

Interactive comment on “A Multiphase CMAQ Version 5.0 Adjoint” by Shunliu Zhao et al.

Anonymous Referee #1

Received and published: 10 January 2020

COMMENT: The authors present a description and evaluation of the implementation of an adjoint methodology into CMAQ version 5.0. This method is compatible with all the major components of the CMAQ model, which is a step forward from previously published implementations in recent versions of CMAQ that only included the implementation of the adjoint approach for inert aerosol species. The authors evaluate the adjoint implementation in each of the major modules of CMAQ which allows for better confidence in the approach and also provides useful information about which modules are best suited to an adjoint. This could guide future decisions about which particular model components (such as inorganic thermodynamics) to include as part of the core model. Components better suited for sensitivity analysis might be a higher priority in situations where multiple choices exist and perform similarly in terms of speed and skill.

The manuscript is generally well organized and written. The use of brute-force sensitivity and finite difference as an evaluation approach is novel. One concern is the illustrative example at the end. It is very helpful to have an illustrative example of the type of information the adjoint provides, but the Figures (Figure 14) related to the illustrative example are confusing to interpret. The Figure caption suggests annual monetized health benefits normalized by emissions are presented. However, it is not clear whether the monetized benefits are normalized by national emissions or emissions from that same grid cell. Further, it is confusing to think about monetized health effects in places where no people reside (over the ocean for instance) and also where there are little to no emissions (northern Ontario near Hudson Bay). Perhaps there is a alternative illustration of the type of information the adjoint provides which would be simpler to interpret such as looking at concentrations relative to some source/region and not even get into converting the concentrations to health effects.

RESPONSE: We appreciate the reviewer’s comment and concern about the clarity of the illustrative example. While we agree with the reviewer that other examples may be more intuitive and easier to follow, we believe that source attribution of health impacts as location-specific BPTs is one of the most relevant, lucid, and practically significant examples of unique capabilities that the adjoint approach offers in the area of policy analysis. To address the reviewer’s concern we have completely revised the section, to better explain the process for adjoint-based source attribution of health impacts, and the meaning of the calculated BPTs.

The reviewer is correct that locations with no population or emissions can have large BPTs. We define the adjoint cost function as the benefit over the entire US domain. What we obtain from

the adjoint simulation (with unit conversion in post-processing) is the location-specific BPTs. In other words, the BPT values shown in Figure 14 are not normalized by emissions and only suggest how much benefit we would gain (or how much damage we would cause) if we cut a ton (or add a ton) of emissions of a pollutant at a specific location. As benefits are considered at the national scale (i.e., for the entire contiguous U.S.), emissions at locations with no population or emission could have an impact on health due to transport. As part of revising this section, and to address this specific point, the following is added to the manuscript”

“While the adjoint cost function is defined based on PM_{2.5} long-term mortality in the US alone, location-specific BPTs also provide a measure of cross-border impact. Finally, we note that BPTs are measures of marginal rather than total societal impact across the U.S., and as such, even areas with little or no emissions may show sizeable BPT estimates.”

Finally, the reviewer’s points about different science modules and their performance with respect to formal sensitivity analysis are well taken. While we agree with these comments, we would also like to point out that evaluation of numerical approaches and algorithms based on their performance in sensitivity analyses and their differentiability is a new and emerging concept in air quality and atmospheric modeling. Historically, these models have not been developed with differentiability in mind, but with accuracy and computational efficiency as the main drivers. As a result of ensuing practical trade-offs, discontinuities abound throughout CMAQ, as well as in other CTMs. These discontinuities are encountered in most science modules such as (in addition to inorganic thermodynamics) cloud processes, advection, mode-merging, SOA formation, native solvers of gaseous and aqueous chemistry, etc. We believe it will be a gradual but continued effort among the modeling community to address the issue of differentiability in future generations of algorithms used in CTMs.

To emphasize this important point we have added the following to the manuscript (section 3.2.1):

“The example given above is one of numerous cases where FDM was found to be inaccurate or inadequate in evaluating adjoint sensitivities. The inadequacy of FDM in producing accurate sensitivity estimates is due to process nonlinearities, as well as discontinuities that exist throughout CMAQ. This is the case in a number of CMAQ processes such as SOA formation, inorganic thermodynamics, clouds, aqueous chemistry, advection, etc. This issue is not limited to CMAQ alone and exists in all air quality models, as providing a smooth solution for the governing equations may be lost in trade-offs for added computational efficiency, improving stability, or reducing numerical artifacts in the development stage.”

And the following to our conclusion:

“We find that the development of adjoint versions of air quality and atmospheric models is often complicated by the abundance of discontinuities throughout these models that make differentiation challenging. Historically, these models have not been developed with differentiability in mind, but with accuracy and computational efficiency as the main drivers. As the development and applications of formal sensitivity analysis tools (such as adjoint models) become more prevalent, there is a need for a gradual but sustained effort by the modeling community to consider differentiability as an additional design constraint in future developments.”

Interactive comment on “A Multiphase CMAQ Version 5.0 Adjoint” by Shunliu Zhao et al.

Anonymous Referee #2

Received and published: 19 February 2020

Summary

This is a very nice model development paper that summarizes a new capability in CMAQ with a potentially wide range of future applications for source attribution, inverse modeling, etc. Having both discrete and continuous adjoints as necessary for different processes along with the use of FDM vs CVM when necessary makes this novel for an air quality model like CMAQ. The authors have systematically broken down the CMAQ model into each of its major atmospheric processes and discussed both the implementation and evaluation of the adjoint technique, along with a policy-relevant illustration at the end.

COMMENT: Given the motivation for this development to go beyond the earlier version of CMAQ Adjoint for gas-phase chemistry and demonstrate capability to model PM_{2.5}, I find it extremely limiting that the evaluation scenario used only a 7-day model simulation, and all evaluation is apparently shown only for a single hour (last hour of a day). While I appreciate the resource requirements for a longer time period, with the growth in computing technologies, it would have been valuable if the evaluation was performed for a one-month period at the minimum to ensure that the results are robust. Also, showing the evaluation for a 24-hour average (in addition to the single hour shown mostly) would also be policy-relevant given the short-term form of the health-based standard for PM_{2.5}.

RESPONSE: We thank the reviewer for pointing out a source of confusion in the manuscript. All evaluations are done for a full day. However, the adjoint model is only forced at the last time-step, and the gradient is then evaluated after 24 hours of backward simulations. While using continuous forcing, for instance one that corresponds to 24-hour average concentration would be more policy-relevant, we have instead used pulse, instantaneous forcing. The choice of pulse forcing for the adjoint evaluation has two reasons. First, this choice allows for a more straightforward testing process (for example, the calculation of finite difference or CVM sensitivities), especially when a large number of tests need to be conducted. Second, using instantaneous forcing provides a more stringent evaluation framework for the adjoint results. The magnitude of the impact of a source often diminishes gradually, or at times precipitously, with time and distance. This means that receptors that are closer to a source, both horizontally and vertically, are likely to be impacted by the source at a higher rate. In adjoint sense, this means that nearby sources are likely to show more prominently when those receptors are forced. The same is true about evolution of influences in time for processes that do not include transport and are only integrated in time; time steps that are closer to the initial forcing, would

have the largest influence. As a result, if forcing is done continuously, then the adjoint estimate is dominated by “local” (in time and space) influences. Under continuous forcing, and with this large “local” component, it would become more difficult to evaluate (for example in comparison with brute-force) how the adjoint influences are sustained over longer times or larger distances. We believe that a daylong simulation (translating to an average transport range of few hundred kilometers) with pulse forcing, as implemented, would provide a signal for the evaluation that has sufficiently evolved but is not overly diminished.

The reviewer’s point about the need for a longer simulation period in our application example, i.e. source attribution of health impacts, is well taken. We have extended our simulation time to two full warm and cold seasons (see revised Figure 14).

COMMENT: Further, the species used for evaluation is very selected (mostly ASO4J) and not robust and comprehensive. I suggest the authors quantify the evaluation metrics for all major PM2.5 constituents for each process when applicable.

RESPONSE: As stated in the manuscript (original, Page 12 Line 31; revised, Page 13 Line 8), our intention is to make, within computational constraints, the evaluation framework more comprehensive in terms of physical pathways and numerical procedures. Testing all source-receptor species pairs is not feasible, and summing up the PM species would obfuscate important processes. We have chosen the source-receptor species pair to best match the process being evaluated. For example, we use AALKJ for the secondary organic aerosol process, ASO4J for the aerosol thermodynamics module ISORROPIA/ANISORROPIA, and O3 for chemistry. For transport processes, the choice does not matter; we use ASO4J, for consistency.

The above-mentioned species are those chosen as receptor species; there are other species used as sources. For example, we use ASO4J as both receptor and source for clouds dynamics, but SO2 is adopted as the source to account for the impact through aqueous chemistry on ASO4J. For aerosols, we choose ASO4J as receptor and ASO4I as source to test intermodal (e.g., coagulation or mode-merging) processes. In cases where we have encountered outstanding discrepancies, we have used the Jacobian of sensitivities (i.e., all possible source-receptor species pairs) related to a process to investigate possible causes, as done for the chemistry process.

Specific Comments:

COMMENT: Page 2 Line 26: Please add full citation for Constantin and Barrett, 2014 to the list of references. It is missing now.

RESPONSE: Added: Constantin, B. V. and Barrett, S. R.: Application of the complex step method to chemistry-transport modeling. Atmos. Environ., 99, 457-465, 10.1016/j.atmosenv.2014.10.017, 2014.

COMMENT: Page 4 Line 5: Can the authors add a brief description of how the CMAQ adjoint capabilities as described in this study are different from the GEOS-Chem adjoint in Henze et al, 2007 or any other updates since then?

RESPONSE: The difference between the two adjoint models are inherited from the two primal models, CMAQ and GEOS-Chem. GEOS-Chem is a global model with coarser resolution, compared to the regional CMAQ model. As a global model, development efforts for GEOS-Chem have had more of a focus (compared to CMAQ) on processes that impact long-range transport (convective mixing, stratospheric intrusion, etc), long-lived species (e.g., CO, methane), and global budgets of atmospheric constituents (e.g., lightning emissions). On the other hand, CMAQ, as a limited area model has seen more development and details in processes that affect surface concentrations. Of course this distinction is rather simplistic, and the scales of the two models have approached in past years, as CMAQ now has a hemispheric version, and GEOS-Chem is capable of higher-resolution simulations in nested configurations. These differences in the base models are also reflected in the adjoint versions.

We do not include a detailed discussion of differences between GEOS-Chem and CMAQ adjoints in the manuscript, as we feel the scale and applications of the two models are quite different such that a comparison would be beyond the scope of a manuscript about CMAQ. However the manuscript is revised to read :” Adjoint of air quality models ..., or were developed for a global model with coarser resolution and varying levels of detail in representation of some of the atmospheric processes (Henze et al., 2007).”

COMMENT: Page 7 Line 6: Should deposition (dry and wet) be added to the list of science processes in CMAQ? I do see later on that the authors justify not developing an adjoint for the deposition process.

RESPONSE: We follow the convention of CMAQ about the science processes. The wet deposition is dealt with in clouds and the dry deposition is part of the vertical diffusion science process. The science processes are first mentioned in the beginning of Section 2. Both wet and dry deposition are included in the development; however, the adjoint for bidirectional deposition of ammonia is not available, and as such ammonia has conventional dry deposition in the current version.

A reference and more details are added to clarify: “As mentioned at the beginning of Section 2, the science processes in CMAQ include advection, horizontal and vertical diffusion (including dry deposition), gas-phase chemistry, aerosols (including thermodynamics and dynamics), and clouds (including aqueous chemistry and wet deposition).”

