reduction of $\sim 281\,000$ retrievals or $\sim 80\%$ per cycle. On application the actual reduction was less because the number of non-retrieval observations was not reduced. As an example when assimilating MET OBS and CHEM OBS we found a reduction of $\sim 35\%$ in the computation cost. That is a wall clock time reduction based on NCAR's 1.5-petaflop high-performance IBM Yellowstone computer with 32 tasks, and 8 tasks per node. We expect similar reductions for other computing configurations.

In Fig. 4b we examine the vertical structure of the CPSRs. The upper panel shows the retrieval domain average of the MOPITT averaging kernel profiles from 18:00 UTC, 28 June 2008. The lower panel shows the domain average of the leading left singular vectors from SVDs of the averaging kernel profiles in the upper panel. Comparison of the upper and lower panels shows that while the singular vector and averaging kernel profiles are similar it is not possible to associate a specific singular vector profile with a specific averaging kernel profile or with a group of profiles. However like the averaging kernel of singular vectors show the sensitivity of the associated CPSR to the true CO

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

⏪

⏩

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

parison of the vertical profiles in Figs. 3 and 7, which shows that MOP CPSR produced greater bias reductions for the forecast and analysis throughout the troposphere. As in MOP QOR the MOP CPSR improvements were greater in the upper troposphere than in the lower troposphere. The MOP CPSR results from Fig. 7 provide further support for our suggestion that the domain average bias reductions in Figs. 2 and 6 were due to bias reductions in the upper troposphere because the greater bias reductions in the upper troposphere of Fig. 7 (compared to Fig. 3) provided greater bias reductions in Fig. 6 (compared to Fig. 2). In summary Figs. 5–7 confirm our analysis of Figs. 1–3 and show that assimilation of CPSRs produced results that were similar to or better than those from the assimilation of QORs at two-thirds the computational cost. We also conducted significance testing for MOP CPSR similar to that for MOP QOR and were able to reject the null hypothesis that there was no difference between the MOP CPSR and MET DA time series in Fig. 6.

7.3 Verification against MOPITT retrievals

We calculated verification statistics (bias, root mean square error (RMSE), and Coefficient of Determination (R^2)) for the 6 h forecasts from all experiments based on the time series results in Figs. 2 and 6. Those statistics are plotted in Fig. 8. Generally Fig. 8 shows that the assimilation of MOPITT CO improved model performance for all metrics when compared against the MOPITT retrievals. Figure 8 also shows that RMSE was dominated by the bias and that the differences in the statistics for the different treatments (assimilating QORs, CPSRs, cross-covariance zeroing, and vertical localization) were generally negligible except for cross-covariance zeroing. We now discuss the cross-covariance zeroing and vertical localization experiments.

formance. Specifically, MOP QOR significantly improved the WRF-Chem CO forecast skill for all three metrics.

Next we focused on making the assimilation of retrievals computationally efficient and introduced “compact phase space retrievals.” CPSRs advance the work of Joiner and DaSilva (1998) and Migliorini et al. (2008) by describing an easily applied methodology to achieve data compression for phase space retrievals. Conceptually, the assimilation of QORs and CPSRs should yield similar results except for the effects of: (i) assimilation sub-processes like localization and inflation and (ii) different observation errors due to the CPSR compression transform. MOP CPSR maps in Fig. 5 showed that assimilation of CPSRs placed CO hot spots in the same locations as MOP QOR but they were of greater magnitude and larger area. The time series and vertical profile plots showed that those differences generally represented analysis and forecast improvements. Skill metrics for MOP CPSR showed that when compared to MOP QOR, the assimilation of CPSRs slightly improved the forecast skill for all metrics, and when compared to MET DA it significantly improved the forecast skill for all metrics. Based on those results we conclude that the assimilation of CPSRs performed as well or better than the assimilation of QORs at a substantially reduced computational cost (~35% reduction in computation time).