We also correct the following reference in the “Conclusions” section:

“Some components of CMAQ that do not yet have an adjoint include the bidirectional dry deposition in vertical diffusion and photolysis rate calculations in gas-phase chemistry.”

COMMENT: Page 10 Section 2.5. This section discusses development of a python-based adjoint forcing pre-processor and mentions ability to calculate local maximum 8-hr average O3 concentration. However, the final policy illustration in Section 4 uses a PM2.5 case study. Please clarify if the python-based pre-processor was enhanced for this application as well, or if a different approach was used for the demonstration case.

RESPONSE: For the PM2.5 application, we used Fortran instead of Python. However, extending the Python tool to include PM2.5 is planned.

COMMENT: Page 11 Section 3: For details about the other inputs used in this study, the reference is to Turner et al (2015a). But that study used 12km modeling, while this is using 36km. Please clarify and reconcile this apparent discrepancy.

RESPONSE: The CMAQ inputs including meteorology, emissions, initial and boundary conditions are prepared using the same configuration, except for the difference in resolution. The manuscript has been revised to clarify:

"More details about meteorological inputs, initial and boundary conditions, and emissions are provided in Turner et al. (2015a), where a version of the dataset with 12-km horizontal resolution was used."

COMMENT: Page 12 Section 3.2.1. Why was AALKJ chosen here for the evaluation of the aerosol module? Can you provide some justification?

RESPONSE: AALKJ is chosen as a relevant example, in this case for the SOA formation process. In this particular case, among various semi-volatile species, AALKJ was chosen because it had poor agreement against FDM, but performed very well in comparison to the CVM. As stated in the manuscript:

"For this process, the FDM behaved well for most of the test cases (results not shown); in the few cases when the ADJ and FDM did not agree and tuning with the perturbation sizes did not help, use of CVM demonstrated good accuracy of adjoint results (i.e., agreement along one-on-one line). This is an example of numerous cases where FDM was found to be inaccurate or inadequate in evaluating adjoint sensitivities."

COMMENT: Page 13 Line 17: Can you quantify this "acceptable accuracy", or provide a reference?

RESPONSE: Revised as (Page 13 Line 30 in the revised manuscript): "Overall, our testing confirms the findings in Capps et al. (2012) that the CVM implementation of ISORROPIA produces approximations that agree with the adjoint results."

COMMENT: Page 16 Line 1: If I understand this right, the adjoint agrees with the CVM, with a relative error less than 10%. Isn't this rather high to be acceptable? If so, for an emissions sector that has a 5% contribution, the results are within the error that the adjoint model produces, and hence cannot be meaningfully used for that range? Please provide some context to this 10% so that future users of this technique do not misuse it, and are aware of the limitations. In fact, I suggest that a separate section of limitations be added to point out other such issues.

RESPONSE: The 10% difference mentioned refers to the agreement between CADJ and the CVM, as DADJ and CVM agree very well. As mentioned in the manuscript, the difference between CADJ and CVM should not be taken as inaccuracy of CADJ, but as a manifestation of different approaches in estimating source-receptor relationships. This point is discussed in the manuscript:

"Results shown in Table 1 suggest that the discrete adjoint has a better agreement with CVM than the continuous adjoint. However, it is important to note that better agreement between the discrete adjoint and CVM should not be understood as better accuracy of the discrete adjoint in

comparison with continuous adjoint. The numerical solution to the advection equation entails inherent truncation errors from discretization schemes. These errors exist in solving the forward or adjoint advection equations; however the discrete adjoint by design remains loyal to, and consistent with the errors in the forward application (CVM in this case), while numerical solution to the continuous adjoint will result in different and inconsistent errors. The continuous adjoint is a different representation of the impacts on the adjoint cost function, but of similar mathematical accuracy, when compared to the forward or tangent linear model; therefore, the numerical solution to it should be considered as accurate as the discrete adjoint, regardless of the agreement with forward-based benchmarks such as CVM.”

The statement is also clarified to better reflect differences between CADJ, DADJ, and CVM:

“In general, the adjoint models, particularly DADJ, agree well with the CVM, while in the case of CADJ a larger relative error exists in comparison with the CVM”

A reference to the above discussion is added to the caption of Table 1:

“Table 1: Evaluation of the full adjoint model with the CADJ (Continuous Adjoint)/DADJ (Discrete Adjoint) against the CVM (Complex Variable Method)/FDM (Finite Difference Method) sensitivities of the concentrations of an accumulation-mode aerosol species ASO4J ($\mu\text{g}/\text{m}^3$) at hour 24 with respect to the concentrations of a gas species SO₂ (ppmV) at hour 23. The cells are arbitrarily picked. The perturbation size for the CVM is 1.E-12 and the one for the FDM 10%. The relation of FDM and CVM sensitivities with CADJ and DADJ results has been discussed in Section 3.3.”

COMMENT: Page 16 Line 3: “The problem with the FDM has been discussed earlier” Please provide exact Section where it was discussed earlier

RESPONSE: Revised: “The problem with the FDM has been discussed at the beginning of Section 3 and is not repeated here. ”

COMMENT: Page 16 Section 3.4: Can you provide this information in a table? It may be easier on the reader

RESPONSE: Table 2 added about checkpoint file size for each science process. Revised on p. 16: “A summary of the checkpointing file sizes is provided in Table 2.”

Table 2: Sizes of checkpoint files for the science processes in CMAQ for a single day. The computational domain has 148 columns, 112 rows, and 24 vertical layers. The synchronization time step of CMAQ is 12 minutes. Shown for horizontal and vertical advections are checkpointing file sizes from the continuous version of adjoint.

Process	Aerosols	Clouds	Chemistry	Vertical Diffusion	Horizontal Advection	Vertical Advection

File Size (GB)	24	24	17	1	0.2	0.2
----------------	----	----	----	---	-----	-----

COMMENT: Page 17: For the model application case, even though it is illustrative, please provide additional information on the source of meteorology and emissions inputs that were used.

RESPONSE: Details added for the new results:” For the backward sensitivity analysis, we run the adjoint for the year of 2016 for the contiguous U.S. domain with 36-km resolution inputs from the Intermountain West Data Warehouse (National Emissions Inventory Collaborative, 2019). The computational domain contains 172 columns and 148 rows with 35 vertical layers.”

COMMENT: Page 17 Line 3: Change “IO” to “I/O”

RESPONSE: Revised as suggested.

COMMENT: Page 17: Line 19: “extrapolating from April to the full year”. Should this really be “7 days in April” as stated in Section 3?

RESPONSE: We have now extended the results to two full 3-month seasons based on a different dataset.

COMMENT: Page 17 Line 24: Why NO and not NOx (NO + NO2)?

RESPONSE: NOx is shown for the new results.

COMMENT: Page 18 Line 18: The language re code availability seems to indicate that the model is not ready for further dissemination given the need for expanded user testing. Please clarify.

RESPONSE: We intend to publicly release the model later in 2020 and after it has gone through limited release and a wider variety of simulations (different domains, scales, etc).

COMMENT: Figures 1- 8: When you say “final concentration”, is the spatial plot for the last hour of one-day? Did that one day have a spinup or is truly the first day of the modeling for this development? Is that same one hour used in the scatter plot, or all hours from the one day? Suggest that all hours for a single day be used in the scatter.

RESPONSE: Plotted in all the plots are the sensitivities of instantaneous concentration at Hour 24 (final) to instantaneous concentrations at Hour 0 (initial). To produce sensitivities to concentrations at each hour, it would require 23 extra runs for the FDM or CVM which is not feasible. As seen in the plots, simulations across the computational space provide sufficient points of comparison.

There is a 6-day spin-up period; the validation is performed on day 7. Revised on p.11:

“The air quality simulation scenario used for the evaluation is for the contiguous U.S. domain with a 36-km horizontal resolution and 24 vertical layers for the first seven days of April 2008, with the first 6 days used for spin-up.”

COMMENT: Figure 9: Which grid-cell was used in showing this Jacobian? Can you provide some context for the choice of this grid-cell and if it is representative for the whole domain? It will be helpful to see how these differences propagate through the 7-day period (or at least a month if feasible) that was modeled. Perhaps, show that as a time-series?

RESPONSE: While the Jacobian includes interesting and valuable information, its creation is time-consuming and resource intensive. In the current scope of work, we have used the Jacobian as a diagnostic tool and only when we encountered problems that we were not able to resolve through conventional means and efforts. The cell for which the results are shown was chosen as it presented a problematic case, where our adjoint results did not agree with CVM, and therefore, we constructed the full Jacobian.

COMMENT: Figure 7 Caption: Change “The perturbation size for the FDM is 0.01/0.001/0.0001 ppb for the three figures” to clarify which figure has what size – top to bottom or bottom to top Figure 7: Can you explain the non-zero values for FDM ranging between 0 – 3000 when ADJ = 0, for all top 3 figures?

RESPONSE: Revised as: “The perturbation size for the FDM is 0.01/0.001/0.0001 ppb for the top three figures (from top to bottom).”

Thresholds are commonly used in CMAQ to make concentrations non-zero for computational purposes, say, to avoid division by zero. When such circumstances occur, the adjoint would give zero sensitivities but the finite difference with an absolute perturbation size would not. This is one of the instances of CMAQ (and other CTMs) having a non-smooth and discontinuous response surface that would not lend itself to accurate finite difference approximation. Considering model complexity, we use sensitivities normalized by the initial concentrations to check if there is an agreement between ADJ and FDM. And there is.

COMMENT: Figure 11: What is the impact of this numerical noise in layers aloft in the adjoint of the transport scheme, and how would this affect the model results? I see a note on page 15 Lines 1-2 that “these are not uncommon and desirable fix does not appear possible.”

RESPONSE: In Figure 11, the y-axis in the tile plots represents the vertical layers. As can be seen, the numerical noise reaches out to all layers.

COMMENT: Figure 11 Caption: There is really no “top plots” here. Please reword.

RESPONSE: Revised as: “For the tile plots, the x- and y-axes represent the horizontal y direction and the vertical layers, respectively.”

COMMENT: Figure 14 Caption: Should “long-term PM2.5 exposure” really be “7-day PM2.5 exposure”?

RESPONSE: We have now extended the results to two full 3-month seasons based on a different dataset.

COMMENT: Table 1: Please define acronyms such as CADJ, DADJ, etc. What are cells 1 – 6 in this table? There is a lot of information in this table for evaluating the full adjoint, but the discussion of this table is very skinny. For e.g., look at cells #3 and #5. The FDM results range so widely (from 21.21 to -23.2), while both CADJ and DADJ results are much closer. What does this mean?

RESPONSE: The acronyms are now added. The cells are picked in an arbitrary way, loosely based on the sensitivity values of the continuous adjoint (CADJ). The values from the FDM merely suggest that a 10% perturbation size is not able to accurately produce the sensitivity information, as encountered many times in the manuscript. This also points to inadequacy of FDM in estimating single-source impacts within a model such as CMAQ that has a fragmented response surface. The Table caption is revised:

“Table 1: Evaluation of the full adjoint model with the CADJ (Continuous Adjoint)/DADJ (Discrete Adjoint) against the CVM (Complex Variable Method)/FDM (Finite Difference Method) sensitivities of the concentrations of an accumulation-mode aerosol species ASO4J ($\mu\text{g}/\text{m}^3$) at Hour 24 with respect to the concentrations of a gas species SO₂ (ppmV) at Hour 23. The cells are arbitrarily picked. The perturbation size for the CVM is 1.E-12 and the one for the FDM 10%.”

A Multiphase CMAQ Version 5.0 Adjoint

Shunliu Zhao¹, Matthew G. Russell¹, Amir Hakami¹, Shannon L. Capps², Matthew D. Turner³, Daven K. Henze⁴, Peter B. Percell⁵, Jaroslav Resler⁶, Huizhong Shen⁷, Armistead G. Russell⁷, Athanasios Nenes^{8,9,10}, Amanda J. Pappin¹¹, Sergey L. Napelenok¹², Jesse O. Bash¹², Kathleen M. Fahey¹², Gregory R. Carmichael¹³, Charles O. Stanier¹³, Tianfeng Chai¹⁴

¹Department of Civil and Environmental Engineering, Carleton University, Ottawa, ON K1S 5B6, Canada

²Civil, Architectural, and Environmental Engineering, Drexel University, Philadelphia, PA 19104, USA

³SAIC, Stennis Space Center, MS 39529, USA

⁴Mechanical Engineering Department, University of Colorado, Boulder, CO 80309, USA

⁵Department of Earth & Atmospheric Sciences, University of Houston, Houston, TX 77204, USA

⁶Institute of Computer Science of the Czech Academy of Sciences, Prague, 182 07, Czech Republic

⁷School of Civil and Environmental Engineering, Georgia Institute of Technology, Atlanta, GA 30331, USA

⁸School of Earth and Atmospheric Sciences, Georgia Institute of Technology, Atlanta, GA 30331, USA

⁹School of Architecture, Civil & Environmental Engineering, Ecole polytechnique fédérale de Lausanne, CH-1015, Lausanne,

Switzerland

¹⁰Institute for Chemical Engineering Sciences, Foundation for Research and Technology Hellas, Patras, GR-26504, Greece

¹¹Air Health Effects Division, Health Canada, Ottawa, ON K1A 0K9, Canada

¹²Atmospheric & Environmental Systems Modeling Division, U.S. EPA, Research Triangle Park, NC 27711, USA

¹³Department of Chemical and Biochemical Engineering, University of Iowa, Iowa City, IA 52242, USA

¹⁴College of Computer, Mathematical, and Natural Sciences, University of Maryland, College Park, MD 20742, USA

Correspondence to: Amir Hakami (amir.hakami@carleton.ca)

Abstract. We present the development of a multiphase adjoint for the Community Multiscale Air Quality (CMAQ) model, a widely used chemical transport model. The adjoint provides location- and time-specific gradients that can be used in various applications such as backward sensitivity analysis, source attribution, optimal pollution control, data assimilation and inverse modeling. The science processes of the CMAQ model include gas-phase chemistry, aerosols, cloud chemistry and dynamics, diffusion and advection. Discrete adjoints are implemented for all the science processes, with an additional continuous adjoint for advection. The development of discrete adjoints is assisted with Algorithmic Differentiation (AD) tools. Particularly, the Kinetic PreProcessor (KPP) is implemented for gas-phase and aqueous chemistry, and two different automatic differentiation tools are used for other processes such as clouds, aerosols, diffusion, and advection. The continuous adjoint of advection is developed manually. For adjoint validation, the brute-force or Finite Difference Method (FDM) is implemented process by process with box- or column-model simulations. Due to the inherent limitations of the FDM caused by numerical round-off errors, the Complex Variable Method (CVM) is adopted where necessary. The adjoint model often shows better agreement with the CVM than with the FDM. The adjoints of all science processes compare favorably with the FDM/CVM. In an example application of the full, multiphase adjoint model, we provide the first estimates of how emissions of PM_{2.5} affect public health across the US.

1 Introduction

Adjoint models generate gradients which can be used directly for backward sensitivity analysis or to provide directions for gradient-based optimization in four-dimensional variational data assimilation (4D-Var) or other inverse problems (Errico, 1997; Navon, 1997; Giles and Pierce, 2000; Wang et al., 2001; Sandu et al., 2005; Griewank, 2012). Applications of adjoint models for data assimilation have a long successful history in meteorology and oceanography (Errico, 1997; Navon, 1997). For atmospheric chemistry, adjoint modeling was used as early as the 1990s (Fisher and Lary, 1995; Elbern et al., 1997). More recently the methods were applied to aerosols in 1D models (Henze et al., 2004; Sandu et al., 2005), 3D models of inert aerosol mass concentrations (Hakami et al., 2005), and 3D models of chemically active aerosol mass concentrations (Henze et al., 2007).

Due to omnipresent uncertainties in emissions, initial and boundary conditions, as well as the underlying complex physical and chemical processes, predicting or accurately simulating air quality poses great challenges; assimilation of chemical data is thus a promising approach in improving the model skill (Carmichael et al., 2008). While applications in data assimilation and inverse modeling form the traditional niche for adjoint applications in atmospheric modeling, adjoint models can also be used to conduct sensitivity analysis. Sensitivity analyses are often performed in air quality studies to estimate the impact of various model inputs, in particular emissions, on model predictions (Dunker et al., 2002; Hakami et al., 2003; Sandu et al., 2005; Cohan et al., 2005; Napelenok et al., 2006; Martien and Harley, 2006; Koo et al., 2007). Among various methods for sensitivity analysis, two general categories are more commonly used: forward and adjoint (Hakami et al., 2007). In the forward approach, sensitivity information is propagated forward in time. The most common forward sensitivity approach is the brute-force or Finite Difference Method (FDM). The FDM requires minimal effort to implement, but the search for a proper step/perturbation size might be needed to produce accurate sensitivities (Iott et al., 1985). The step size selection process could be resource demanding and repeatedly required, especially when the numerical model contains highly nonlinear processes that are routinely encountered in atmospheric models. Another forward approach is the Complex Variable Method (CVM) (Squire and Trapp, 1998; Anderson and Nielsen, 2001). Unlike the FDM, the CVM is not subject to subtraction errors and can provide accurate sensitivities by using a perturbation size as small as allowed in floating-point calculations, but this approach has only been implemented in one atmospheric chemical transport model (Giles et al., 2003; Constantin and Barrett, 2014). A third forward approach is the Decoupled Direct Method (DDM) or the Tangent Linear Model (TLM), in which differentiation is directly applied to the governing equations or algorithmically to the primal computer model. DDM can generate exact sensitivities (i.e., subject to numerical errors, with no perturbations required) at the cost of a significant amount of model development (Dunker et al., 2002; Napelenok et al., 2006).

In the adjoint approach, the sensitivity information is propagated backward in time and a single model run generates sensitivities of a metric of model outputs with respect to all model inputs (Errico, 1997; Giles and Pierce, 2000). The adjoint

and the forward sensitivity approaches complement each other, in the sense that the forward approach calculates sensitivities of all model outputs to a single model input, combined or individual, in one model run (an extra baseline run required for the FDM), while the adjoint method provides sensitivities of a single model output (individual, or integrated) to all model inputs (Hakami et al., 2007).

5
In addition to 4D-Var and sensitivity analysis, adjoint models have been implemented for source attribution of air pollutants (Zhang et al., 2015; Qi et al., 2017). Compared to the zero-out method which computes contributions by switching an emission sector on and off, the adjoint approach has the advantage of not changing the chemical environment which could lead to inaccuracies in estimates (Koo et al., 2009). Furthermore, the adjoint approach can be readily extended to include more
10 emission sectors for investigation with a marginal increase in computational cost; as mentioned earlier, the computational cost of adjoint models is practically independent of the number of input parameters. Source attribution by adjoint models has its own limitation due to the inherent linear assumption in adjoint formulation. Koo et al. (2009) found that the linear assumption held for a 20% emission reduction in general for secondary inorganic aerosols; for secondary organic aerosols, the linearity assumption was valid for up to 100% reductions in anthropogenic emissions. Although these bounds are based on the DDM
15 of the three-dimensional Comprehensive Air quality Model with extensions (CAMx; ENVIRON, 2005; Koo et al., 2007), they are applicable to the bounds for adjoint models for source attribution, as both the DDM and the adjoint are tangent estimations and based on the same assumption of linearity.

The U.S. Environmental Protection Agency's Community Multiscale Air Quality (CMAQ) model is a regional-to-hemispheric
20 air quality model, which is widely used due to its community-driven development and state-of-the-art science components (Byun and Schere, 2006; Foley et al., 2010). Limited adjoint versions of CMAQ have been developed before; gas-phase adjoint model was previously developed for CMAQ 4.5.1 (Hakami et al., 2007) and has been used in several applications related to ozone (Resler et al., 2010; Mesbah et al., 2012; Pappin and Hakami, 2013; Zhao et al., 2013; Pappin et al., 2015; Pappin et al., 2016; Park et al., 2016). Turner et al. (2015a; 2015b) developed and applied the adjoint of the black carbon (BC) aerosol for
25 CMAQ 4.7.1 but did not include other aerosol species or gas-phase chemistry.

The lack of chemically comprehensive aerosol and cloud processes in the adjoint model has so far prevented applications related to aerosols, which in turn has imposed significant limitation on adjoint-based multi-pollutant studies on topics such as human health and climate. Analogous to ozone, exposure to fine particulate matter ($PM_{2.5}$) poses risks to human health (Brook
30 et al., 2010; Crouse et al., 2012, 2015; West et al., 2016; Turner et al., 2016; Di et al., 2017; Pinault et al., 2017). Particulate matter also plays a significant role in climate change by influencing the radiative budget of the atmosphere (Tai et al., 2010, 2012; Fiore et al., 2012; Fuzzi et al., 2015). A multiphase adjoint model has been shown to better delineate the influence of model inputs such as emissions on human health (Lee et al., 2015; Koplitz et al., 2016) and climate (Henze et al., 2012; Karydis et al., 2012; Lacey et al., 2017).

Deleted: : Henze et al., 2007

Adjoint of air quality models have been developed before, but many of these adjoint models were for the gas-phase chemistry and/or contain a less detailed treatment (lacking microphysics or thermodynamics) of aerosols than that of CMAQ (Elbern et al., 2000; Henze et al., 2004; Sandu et al., 2005; Hakami et al., 2005; Martini and Harley, 2006; Dubovik et al., 2008; Huneus et al., 2009), or were developed for a global model with coarser resolution and varying levels of detail in representation of some of the atmospheric processes (Henze et al., 2007). This work aims to fill in these gaps by developing a full adjoint for a widely used regional air quality model with detailed, multiphase, size-resolved treatment of aerosols, and to modify the adjoint to reflect more recent science process updates present in CMAQ version 5.0.

2 Adjoint Model Development

10 The CMAQ modeling system solves the atmospheric diffusion equations (ADEs, Byun, 1999; Jacobson, 2005; Seinfeld and Pandis, 2006)

$$\frac{\partial c_i}{\partial t} = -\mathbf{u} \cdot \nabla c_i + \frac{1}{\rho} \nabla \cdot (\rho \mathbf{K} \nabla c_i) + r_i + e_i + s_i, \quad (1)$$

where c_i is the mixing ratio of species i , \mathbf{u} is the wind velocity, ρ is the air density, \mathbf{K} is the diffusivity tensor, and r_i , e_i and s_i represent the change rates from chemical/thermodynamic transformations, emissions, and the loss process for species i , respectively. The first two terms on the right-hand side represent the transport process, namely advection and diffusion. Given the proper initial conditions and boundary conditions, the CMAQ model simulates the fate of air pollutants and their precursors, emitted to or produced via chemical reactions in the atmosphere.

Integration of the ADE in CMAQ is accomplished through operator splitting, which facilitates the modular structure of the model (Byun and Schere, 2006; McRae et al., 1982). CMAQ includes sub-modules implemented for all the science processes:

- VDIFF for vertical diffusion,
- HADV for horizontal advection,
- ZADV for vertical advection,
- HDIFF for horizontal diffusion,
- CLDPROC for cloud dynamics and aqueous chemistry,
- CHEM for gas-phase chemistry,
- AERO for aerosol dynamics and thermodynamics.