Collectively our analysis of the MOP QOR and CPSR results in Figs. 1–3 and 5–7 suggested that: (i) in the lower troposphere MET DA had a negative CO bias in polluted areas and a positive bias in nonpolluted areas (Figs. 1 and 5) and (ii) bias reductions in the domain average retrieval space CO were due to reductions in the negative CO bias in the upper troposphere (Figs. 2, 3, 6, and 7). We proposed three causes for the CO biases: (i) emission errors – overestimation of CO emissions in polluted areas and underestimation in nonpolluted areas, (ii) transport errors – too much CO transport from the near surface to the lower troposphere and too little transport from the lower to upper troposphere, and (iii) chemistry errors – too little CO destruction in the near surface and lower troposphere and too much destruction in the upper troposphere.

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion



Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

weather forecasting models in Europe, *Atmos. Chem. Phys.*, 12, 1–87, doi:10.5194/acp-12-1-2012, 2012.

MACC-II Final Report: Monitoring Atmospheric Composition and Climate – Interim Implementation, available at: <https://www.wmo.int/pages/prog/arep/gaw/documents/GAW-2013-Peuch-MACCII.pdf>, October 2014.

Migliorini, S., Piccolo, C., and Rodgers, C. D.: Use of the information content in satellite measurements for an efficient interface to data assimilation, *Mon. Weather Rev.*, 136, 2633–2650, 2008.

Pagowski, M. and Grell, G. A.: Experiments with the assimilation of fine aerosols using an ensemble Kalman filter, *J. Geophys. Res.*, 117, D21302, doi:10.1029/2012JD018333, 2012.

Pfister, G. G., Avise, J., Wiedinmyer, C., Edwards, D. P., Emmons, L. K., Diskin, G. D., Podolske, J., and Wisthaler, A.: CO source contribution analysis for California during ARCTAS-CARB, *Atmos. Chem. Phys.*, 11, 7515–7532, doi:10.5194/acp-11-7515-2011, 2011.

Pfister, G. G., Walters, S., Emmons, L. K., Edwards, D. P., and Avise, J.: Quantifying the contribution of inflow on surface ozone over California during summer 2008, *J. Geophys. Res. Atmos.*, 118, 12282–12299, doi:10.1002/2013JD020336, 2013.

Rodgers, C. D.: *Inverse Methods for Atmospheric Sounding, Theory and Practice*, World Sci., Singapore, 2000.

Smith, W. L. and Woolf, H.: The use of eigenvectors of statistical covariance matrices for interpreting satellite sounding radiometer observations, *J. Atmos. Sci.*, 33, 1–7, 1976.

Thompson, O.: Regularizing the satellite temperature-retrieval problem through singular-value decomposition of the radiative transfer physics, *Mon. Weather Rev.*, 120 2314–2328, 1992.

Twomey, S.: Information content in remote sensing, *Appl. Opt.*, 13, 942–945, 1974.

Wiedinmyer, C., Akagi, S. K., Yokelson, R. J., Emmons, L. K., Al-Saadi, J. A., Orlando, J. J., and Soja, A. J.: The Fire INventory from NCAR (FINN): a high resolution global model to estimate the emissions from open burning, *Geosci. Model Dev.*, 4, 625–641, doi:10.5194/gmd-4-625-2011, 2011.

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

Table 1. Summary of the WRF-Chem/DART Forecast/Data Assimilation Experiments.

| Experiment | Assimilate meteorology observations | Assimilate MOPITT CO QORs | Assimilate MOPITT CO CPSRs | Use error covariance zeroing | Use vertical localization |
|------------|-------------------------------------|---------------------------|----------------------------|------------------------------|---------------------------|
| MET DA | Yes | No | No | No | No |
| MOP QOR | Yes | Yes | No | No | No |
| MOP CPSR | Yes | No | Yes | No | No |
| MOP NROT | Yes | Yes | No | Yes | No |
| MOP LOC | Yes | Yes | No | No | Yes |