In the CMAQ model, the science processes are executed one after another at every synchronization time step that is dictated by the stability criteria for horizontal advection (Byun and Schere, 2006). To guarantee accuracy and to meet stability criteria, internal time steps specific to each process are also employed.

The adjoint equations corresponding to the ADEs can be written as

$$-\frac{\partial \lambda_i}{\partial t} = \nabla \cdot (\mathbf{u} \lambda_i) + \nabla \cdot \left(\rho \mathbf{K} \nabla \frac{\lambda_i}{\rho} \right) + r_i + \varphi_i, \quad (2)$$

where λ_i represents the adjoint variable of species i ; r_i represents the contributions from r_i , e_i and s_i ; and φ_i denotes adjoint forcing (Elbern et al., 2000; Sandu et al., 2005; Martien and Harley, 2006; Henze et al., 2007; Hakami et al., 2005, 2007). In the following sub-sections, the adjoint model development techniques and strategies are introduced and the details of the challenges and treatment for each science module are discussed.

2.1 Continuous and discrete adjoints

There are two approaches in developing an adjoint model, the discrete and the continuous (Giles and Pierce, 2000). The discrete approach starts with a numerical model of the primal equation and differentiates it directly line by line, or differentiates the numerical algorithm used to solve the continuous primal equation. One significant advantage of the discrete approach is that the model building process can be automated, at least partially (Giering and Kaminski, 1998; Griewank, 2003). A variety of Automatic Differentiation (AD; also referred to as Algorithmic Differentiation) tools for various programming languages are available (e.g., <http://www.autodiff.org/>). It should be noted that the sensitivities from the discrete adjoint model are exact in the sense that they are the exact (to machine precision) first-order derivatives of the forward model unless approximations are made (Errico, 1997). It is expected that the sensitivities are comparable with those from the FDM (with properly chosen perturbation sizes) or the CVM as the FDM and the CVM are both based on the same forward model (see Section 3 for the details of the CVM).

The continuous approach takes a governing equation, derives its adjoint equation, and numerically solves the adjoint equation. The continuous adjoint model is not constrained to using the same numerical scheme as the forward or the discrete adjoint model (Sirkes and Tziperman, 1997). Take horizontal advection as an example. The forward equation and the adjoint equation share the same form; the only difference is that the adjoint equation runs backward in time (Hakami et al., 2007; Gou and Sandu, 2011). Implemented for advection in CMAQ is the Piecewise Parabolic Method (PPM), which is a higher-order Godunov-type method and uses intrinsic dissipation to improve stability and accuracy (Colella and Woodward, 1984; Byun and Schere, 2006). With the PPM employed for the adjoint equation, the corresponding continuous adjoint model bears the same desirable numerical features.

2.2 The backward nature of an adjoint

Adjoint models are integrated backward in time. The nature of the backward propagation of adjoint sensitivities could be best demonstrated from a discrete perspective, which is detailed in Wang et al. (2001). Suppose we have the forward primal model written as

$$\mathbf{c}_t = \mathbf{G}_{t-1}(\mathbf{c}_{t-1}), \quad (3)$$

where \mathbf{c} is the vector of state variables (e.g., concentrations), subscript 't' indicates time, and \mathbf{G} denotes the primal model.

Linearizing the above equation one can obtain the corresponding TLM (Talagrand and Courtier, 1987) as

$$\delta \mathbf{c}_t = \mathbf{G}'_{t-1} \delta \mathbf{c}_{t-1}, \quad (4)$$

- 5 where $\delta \mathbf{c}$ represents perturbations to the state variables, and \mathbf{G}'_{t-1} is the Jacobian matrix. The Jacobian matrix is the first derivative of \mathbf{G}_{t-1} with respect to the input vector and has the following form (Wang et al., 2001),

$$\mathbf{G}'_{t-1} = \begin{pmatrix} \frac{\partial G_{11}}{\partial c_1} & \frac{\partial G_{12}}{\partial c_2} & \dots & \frac{\partial G_{1n}}{\partial c_n} \\ \frac{\partial G_{21}}{\partial c_1} & \frac{\partial G_{22}}{\partial c_2} & \dots & \frac{\partial G_{2n}}{\partial c_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial G_{n1}}{\partial c_1} & \frac{\partial G_{n1}}{\partial c_2} & \dots & \frac{\partial G_{nn}}{\partial c_n} \end{pmatrix}_{t-1}. \quad (5)$$

- The Jacobian matrix is not readily explicit in the model; by perturbing any one of the elements of the input vector \mathbf{c}_0 , we obtain the values of the corresponding column of the Jacobian matrix. (Depending on the problem at hand, a combination of
10 perturbations could be useful and feasible.) The obtained sensitivities are with respect to the perturbed input of all output variables or a metric defined over the output variables.

To resolve the values of a row of the Jacobian matrix, the Jacobian matrix can be transposed to construct the following adjoint model,

$$15 \quad \boldsymbol{\lambda}_{t-1} = \mathbf{G}'_{t-1}^T \boldsymbol{\lambda}_t, \quad (6)$$

where $\boldsymbol{\lambda}$ is the vector of adjoint variables. The row of the Jacobian matrix holds the sensitivities with respect to all input variables.

For simulations from time 0 to t, the adjoint could be written

$$20 \quad \boldsymbol{\lambda}_0 = \mathbf{G}'^T \boldsymbol{\lambda}_t = (\mathbf{G}'_{t-1} \circ \mathbf{G}'_{t-2} \circ \dots \circ \mathbf{G}'_0)^T \boldsymbol{\lambda}_t = \mathbf{G}'_0^T \circ \mathbf{G}'_1^T \circ \dots \circ \mathbf{G}'_{t-1}^T \boldsymbol{\lambda}_t, \quad (7)$$

- where subscript 't' indicates the last time step used in calculating the Jacobian matrix. Because of the transposition, the order of the composition of the Jacobian matrices for different time steps is reversed. The Jacobian matrix at the last time step is applied first, instead of the one at the first time step. In other words, the adjoint sensitivities are propagated backward in time. A description from the continuous adjoint perspective of the backward propagation nature is found in Giles and Pierce (2000),
25 Hakami et al., (2007) or Henze et al. (2007).

Propagating backward in time introduces another prominent challenge for the adjoint model, i.e., the checkpointing of intermediate values of state variables for nonlinear processes. Unless the forward primal model is linear, the intermediate

values of the state variables are needed to calculate the transposed Jacobian matrix at each time step. Strategies of checkpointing are discussed in Section 2.4.

2.3 Automatic differentiation for the CMAQ science processes

5 The adjoint model development has been assisted with several AD tools. As mentioned earlier, the science processes in CMAQ include transport (advection and diffusion), gas-phase chemistry, aerosols and clouds. For the transport processes, TAMC (Tangent linear and Adjoint Model Compiler) was employed for the adjoint (Giering, 1999). The Kinetic PreProcessor (KPP) was adopted for the gas-phase chemistry and the aqueous chemistry of clouds (Damian et al., 2002). For aerosols and cloud dynamics, Tapenade was used to generate the adjoint (Hascoët and Pascual, 2013).

10 The AD approach is used for all the CMAQ science processes. The CMAQ code is written in Fortran 77/90 and is not ready for AD in general; significant modifications are required to process the original CMAQ code. For Tapenade 3.10, those modifications include preprocessing directives and macros, defining a proper cost function, constructing a root subroutine, and dealing with pointers and black-box subroutines, to name a few. Without the revisions, the AD tool either fails to generate the
15 adjoint files or produces one that requires excessive manual intervention.

In addition to the changes mentioned above, an important numerical procedure, the bisection method, needs some special treatment for the adjoint. The bisection procedure does not provide a passage for the propagation of sensitivity information. In the current work, the post-differentiation technique (Bartholomew-Biggs, 1998) is implemented before the adjoint is derived
20 so as to obtain sensitivity information in a manner consistent with the underlying algorithm. The extensively used bisection procedure remains in the CMAQ code; after the solution is converged, an extra step of the Newton-Raphson method is attached to the bisection procedure to facilitate calculation of the gradient of the equation at the root (Capps et al., 2012).

The adjoint models generated by the AD tools do not run right out of the box. A post-processing step is necessary to prepare
25 the code for testing, including modifications based on the error/warning messages issued by the tools and checking against the original CMAQ code, for example, for missing code parts. In general, the full cycle of development involves code preparation, AD, post-processing, and validation, and is repeated as many times as needed.

2.3.1 Aerosols

The CMAQ model uses the modal approach to treat aerosols (Binkowski and Roselle, 2003). Specifically, the size range of an
30 aerosol species is divided into three modes: the Aitken mode of aerosols with a geometric diameter less than 0.1 μm , the accumulation mode with a geometric diameter between 0.1 μm and 2.5 μm , and the coarse mode with a geometric diameter greater than 2.5 μm . Fine particles include those with a geometric diameter less than 2.5 μm , which include all particles in the

Aitken and accumulation modes. Although the geometric diameter is important for distinguishing the modes, it is a derived quantity in CMAQ. The aerosols are represented by mass, number and surface area by design; all the other quantities required for simulation are derived from these three representative moments based on the lognormal assumption of the size distributions (Binkowski and Roselle, 2003).

5 The key components of the CMAQ aerosol module include the production of secondary organic aerosols (SOA), new particle formation from nucleation, particle coagulation and condensational growth, heterogeneous chemical reactions to generate nitric acid, mode merging and aerosol thermodynamics (see Subsection 2.3.2 for cloud-related aerosol processes). Aerosol thermodynamics is treated with the ISORROPIA thermodynamic equilibrium model (Nenes et al., 1998), for which the adjoint, 10 ANISORROPIA, has been developed and documented in Capps et al. (2012). Bisection procedures are extensively used in ISORROPIA and the post-differentiation technique mentioned in Section 2.3 was employed to ensure the propagation of sensitivity information. Capps et al. (2012) discuss how [challenges posed by the non-linearity and solution discontinuity in ISORROPIA are handled in the adjoint development](#).

Deleted: the non-linearity and solution discontinuity in ISORROPIA poses great challenges in validating the adjoint model.

15 To generate the adjoint by AD, the rest of the processes in the aerosol module are lumped into a box model. The corresponding Fortran code is preprocessed for the AD tool. Although a Newton-type iterative procedure, instead of a bisection method, is used in the SOA process, the post-differentiation technique is implemented to improve computational efficiency. In other words, the original Newton routines are used to obtain a converged solution, which is then used to propagate the sensitivity information through the corresponding adjoint routine generated from the 1-step Newton method. It is worth noting that post- 20 differentiation is applied when post-processing the adjoint code; for AD, the original Newton routine is revised to iterate only once.

2.3.2 Clouds

The cloud module in CMAQ deals with cloud dynamics and aqueous chemistry. Depending on the cloud size, one of two solution techniques is employed. The resolved cloud sub-module (RESCLD) is invoked when the cloud size is larger than the 25 grid size. Under this circumstance, cloud dynamics become part of the transport process and need not to be treated separately in the cloud module. When clouds partially exist in a cell, the subgrid module for convective clouds (CONVCLD) is invoked, which resolves the vertical convective mixing in the boundary layer based on the Asymmetrical Convective Model (Pleim and Chang, 1992). The mixing process computes mixing ratios for each individual species inside and outside a cloud. The obtained mixing ratios are then redistributed to each layer according to its initial value using a weighting function. The nonlinearity 30 introduced by the weighting function has proven to be problematic in differentiation as discussed in Section 3 during testing. Both sub-modules, RESCLD and CONVCLD, simulate in-cloud scavenging, wet deposition, and aqueous chemical reactions. An exponential decay formulation is used for the in-cloud scavenging processes. For cloud dynamics, the same AD tool and procedure as for the aerosol module are employed for the adjoint development.