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

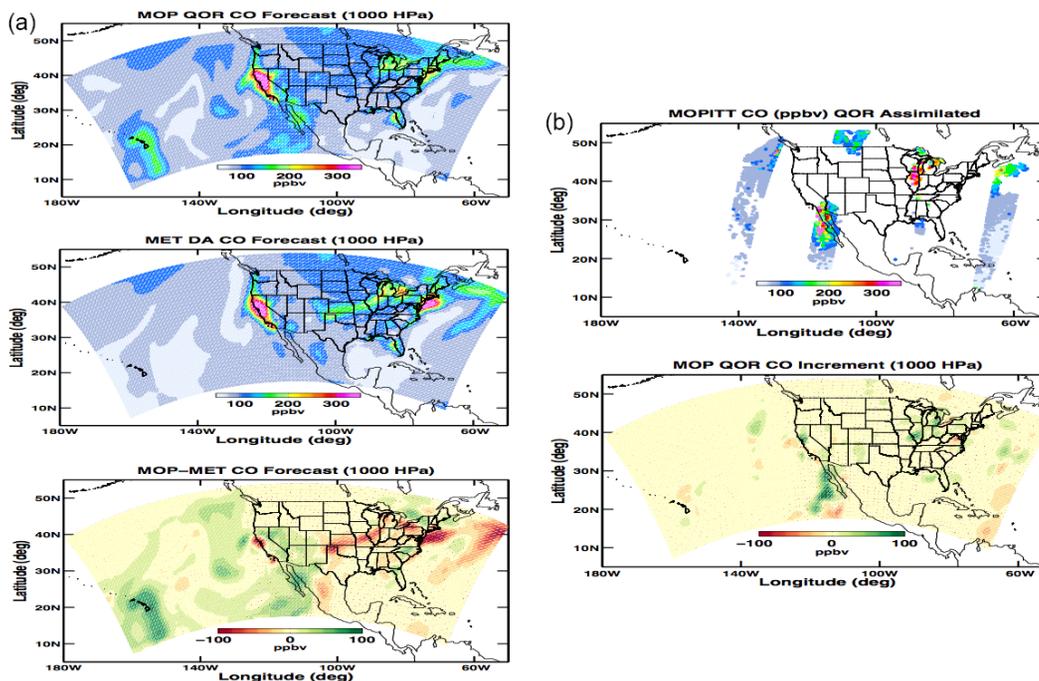


Figure 1. (a) Shaded contours of CO in ppbv for the MOP QOR (upper panel) and MET DA (middle panel) experiments for the first model level above the surface (~ 1000 hPa) from the 6 h forecast valid on 18:00 UTC, 28 June 2008. The lower panel shows the difference contours for those experiments (MOP QOR – MET DA). The shaded area represents the WRF-Chem. (b) The upper panel shows the assimilated MOPITT CO retrievals between the surface and 900 hPa for 18:00 UTC, 28 June 2008. The lower panel shows the associated assimilation increment.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Assimilating compact
phase space
retrievals of
atmospheric
composition

A. P. Mizzi et al.

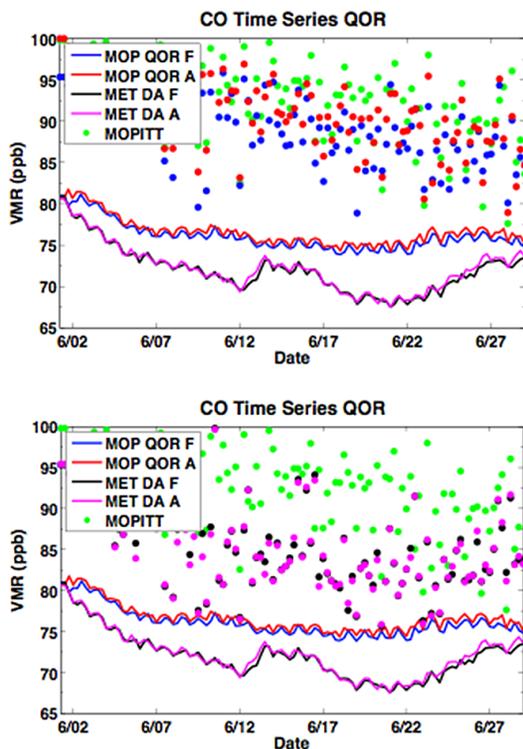


Figure 2. Time series of domain average CO from the MOP QOR and MET DA experiments. The red and magenta dots show the MOP QOR and MET DA analyses (A) in retrieval space, and the blue and black dots show the respective forecasts (F) in retrieval space. The green dots represent the MOPITT retrievals. The lines represent the results in state space.

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

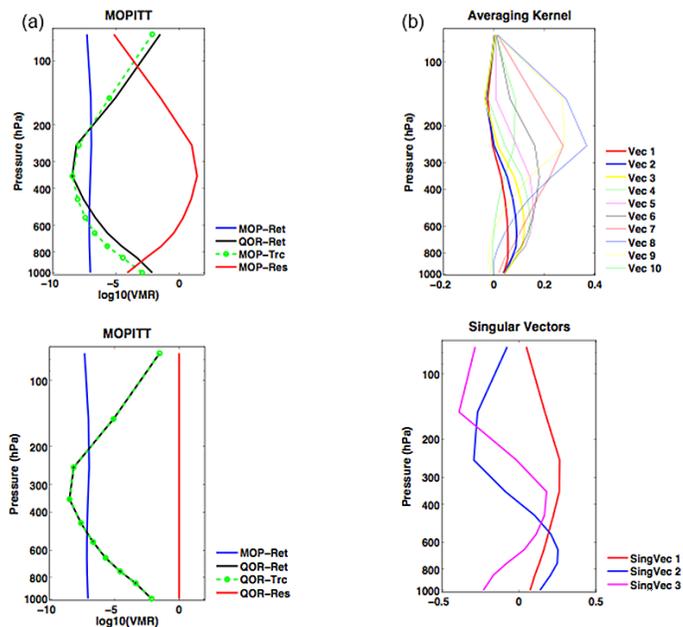


Figure 4. (a) Horizontal domain average of the full and truncated terms in the retrieval equation for 18:00 UTC, 28 June 2008. MOP-Ret, MOP-Trc, and MOP-Res are the MOPITT retrieval, truncated retrieval, and residual profiles respectively. QOR-Ret is the MOPITT QOR profile, QOR-Trc is the truncated MOPITT QOR profile, and QOR-Res is the MOPITT QOR residual profile. (b) Horizontal domain average of the MOPITT averaging kernel profiles in the upper panel and leading left singular vectors of those averaging kernels in the lower panel for 18:00 UTC, 28 June 2008.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

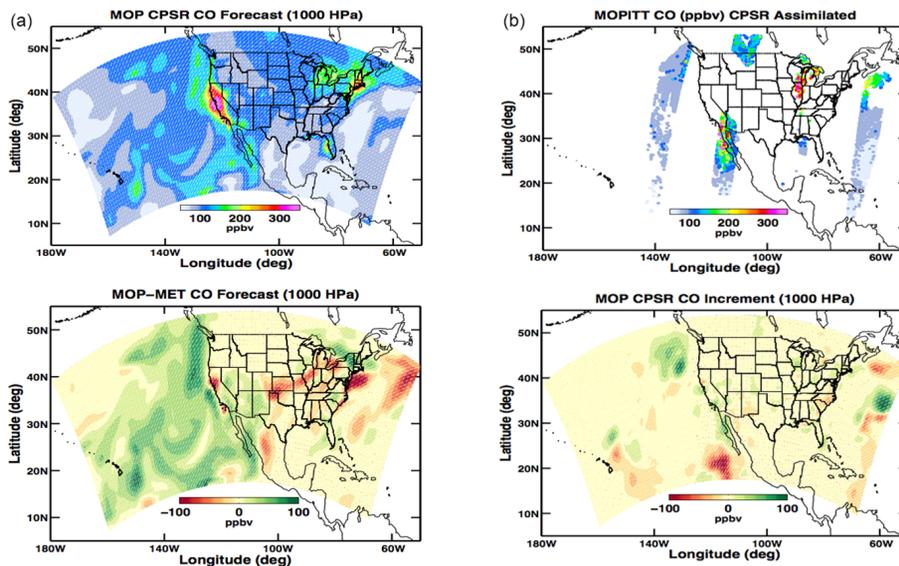


Figure 5. (a) Same as Fig. 1a except for the MOP CPSR experiment and the middle panel from Fig. 1a, the MET DA experiment is not plotted. (b) Same as Fig. 1b except for the MOP CPSR experiment.