For aqueous chemistry, KPP version 2.2.3 is used (Damian et al., 2002). Unlike the other AD tools used in this study, KPP operates on the algorithmic level instead of directly on the code; given a chemical mechanism, the KPP kinetic preprocessor can generate the corresponding computer code in several languages including Fortran 77/90. The KPP has the capability to generate the forward model, the DDM/TLM, and the adjoint through separate runs. A detailed treatment of cloud chemistry with KPP can be found elsewhere (Fahey et al., 2017). The species treated in aqueous chemistry are consistent with the CB05 chemical mechanism and the AERO5 aerosol module in the current adjoint implementation.

2.3.3 Gas-Phase Chemistry

As done for the previous version of CMAQ-ADJ (Hakami et al., 2007), the KPP kinetic preprocessor is used to generate the subroutines required for constructing the adjoint. The chemical kinetic mechanism implemented is updated to CB05 from CB-IV for the previous adjoint (Yarwood et al., 2005).

2.3.4 Transport

The transport module in CMAQ 5.0 consists of four components including horizontal advection, vertical advection, horizontal diffusion, and vertical diffusion. Horizontal advection is further divided into two x - and y -direction components with the order of the two advectons alternating to maintain a symmetric form (Byun and Schere, 2006). As discussed in Subsection 2.1, the PPM scheme is implemented for the advection process. PPM is monotonic and enforces positivity (Byun and Schere, 2006). As the sensitivity values could be either positive or negative, the positivity enforcing feature should be disabled when the method is applied for the adjoint equation to develop a continuous version of the adjoint (Gou and Sandu, 2011). The discrete adjoints of all the four components of transport were developed using TAMC version 5.3.2, which is an AD tool for Fortran-77 programs with partial Fortran 90 support (Giering, 1999).

2.4 Manual interventions

Manual interventions are necessary to revise and assemble the adjoint source code generated by AD tools. First, the warning and error messages issued by AD tools are checked. Once the code is successfully compiled, the forward sweep of the adjoint routines is checked against the original CMAQ for completeness. For post-differentiation, the iterative Newton-Raphson method originally implemented for SOA is added back to the forward sweep, to replace the one-step version created specifically for AD. The solution is saved (pushed into stack) after convergence and restored (popped out of stack) for the one-step adjoint routine of the Newton-Raphson method. Last, but not least, the adjoint code is refined according to the coding guidelines of CMAQ including indentation styles and the usage of macros and directives.

As mentioned in Section 2.2, the adjoint model is integrated backward in time and checkpointing is required for active variables used in the propagation of sensitivity information through nonlinear processes. There are two strategies to retrieve intermediate

values: repeatedly run the forward model or run the forward model once and store all the intermediate values. In practice, a trade-off between the two strategies is employed, as repeatedly running the forward model is prohibitively expensive and saving all intermediate values is impractical due to limited storage space (Wang et al., 2009). Intermediate values can be saved to memory, allowing for faster access, but subject to physical memory constraints.

5
The AD tools automatically employ a combined strategy of rerunning and checkpointing for intermediate values of active variables. For example, Tapenade performs a live analysis to determine if a variable is active and automatically apply a strategy that combines checkpointing and rerunning for those active variables at the subroutine level. Every adjoint routine contains a forward sweep and a backward sweep. During the forward sweep (rerunning), the values of active variables (including control
10 variables) are pushed into a stack (checkpointing), and then are popped out and used during the backward sweep. There are different stacks for different data types such as integer (for conditional/branch control), real, and double precision variables. The values of active variables are checkpointed to a stack of its own type accordingly. Stacks operate on the last-in first-out principle and are well-suited for the checkpointing purpose. Attention should be paid to the live analysis process; the “save” attribute or a bug in the AD tool could unexpectedly cause corrupted checkpointing of an active variable.

15
This implementation of internal checkpointing entails saving of numerous intermediate variables and requires far more computer memory than available for a typical regional air quality application with all the science processes involved for a multiple-day simulation. The strategy adopted for the current full adjoint model is the same as the one used for the previous version (Hakami et al., 2007). A forward run is carried out with the original CMAQ model revised to checkpoint to files the
20 values of active variables for each nonlinear science process at every synchronization time step (i.e., the external time step for horizontal advection). The backward adjoint run then reads the checkpointing file at the beginning of the corresponding science process.

2.5 Adjoint forcing pre-processor

The adjoint model calculates the sensitivity of a cost function, J , with respect to model parameters. The adjoint forcing, φ_i , in
25 Equation 2 corresponds to $\frac{\partial J}{\partial c_i}$. Like emissions in the forward sweep, the adjoint forcing is applied at every time step in the backward sweep and may be applied in any grid cell. Typically, the cost function depends on the concentration-based results of the forward sweep. The derivative of the cost function with respect to the concentration of interest must be calculated for each time step and added to an IOAPI-compliant netCDF file. To ease the burden of introducing different cost functions for users, a Python-based adjoint forcing pre-processor has been developed. As a model, the pre-processor includes the calculation
30 of the local maximum daily 8-hour average ozone concentration, which is a regulatory metric in the U.S. This example addresses issues of shifting to local time and a forcing dependent on an average in every grid cell. The example produces an adjoint forcing file corresponding to this cost function ready for use in the backward sweep. The Python-based pre-processor

will make implementation of additional cost functions such as 24-hour average aerosol constituents or observation operators for satellite-based atmospheric composition straightforward.

3 Model Evaluation

The adjoint model is evaluated on a process-by-process basis against the brute force Finite Difference Method (FDM) and the
5 Complex Variable Method (CVM) (Squire and Trapp, 1998). Box or column models are used when applicable to maximize the number of comparison pairs from the backward and forward sensitivity test runs (Hakami et al., 2007). The finite difference method is straightforward to implement and has been used for evaluating other sensitivity methods (Dunker et al., 2002; Napelenok et al., 2006; Hakami et al., 2007; Henze et al., 2007), but in some of the cases requires finding optimal step sizes to obtain accurate results. The searching process is time-consuming and could be impractical when a large data set such as that
10 of a three-dimensional atmospheric model is involved. With the CVM, the results are practically insensitive to the perturbation size with the exception of some rare circumstances discussed in Section 3.2.1.

The air quality simulation scenario used for the evaluation is for the contiguous U.S. domain with a 36-km horizontal resolution and 24 vertical layers for the first seven days of April 2008, with the first six days used for spin-up. More details about
15 meteorological inputs, initial and boundary conditions, and emissions are provided in Turner et al. (2015a), where a version of the dataset with 12-km horizontal resolution was used. Unless otherwise noted, evaluations are done for daylong adjoint simulations, with single forcing at the last time step.

3.1 The complex variable method

The CVM for the first-order sensitivities is formulated as follows, which is a Taylor series expansion about an imaginary
20 perturbation step,

$$J(p + ih) = J(p) + J'(p)(ih) - J''(p)h^2/2 + \dots \quad (8)$$

where, J is the cost function, p is the parameter to which the gradient is evaluated, i the imaginary unit, and h is the perturbation step.

25 Extracting only the imaginary part and rearranging,

$$J'(p) = \Im(J(p + ih))/h + O(h^2) \quad (9)$$

where, \Im represents the operator to extract the imaginary part of a complex number. As seen in Equation 8, the CVM has a second-order accuracy, which is the same as the central finite difference formulation. The CVM, however, is not subject to

subtractive errors and therefore permits the use of as small a step size as allowed in floating-point calculations to achieve much better accuracy, which helps in situations when the brute force FDM fails or proves inaccurate.

5 To construct a CVM version of CMAQ, several guidelines are followed which include changing the data types of all the active variables from REAL/DOUBLE PRECISION to the corresponding complex type, creating a complex version of intrinsic functions such as MAX, MIN and ABS, and evaluating only the real part of complex variables used in conditionals (Giles and Piece, 2000). The original 3D CMAQ framework is set up for testing with the processes or sub-processes under investigation enabled and the rest commented out. To run the CVM, a perturbation is added to the imaginary part of a source variable at a time step of interest and then the imaginary part of a receptor is extracted and divided by the perturbation size to obtain the
10 CVM sensitivity.

3.2 Process-by-process model evaluation

3.2.1 Aerosols

As mentioned before, the CMAQ aerosol module incorporates the following science processes: SOA formation, nucleation, condensation, coagulation, heterogeneous chemistry, mode merging, and aerosol thermodynamics. The sub-processes are
15 evaluated individually and eventually as a whole in simulations in which other processes (e.g., advection) are turned off. Adjoint sensitivities are first compared with those from the FDM, and if a mismatch persists, the CVM method is implemented (if feasible) for that process for further evaluation.

Figure 1 shows the adjoint (ADJ) and CVM sensitivities of an example SOA process. The sensitivities are of the final
20 concentrations of an accumulation-mode aerosol species, AALKJ ($\mu\text{g}/\text{m}^3$), with respect to the initial concentrations of a semi-volatile species, SV_ALK (ppmV), from a one-day test run. For this process, the FDM behaved well for most of the test cases (results not shown); in the few cases when the ADJ and FDM did not agree and tuning with the perturbation sizes did not help, use of CVM demonstrated good accuracy of adjoint results (i.e., agreement along one-on-one line). It should be noted that the adopted perturbation size for the CVM is generally 10^{-12} ; a smaller perturbation size usually does not improve accuracy of the
25 obtained sensitivities, but risks diminishing the sensitivity information through propagation due to the single precision nature of some of the variables within CMAQ. A smaller perturbation size could be used if double-precision data types were adopted for the whole code. For the SOA process, however, a smaller perturbation size of 10^{-24} does improve the accuracy. Test results of the other organic aerosol species show similar accuracy.

30 The example given above is one of numerous cases where FDM was found to be inaccurate or inadequate in evaluating adjoint sensitivities. The inadequacy of FDM in producing accurate sensitivity estimates is due to process nonlinearities, as well as discontinuities that exist throughout CMAQ. This is the case in a number of CMAQ processes such as SOA formation,

Deleted: This is an example of numerous cases where FDM was found to be inaccurate or inadequate in evaluating adjoint sensitivities. ...

inorganic thermodynamics, clouds, aqueous chemistry, advection, etc. This issue is not limited to CMAQ alone and exists in all air quality models, as providing a smooth solution for the governing equations may be lost in trade-offs for added computational efficiency, improving stability, or reducing numerical artifacts in the development stage.

5 Generally good spatial agreement is observed between the ADJ and the CVM when all aerosol sub-processes except thermodynamics (ISORROPIA) are included in a one-day run (Figure 2). The cost function is the final concentrations of an accumulation-mode aerosol species ASO4J ($\mu\text{g}/\text{m}^3$) and the perturbation variable is the initial concentrations of an Aitken-mode aerosol species ASO4I ($\mu\text{g}/\text{m}^3$) with a perturbation of 10^{-12} for the CVM. One of the reasons to choose the two model species for testing was that sulphate aerosol is a crucial component of fine aerosols. The other reason was that the pathway
10 from ASO4I to ASO4J covers the size range for the important processes of coagulation and the numerical mode merging that handle interactions between the Aitken and accumulation modes.

With the addition of the aerosol thermodynamics (ISORROPIA/ANISORROPIA), the degree of agreement is degraded for few points across the domain (Figure 3, left panel). The disagreement is likely caused by inherent nonlinearities, as well as
15 discontinuities in the solution surfaces of ISORROPIA, which at the code level is manifested as a series of executive branches (Capps et al., 2012). As shown in Equation 8, the introduction of the imaginary part, ih , changes the value of the real part of a complex variable. Although the change is of $O(h^2)$, different executive branches can be invoked as a result. When extra care is taken for the ADJ and CVM calculations to follow consistent branches, the agreement is much improved (Figure 3, right panel).

20 As an example for the full aerosol model application, the sensitivities of final ASO4J ($\mu\text{g}/\text{m}^3$) with respect to initial NH_3 (ppmV) from a one-day simulation are evaluated against CVM estimates (Figure 4). Although the branches in ISORROPIA are made identical between the ADJ and the CVM, the agreement, while still good, deteriorates for few points when compared to Figure 3; reducing the perturbation size does not affect the agreement. One possible explanation is the high nonlinearity
25 associated with the aerosol thermodynamics (which increases as the relative humidity drops), that renders the CVM with a finite perturbation size ineffective in producing accurate sensitivities. A more detailed discussion of the discontinuities and nonlinearity of ISORROPIA is given in Capps et al. (2012). Overall, our testing confirms the findings in Capps et al. (2012) that the CVM implementation of ISORROPIA produces approximations that agree with the adjoint results.

3.2.2 Cloud dynamics and chemistry

30 The adjoint of mixing due to formation of sub-grid convective clouds was tested by comparing FDM sensitivities (CVM was not available for all cloud processes) to adjoint-calculated sensitivities in a one-day test run. The test, shown in Figure 5 for a perturbation of $0.01 \mu\text{g}/\text{m}^3$ in initial ASO4J, was successful but exhibited sensitivity to the type and size of the perturbation as discussed below. For the FDM, either a percentage perturbation or a perturbation with an absolute small value can be used.

Deleted: bottom

Deleted: the discontinuities abundant

Deleted: When the ADJ and CVM sensitivities from the same branches are compared

Deleted: is worse than what is seen in

Deleted: inherent within the aerosol processes including ISORROPIA

Deleted: Overall, the sensitivities generated from the aerosol adjoint show acceptable accuracy when compared with the CVM.

For test cases with small initial concentration values, however, percentage perturbations are sometimes not detectable at the end of run (due to round-off errors), which leads to diminished sensitivity values. The non-linearity introduced by the weighting functions implemented for re-distribution of gas and aerosol concentrations in convective mixing made this issue stand out; although not apparent in Figure 5, the FDM constantly failed to produce some larger values observed in the adjoint sensitivity field. On the other hand, perturbations with an absolute value tend to undermine the accuracy of the FDM sensitivities by causing truncation errors that are comparable to the size of perturbation. Absolute perturbations have been generally favoured and used in the evaluations of cloud processes. The adjoint of the resolved clouds module shows a better agreement with the FDM (not shown) than that of the convective clouds.

10 The adjoint of the aqueous chemistry science process was tested similarly, with a $0.01 \mu\text{g}/\text{m}^3$ perturbation in initial ASO4J. Figure 6 shows a successful test, with close agreement between the ADJ and the FDM sensitivities. Aqueous chemistry was further tested by examining the sensitivity of final ASO4J to small (0.01/0.001/0.0001 ppb) perturbations in gas-phase SO_2 and results are shown in Figure 7. Mismatches were apparent in all three plots between the ADJ and the FDM, especially the column with high values of FDM sensitivities and zero ADJ sensitivities. As discussed above, small initial values could render
15 the FDM approximation difficult. A way to quickly check whether these small values are causing disagreement is to semi-normalize the sensitivities by multiplying with the initial values. The results from semi-normalization with initial SO_2 concentrations confirm that the deviations were caused by small initial values (Figure 7, bottom).

3.2.3 Gas-phase chemistry

ADJ and FDM sensitivities of the final O_3 concentrations (ppmV) with respect to the initial NO_2 concentrations (ppmV) from
20 a one-day test with the updated CB05 chemical mechanism show good agreement, depending on the choice of appropriate perturbation size (Figure 8). Results of three perturbation sizes 0.1, 0.01 and 0.001 ppb are shown for the FDM, which demonstrate the impact of perturbation sizes on accuracy.

Mismatches in some test cases led to the development of a CVM for the gas-phase chemistry. During the testing, Jacobians of
25 ADJ and CVM sensitivities are created to provide a means for visual examination of all gas-phase species. An example is shown in Figure 9, where the x and y axes represent all species involved in gas-phase chemical reactions, and each point represents the sensitivity of an x -axis species with respect to a y -axis species. Presented in Figure 10 is the corresponding scatter plot which compares the ADJ with the CVM and shows an excellent match between the two methods. The absolute sensitivity values are used in the tile plot in Figure 9 for better visualization.

30 3.2.4 Transport

The process of transport includes advection and diffusion. For the advection process, the nonlinear PPM scheme is implemented in CMAQ as discussed in 2.3.4.

The adjoint of the advection equation could be written as,

$$-\frac{\partial \lambda_i}{\partial t} = \nabla \cdot (\mathbf{u} \lambda_i). \quad (10)$$

5 Compared to the corresponding terms in Equation 1, the signs of the two terms in the above equation are reversed, which implies that the adjoint values are propagated in an opposite direction (Giles and Pierce, 2000). With a reversal in wind direction, the PPM could be used to integrate the adjoint equation backward in time and solve for continuous adjoint sensitivities.

10 The discrete adjoint of advection is the result of direct differentiation of the numerical model and can be validated against the FDM/CVM component by component. Figure 11 of horizontal advection in x direction shows a good agreement between the discrete adjoint and the CVM of sensitivities of final ASO4J ($\mu\text{g}/\text{m}^3$) with respect to initial ASO4J ($\mu\text{g}/\text{m}^3$). However, numerical noises are clearly in sight; unfortunately, such spurious oscillations from discrete adjoints derived from a nonoscillatory advection scheme are not uncommon and a desirable fix does not appear possible (Thuburn and Haine, 2001).

15 For the continuous adjoint (CADJ), the horizontal and vertical advection processes were tested as a whole, as in forward CMAQ these processes are linked together for mass conservation and consistency between transport processes in CMAQ and the underlying meteorological model. One issue with testing the advection processes altogether is that pointwise comparison of sensitivities becomes much more computationally expensive as the models are not row or column models anymore, and running the adjoint and the FDM/CVM once would generate only one pair of sensitivities for comparison. To partially remediate this situation, we defined the cost function for the adjoint as the final average ASO4J concentration across the entire surface layer (instead of concentration at a single cell which would lead to a small number of cells with sensitivity signals over time, not sufficient for validation), and then randomly selected a number of cells at the surface for the CVM runs. The CADJ and the CVM agree well as shown in Figure 12 where the regression line has a slope close to unity and y -intercept close to zero, and the value of R^2 is 0.957.

25 The choice between the continuous and discrete adjoint would depend on the type of problem at hand. For instance, the continuous adjoint is generally desirable when performing backward sensitivity analysis as an oscillating sensitivity field (visible in Figure 11) may defy physical justification (Hakami et al., 2007; Henze et al., 2007). For optimization problems, the discrete adjoint would be preferable as it produces exact gradients (subject to round-off errors) that could help the optimization process converge (Giles and Pierce, 2000). However, it was reported that the noisy gradient field obtained from the discrete adjoint could cause the optimization to converge to local minimums (Vukićević et al., 2001). The continuous adjoint may outperform the discrete in terms of computational efficiency and accuracy, as found by Gou and Sandu (2011) with their 4D-Var experiments.

Air pollutant emissions are processed in vertical diffusion in CMAQ. As shown in Figure 13a, the adjoint of vertical diffusion compares well with the finite difference method (Figure 13a). The adjoint of emissions also works well as demonstrated in Figure 13b, where the adjoint sensitivities of final NO₂ concentrations (ppmV) to initial NO₂ emissions (moles/s) compare favourably with the FDM.

5 3.3 Full-model evaluation

For the full adjoint model, interactions between cells through transport make it prohibitively costly to generate sufficient sensitivity pairs for an extensive comparison as conducted for the process-by-process evaluation with box or column models. Presented in Table 1 are the FDM, CVM, adjoint with continuous treatment for transport and adjoint with discrete treatment for transport (DADJ) sensitivities obtained for a few grid cells.

In general, the adjoint models, particularly DADJ, agree well with the CVM, while in the case of CADJ a larger relative error exists in comparison with the CVM. The FDM sensitivities with a 10% perturbation step on the other hand are not quite in accordance, which is why a full CVM was created. The problem with the FDM has been discussed at the beginning of Section 3.4 and is not repeated here.

Deleted: agrees well with the CVM, especially the DADJ with a relative error less than 10%

Deleted: earlier

Results shown in Table 1 suggest that the discrete adjoint has a better agreement with CVM than the continuous adjoint. However, it is important to note that better agreement between the discrete adjoint and CVM should not be understood as better accuracy of the discrete adjoint in comparison with continuous adjoint. The numerical solution to the advection equation entails inherent truncation errors from discretization schemes. These errors exist in solving the forward or adjoint advection equations; however the discrete adjoint by design remains loyal to, and consistent with the errors in the forward application (CVM in this case), while numerical solution to the continuous adjoint will result in different and inconsistent errors. The continuous adjoint is a different representation of the impacts on the adjoint cost function, but of similar mathematical accuracy, when compared to the forward or tangent linear model; therefore, the numerical solution to it should be considered as accurate as the discrete adjoint, regardless of the agreement with forward-based benchmarks such as CVM.

25 3.4 Computational system requirement

Adjoint simulations entail a significantly higher computational demand than forward CMAQ. First, the checkpointing files required for the adjoint simulations need a significant amount of storage space. For each science process the amount of storage can be estimated as $N_c \times N_r \times N_l \times N_s \times N_t \times N_b$ bytes, where N_c , N_r , N_l , N_s , N_t , and N_b represent the numbers of columns, rows, layers, chemical species, synchronization time steps, and bytes for a single-precision number ($N_b = 4$), respectively. For our computational domain with 148 columns, 112 rows, 24 vertical layers and 12-minute synchronization (i.e., 120 time steps/day), the checkpointing file for each adjoint simulation for aerosols with 137 chemical species takes approximately 24 GB storage space for a day. For the other science processes the sizes of checkpointing files are approximately as follows:

clouds, 24 GB; chemistry with 96 species, 17 GB; vertical diffusion with 1 layer, 1 GB; the continuous adjoint of horizontal and vertical advections with 1 species, 0.2 GB for each process. [A summary of the checkpointing file sizes is provided in Table 2.](#) For a one-month adjoint simulation, the checkpointing files occupy about 2 TB of storage space, which is about 10 times the storage needed for typical CMAQ simulations. For higher resolution simulations, the storage needs would increase proportionally due to the number of grid cells, but also due to increasingly smaller time steps dictated by the Courant number for smaller horizontal grid sizes. For very high resolution simulations (e.g., 1 km horizontal resolution), the required checkpoint storage space can be as large as 1 TB/day. To mitigate the burden on storage, it is plausible to run the adjoint segment by segment, i.e., by generating the checkpointing files only for a few days at a time when running the forward CMAQ model. Since the adjoint runs backward in time, this strategy works at the expense of CPU time as the forward model must be repeated.

Adjoint simulations also require significantly higher CPU times, as the re-computation of intermediate values of adjoint simulations as discussed in Section 2.4 is an additional computational overhead. Furthermore, the significant amount of input/output (I/O) operation associated with the checkpointing, leads to additional CPU time and can result in noticeable loss of computational efficiency in systems. Typically, the adjoint simulation takes approximately 4 times as long as the forward CMAQ. Intensive I/O for checkpointing files can also result in reduced scalability of the adjoint model, as the I/O libraries currently implemented are serial.