Title Page

Abstract

Introduction

Conclusions

References

Tables

Figures

◀

▶

◀

▶

Back

Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

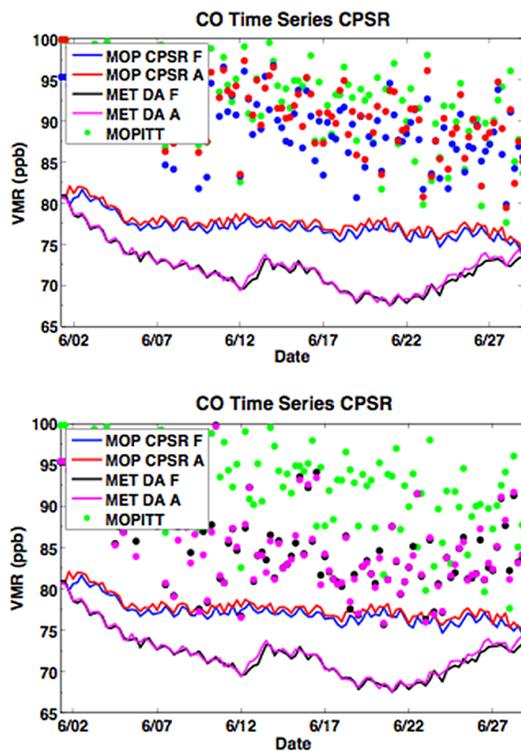


Figure 6. Same as Fig. 2 except for the MOP CPSR experiment.

[Title Page](#)[Abstract](#)[Introduction](#)[Conclusions](#)[References](#)[Tables](#)[Figures](#)[⏪](#)[⏩](#)[◀](#)[▶](#)[Back](#)[Close](#)[Full Screen / Esc](#)[Printer-friendly Version](#)[Interactive Discussion](#)

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

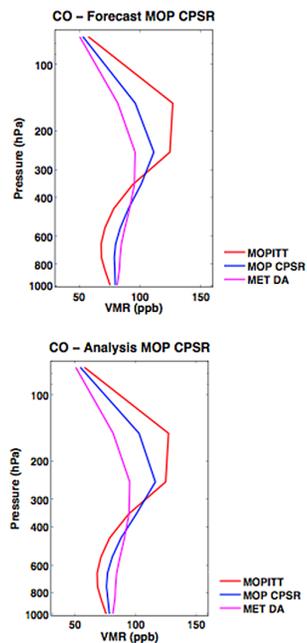


Figure 7. Same as Fig. 3 except for the MOP CPSR experiment.

Assimilating compact phase space retrievals of atmospheric composition

A. P. Mizzi et al.

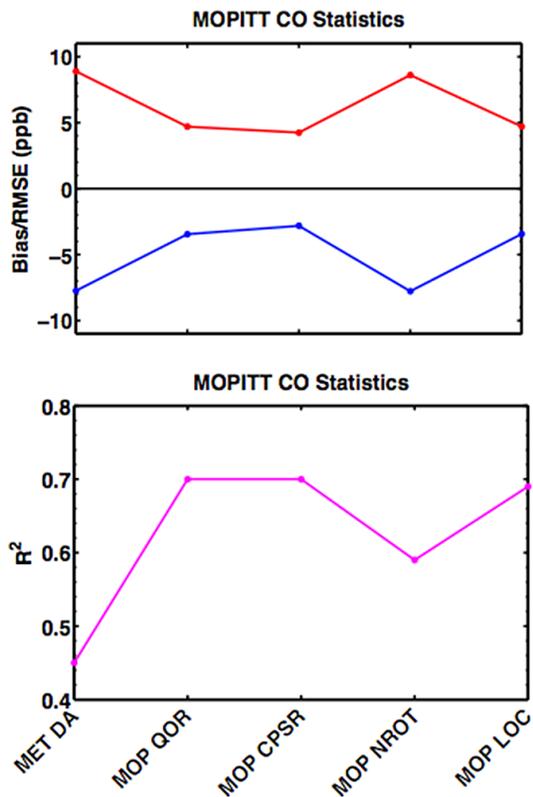


Figure 8. Verification statistics for all experiments in MOPITT retrieval space. The blue curve is the bias (model – observation). The red curve is the root mean square error (RMSE), and the magenta curve is the Coefficient of Determination. The experiments are described in the text and summarized in Table 1.