4 Model Application: Backward Sensitivity Analysis

To demonstrate a policy-relevant application of the multiphase adjoint of CMAQ, we estimate [Marginal Benefits \(MB\) or Benefit-Per-Ton \(BPT\) of emission sources](#). MBs or BPTs are a commonly used measure of source impacts on population health, and are defined as the monetized societal benefits associated with the reduction of one metric ton of precursor or primary emissions. [More details on the use of adjoint models in source attribution of health impacts can be found elsewhere](#) (Pappin et al., 2013 and 2016; [Turner et al., 2015a and 2015b](#)). Mortality counts, or nationwide valuation of mortality induced by air pollution is a scalar well suited for formulating as the adjoint cost function. We define an adjoint cost function, J, which represents the monetized valuation of annual deaths due to long-term PM_{2.5} exposure within the U.S., as below, [using the Global Exposure Mortality Model \(GEMM; Burnett et al., 2018\)](#):

$$J = \sum_{x,y} V_{SL} M_{0,x,y} P_{x,y} (1 - e^{-\theta T(z)})$$

$$T(z) = \frac{\log(1 + \frac{z}{\alpha})}{1 + \exp(-(z-\mu)/\nu)} \quad (11)$$

$$z = \max(0, C_{x,y} - cf)$$

where J (\$/yr) is calculated using the location-specific baseline mortality rate ($M_{0,x,y}$; yr⁻¹) and population ($P_{x,y}$); θ as the risk estimate that represents the slope of the [non](#)linear concentration-mortality curve; the time-averaged and location-specific PM_{2.5}

Deleted: the impacts of emission sources on adverse PM_{2.5} health impacts across the US in an adjoint sensitivity analysis; more details of an adjoint-based approach for

Deleted: $J = \sum_{x,y} V_{SL} M_{0,x,y} P_{x,y} (1 - e^{-\beta C_{x,y}}) \quad (11)$

Deleted: β

Deleted: log-

concentration, $C_{x,y}$ ($\mu\text{g}/\text{m}^3$); cf is the counterfactual $\text{PM}_{2.5}$ concentration; and the value of a statistical life, V_{SL} for monetizing outcomes. We use the following values for the parameters of GEMM for adults aged 25-99 and for deaths due to noncommunicable diseases and lower respiratory infections: $\theta=0.1231$, $\alpha = 1.5$, $\mu=10.4$, $\nu=2.5.9$, and $cf=2.4\mu\text{g}/\text{m}^3$.

Deleted: a β of 0.5827% per $\mu\text{g}/\text{m}^3$ from Krewski et al. (2009) for adults aged 30-99 and for deaths due to all causes. We use the simulation period-average $\text{PM}_{2.5}$ concentration in place of the annual average due to our shorter episode duration, while recognizing the limitations of extrapolating from April to the full year.

5 For the backward sensitivity analysis, we run the adjoint for the year of 2016 for the contiguous U.S. domain with 36-km resolution inputs from the Intermountain West Data Warehouse (National Emissions Inventory Collaborative, 2019). The computational domain contains 172 columns and 148 rows with 35 vertical layers. The average $\text{PM}_{2.5}$ concentration is obtained from the one-year forward run. Backward runs are conducted for two full seasons of winter and summer (Jan-Mar, and Jul-Sep, respectively).

10 The use of this cost function results in gradient estimates that can be presented as location-specific BPTs ($\$/\text{yr}$). BPTs for primary $\text{PM}_{2.5}$ emitted across source locations in the US, estimated from the two seasons are shown in Figure 14, as well as those for the $\text{PM}_{2.5}$ inorganic precursors NH_3 , SO_2 , and NO_x . The seasonal differences in BPTs, particularly for precursor emissions such as ammonia, are apparent and significant. NO_x sensitivities are negative in some regions with more frequent

Deleted: provide a measure of how reducing emissions by 1 metric ton would result in societal benefits across the U.S., referred to as marginal benefits or benefits-per-ton (BPT, $\$/\text{yr}$). Benefits-per-ton

15 NO_x -inhibited regimes, mainly due to the role that ozone plays in night time nitrate formation. BPTs show a great deal of spatial variability but generally follow the population distribution for primary $\text{PM}_{2.5}$ emissions, while for inorganic precursor emissions areas of higher influence are dictated by transport patterns, secondary (inorganic) aerosol formation dynamics, and lifetime of secondary particles. In other words, BPTs are generally highest in emission locations that have large potential for affecting downwind population centers. While the adjoint cost function is defined based on $\text{PM}_{2.5}$ long-term mortality in the
20 US alone, location-specific BPTs also provide a measure of cross-border impact. Finally, we note that BPTs are measures of marginal rather than total societal impact across the U.S., and as such, even areas with little or no emissions may show sizeable BPT estimates.

Deleted: . NO

25 BPTs values such as those shown in Figure 14 have the potential to form important quantitative decision metrics, as they provide a means to squarely compare societal benefits of emission reductions with control costs associated with those reductions. It is worth emphasizing that given the length and the coarse resolution of the simulations, these results should not be regarded as BPT values applicable in policy development and benefit assessment; instead they are meant to serve as a demonstration of the utility and efficacy of the adjoint model to attribute health impacts to individual sources.

Deleted: the short simulation period

Deleted: viewed

Deleted: that can be used

Deleted: ,

30 5 Conclusions

In this paper, we develop a multiphase adjoint of CMAQ. A rigorous point-to-point evaluation against the brute force FDM and CVM is conducted for each individual process and the full model with all processes included. Overall, the adjoints appear

to produce sensitivities comparable to those generated by either the FDM or the CVM. The choice of the discrete or continuous version of the advection adjoint would depend on the type of problem to be solved. The continuous adjoint is preferred if the sensitivity field itself is of interest, as spurious oscillations would create intricate obstacles for exploring the underlying physical significances. For gradient-based optimization and data assimilation, the discrete adjoint might be advantageous for faster convergence but could risk the minimization settling upon some local minima. Some components of CMAQ that do not yet have an adjoint include the calculations of bidirectional dry deposition velocities in vertical diffusion and photolysis rate calculations in gas-phase chemistry. The development of an adjoint for these two components is not considered essential. The CMAQ adjoint provides backward sensitivity analysis capabilities for a widely used model with detailed aerosol treatment, and enables a range of applications data assimilation, emission inversions, policy analysis, and source attribution of health impacts.

Deleted: s

We find that the development of adjoint versions of air quality and atmospheric models is often complicated by the abundance of discontinuities throughout these models that make differentiation challenging. Historically, these models have not been developed with differentiability in mind, but with accuracy and computational efficiency as the main drivers. As the development and applications of formal sensitivity analysis tools (such as adjoint models) become more prevalent, there is a need for a gradual but sustained effort by the modeling community to consider differentiability as an additional design constraint in future developments.

6 Code Availability

Upon completion of expanded user testing, the CMAQ Adjoint code will be hosted and distributed by the U.S. EPA.

20 7 Author Contribution

SZ contributed to the development of aerosol dynamics, clouds adjoint, aqueous chemistry, coordinated adjoint code assembly and testing, and prepared the manuscript. AH led the team and developed chemistry module. SC developed the ANIORROPIA, and contributed to development of the forcing generator and preparation of user manual. MT contributed to the development of the adjoint of aerosol dynamics. PP developed the adjoint of transport. KF provided the KPP configuration of aqueous chemistry. JR developed the checkpointing modules. MR developed the adjoint forcing generator. SH contributed 25 to the debugging of vertical transport. DH, AR, AN, TC, CS, and GC provided feedback on adjoint development. AP performed adjoint health impact analysis. SN and JB provided support for existing and new CMAQ modules. All co-authors contributed to the preparation of the manuscript.

8 Acknowledgements

Hakami acknowledges support from the Natural Sciences and Engineering Research Council of Canada Grant RGPIN-2016-06181, Health Canada Contract 4500383370, and the Health Effects Institute Agreement 4962-RFA17-2/18-4. Henze received support from NASA grant NNX16AQ26G. Simulations were performed on computational resources of Compute Canada and

5 Compute Ontario.

References

- Anderson, W. K., and Nielsen E.: Sensitivity Analysis for Navier–Stokes Equations on Unstructured Grids Using Complex Variables, *AIAA J.*, 39(1), 56–63, doi:10.2514/2.1270, 2001.
- Bartholomew-Biggs, M. C.: Using forward accumulation for automatic differentiation of implicitly-defined functions, *Comput. Optim. Appl.*, 9(1), 65–84, doi:10.1023/A:1018382103801, 1998.
- Baydin, A. G., Pearlmutter, B. A., and Siskind, J. M.: Tricks from deep learning, arXiv:1611.03777, 2016.
- Binkowski, F. S., and Roselle, S. J.: Models-3 Community Multiscale Air Quality (CMAQ) model aerosol component 1. Model description, *J. Geophys. Res. Atmos.*, 108(D6), 4183, doi:10.1029/2001JD001409, 2003.
- Brook, R. D., Rajagopalan, S., Pope III, C. A., Brook, J. R., Bhatnagar, A., Diez-Roux, A. V., Holguin, F., Hong, Y., Luepker, R. V., Mittleman, M. A., Peters, A., Siscovick, D., Smith, Jr, S. C., Whitsel, L., Kaufman, J. D.: Particulate matter air pollution and cardiovascular disease: an update to the scientific statement from the American Heart Association, *Circ.*, 121(21), 2331–2378, doi:10.1161/CIR.0b013e3181dbee1, 2010.
- [Burnett R., Chen H., Szyszkwicz M., Fann N., Hubbell B., Pope III C. A., Apte J. S., Brauer M., Cohen A., Weichenthal S., Coggin J., Di Q., Brunekreef B., Frostad J., Lim S. S., Kan H., Walker K. D., Thurston G. D., Hayes R. B., Lim C. C., Turner M. C., Jerrett M., Krewski D., Gapstur S. M., Diver W. R., Ostro B., Goldberg D., Crouse D. L., Martin R. V., Peters P., Pinault L., Tjepkema M., van Donkelaar A., Villeneuve P. J., Miller A. B., Yin P., Zhou M., Wang L., Janssen N. A. H., Marra M., Atkinson R. W., Tsang H., Quoc Thach T., Cannon J. B., Allen R. T., Hart J. E., Laden F., Cesaroni G., Forastiere F., Weinmayr G., Jaensch A., Nagel G., Concin H., Spadaro J. V.: Global estimates of mortality associated with long-term exposure to outdoor fine particulate matter. *Proc. Natl. Acad. Sci.*, 115\(38\), 9592–9597, doi: 10.1073/pnas.1803222115, 2018.](#)
- Byun, D. W.: Dynamically Consistent Formulations in Meteorological and Air Quality Models for Multi-Scale Atmospheric Applications, Part I: Governing Equations in Generalized Coordinate System, *J. Atmos. Sci.*, 56, 3789–3807, doi: 10.1175/1520-0469(1999)056<3789:DCFIMA>2.0.CO;2, 1999.
- Byun, D. W., and Schere, K. L.: Review of the governing equations, computational algorithms, and other components of the Models-3 Community Multiscale Air Quality (CMAQ) modeling system, *Appl. Mec. Rev.*, 59(2), 51–77, doi:10.1115/1.2128636, 2006.
- Capps, S. L., Henze, D. K., Hakami, A., Russell, A. G., and Nenes, A.: ANISORROPIA: the adjoint of the aerosol thermodynamic model ISORROPIA, *Atmos. Chem. Phys.*, 12, 527–543, doi:10.5194/acp-12-527-2012, 2012.
- Carmichael, G. R., Sandu, A., Chai, T., Daescu, D. N., Constantinescu, E. M., and Tang, Y.: Predicting air quality: improvements through advanced methods to integrate models and measurements, *J. Comput. Phys.*, 227(7), 3540–3571, doi:10.1016/j.jcp.2007.02.024, 2008.
- Chai, T., Carmichael, G. R., Tang, Y., Sandu, A., Hardesty, M., Pilewskie, P., Whitlow, S., Browell, E. V., Avery, M. A., Nédélec, P., Merrill, J. T., Thompson, A. M., and Williams, E.: Four-dimensional data assimilation experiments with

- International Consortium for Atmospheric Research on Transport and Transformation ozone measurements, *J. Geophys. Res. Atmos.*, 112(D12S15), doi:10.1029/2006JD007763, 2007.
- Cohan, D. S., Hakami, A., Hu, Y., and Russell, A. G.: Nonlinear response of ozone to emissions: source apportionment and sensitivity analysis, *Environ. Sci. Technol.*, 39(17), 6739-6748, doi:10.1021/es048664m, 2005.
- 5 Colella, P., and Woodward, P. R.: The Piecewise Parabolic Method (PPM) for gas-dynamical simulations, *J. Comput. Phys.* 54(1), 174-201, doi:10.1016/0021-9991(84)90143-8, 1984.
- [Constantin, B. V. and Barrett, S. R.: Application of the complex step method to chemistry-transport modeling. *Atmos. Environ.*, 99, 457-465, 10.1016/j.atmosenv.2014.10.017, 2014.](#)
- Crouse, D. L., Peters, P. A., Hystad, P., Brook, J. R., van Donkelaar, A., Martin, R. V., Villeneuve, P. J., Jerrett, M., Goldberg, M. S., Pope III, C. A., and Brauer, M.: Ambient PM_{2.5}, O₃, and NO₂ exposures and associations with mortality over 16 years of follow-up in the Canadian Census Health and Environment Cohort (CanCHEC), *Environ. Health Perspect.*, 123(11), 1180-1186, doi: 10.1289/ehp.1409276, 2015.
- 10 Crouse, D. L., Peters, P. A., van Donkelaar, A., Goldberg, M. S., Villeneuve, P. J., Brion, O., Khan, S., Atari, D. O., Jerrett, M., Pope III, C. A., and Brauer, M.: Risk of nonaccidental and cardiovascular mortality in relation to long-term exposure to low concentrations of fine particulate matter: a Canadian national-level cohort study, *Environ. Health Perspect.*, 120(5), 708-714, doi: 10.1289/ehp.1104049, 2012.
- Damian, V., Sandu, A., Damian, M., Potra, F., and Carmichael, G. R.: The kinetic preprocessor KPP-a software environment for solving chemical kinetics, *Comput. Chem. Eng.*, 26(11), 1567-1579, doi:10.1016/S0098-1354(02)00128-X, 2002.
- Di, Q., Wang, Y., Zanobetti, A., Wang, Y., Koutrakis, P., Choirat, C., Dominici, F., and Schwartz, J. D.: Air pollution and mortality in the Medicare population, *N. Engl. J. Med.*, 376(26), 2513-2522, doi: 10.1056/NEJMoa1702747, 2017.
- 20 Dubovik, O., Lapyonok, T., Kaufman, Y. J., Chin, M., Ginoux, P., Kahn, R. A., and Sinyuk, A.: Retrieving global aerosol sources from satellites using inverse modeling, *Atmos. Chem. Phys.*, 8(2), 209-250, doi:10.5194/acp-8-209-2008, 2008.
- Dunker, A. M., Yarwood, G., Ortman, J. P., and Wilson, G. M.: The decoupled direct method for sensitivity analysis in a three-dimensional air quality model implementation, accuracy, and efficiency, *Environ. Sci. Technol.*, 36(13), 2965-2976, doi:10.1021/es0112691, 2002.
- 25 Elbern, H., Schmidt, H. and Ebel, A.: Variational data assimilation for tropospheric chemistry modeling, *J. Geophys. Res. Atmos.*, 102(D13), 15967-1598, doi:10.1029/97JD01213, 1997.
- Elbern, H., Schmidt, H., Talagrand, O., and Ebel A.: 4D-variational data assimilation with an adjoint air quality model for emission analysis, *Environ. Model. Softw.*, 15, 539-548, doi:10.1016/S1364-8152(00)00049-9, 2000.
- 30 Errico, R. M.: What is an adjoint model?, *Bull. Amer. Meteor. Soc.*, 78(11), 2577-2591, doi:10.1175/1520-0477(1997)078<2577:WIAAM>2.0.CO;2, 1997.
- ENVIRON: User's Guide, Comprehensive Air Quality Model with Extensions (CAMx), Version 4.20, ENVIRON International Corporation, Novato, CA, available at <http://www.camx.com>, 2005.

- Fahey, K. M., Carlton, A. G., Pye, H. O. T., Baek, J., Hutzell, W. T., Stanier, C. O., Baker, K. R., Appel, K. W., Jaoui, M., and Offenberg, J. H.: A framework for expanding aqueous chemistry in the Community Multiscale Air Quality (CMAQ) model version 5.1, *Geosci. Model Dev.*, 10(4), 1587–1605, doi:10.5194/gmd-10-1587-2017, 2007.
- Fiore, A. M., Naik, V., Spracklen, D. V., Steiner, A., Unger, N., Prather, M., Bergmann, D., Cameron-Smith, P. J., Cionni, I.,
5 Collins, W. J., Dalsøren, S., Eyring, V., Folberth, G. A., Ginoux, P., Horowitz, L. W., Josse, B., Lamarque, J.-F., MacKenzie, I. A., Nagashima, T., O'Connor, F. M. O., Righi, M., Rumbold, S. T., Shindell, D. T., Skeie, R. B., Sudo, K., Szopa, S., Takemura, T., and Zeng, G.: Global air quality and climate, *Chem. Soc. Rev.*, 41(19), 6663–6683, doi:10.1039/C2CS35095E, 2012.
- Fisher, M., and Auvinen, H.: Long window 4D-Var, *Proc. ECMWF Seminar on Data Assimilation for Atmosphere and Ocean*,
10 189–202, 2011.
- Fisher, M., and Lary, D. J.: Lagrangian four-dimensional variational data assimilation of chemical species, *Q. J. R. Meteorol. Soc.*, 121(527), 1681–170, doi:10.1002/qj.49712152709, 1995.
- Foley, K. M., Roselle, S. J., Appel, K. W., Bhave, P. V., Pleim, J. E., Otte, T. L., Mathur, R., Sarwar, G., Young, J. O., Gilliam, R. C., Nolte, C. G., Kelly, J. T., Gilliland, A. B., and Bash, J. O.: Incremental testing of the Community Multiscale Air Quality
15 (CMAQ) modeling system version 4.7, *Geosci. Model Dev.*, 3(1), 205–226. doi:10.5194/gmd-3-205-2010, 2010.
- Fuzzi, S., Baltensperger, U., Carslaw, K., Decesari, S., Denier Van Der Gon, H., Facchini, M. C., Fowler, D., Koren, I., Langford, B., Lohmann, U., Nemitz, E., Pandis, S., Riipinen, I., Rudich, Y., Schaap, M., Slowik, J. G., Spracklen, D. V., Vignati, E., Wild, M., Williams, M., and Gilardoni, S.: Particulate matter, air quality and climate: lessons learned and future needs, *Atmos. Chem. Phys.*, 15(14), 8217–8299, doi:10.5194/acp-15-8217-2015, 2015.
- 20 Giering, R., and Kaminski, T.: Recipes for adjoint code construction, *ACM Trans. Math. Softw.*, 24(4), 437–474, doi:10.1145/293686.293695, 1998.
- Giering, R.: Tangent linear and Adjoint Model Compiler, Users manual 1.4, <http://www.autodiff.com/tame>, 1999.
- Giles, M. B., and Pierce, N. A.: An introduction to the adjoint approach to design, *Flow Turbul. Combust.*, 65(3–4), 393–415, doi:10.1023/A:101143041, 2000.
- 25 Giles, M. B., Duta, M. C., Müller, J. D., and Pierce, N. A.: Algorithm developments for discrete adjoint methods, *AIAA J.*, 41(2), 198–205, doi:10.2514/2.1961, 2003.
- Gou, T., and Sandu, A.: Continuous versus discrete advection adjoints in chemical data assimilation with CMAQ, *Atmos. Environ.*, 45(28), 4868–4881, doi:10.1016/j.atmosenv.2011.06.015, 2011.
- Griewank, A.: A mathematical view of automatic differentiation, *Acta Num.*, 12, 321–398, doi:10.1017/S0962492902000132,
30 2003.
- Griewank, A.: Who Invented the Reverse Mode of Differentiation?, *Documenta Math.*, Extra Volume ISMP, 389–400, 2012.
- Hakami, A., Odman, M. T. and Russell, A. G.: High-order, direct sensitivity analysis of multidimensional air quality models, *Environ. Sci. Technol.*, 37(11), 2442–2452, doi:10.1021/es020677h, 2003.

- Hakami, A., Henze, D. K., Seinfeld, J.H., Chai, T., Tang, Y., Carmichael, G. R., and Sandu, A.: Adjoint inverse modeling of black carbon during the Asian Pacific Regional Aerosol Characterization Experiment, *J. Geophys. Res. Atmos.*, 110(D14301), doi:10.1029/2004JD005671, 2005.
- Hakami, A., Henze, D. K., Seinfeld, J. H., Singh, K., Sandu, A., Kim, S., Byun, D. and Li, Q.: The adjoint of CMAQ, *Environ. Sci. Technol.*, 41(22), 7807-7817, doi:10.1021/es070944p, 2007.
- Hascoët, L., and Pascual, V.: The Tapenade Automatic Differentiation tool: principles, model, and specification. *ACM Trans. Math. Softw.*, 39(3), 20, doi:10.1145/2450153.2450158, 2013.
- Henze, D. K., Seinfeld, J. H., Liao, W., Sandu, A., and Carmichael, G. R.: Inverse modeling of aerosol dynamics: condensational growth, *J. Geophys. Res. Atmos.*, 109(D14201), doi:doi:10.1029/2004JD004593, 2004.
- Henze, D. K., Hakami, A., and Seinfeld, J. H.: Development of the adjoint of GEOS-Chem, *Atmos. Chem. Phys.*, 7(9), 2413-2433, doi:10.5194/acp-7-2413-2007, 2007.
- Henze, D. K., Shindell, D. T., Akhtar, F., Spurr, R. J., Pinder, R. W., Loughlin, D., Kopacz, M., Singh, K., and Shim, C.: Spatially refined aerosol direct radiative forcing efficiencies, *Environ. Sci. Technol.*, 46(17), 9511-9518, doi:10.1021/es301993s, 2012.
- Huneeus, N., Boucher, O., and Chevallier, F.: Simplified aerosol modeling for variational data assimilation. *Geosci. Model Dev.*, 2(2), 213-229, doi:10.5194/gmd-2-213-2009, 2009.
- Iott, J., Hafka, R. T., and Adelman, H. M., Selecting step sizes in sensitivity analysis by finite differences, NASA TM-86382, 1985.
- Jacobson, M. Z.: *Fundamentals of Atmospheric Modeling*, Cambridge University Press, New York, 2005.
- Jiang, W., Smyth, S., Giroux, É., Roth, H., and Yin, D.: Differences between CMAQ fine mode particle and PM_{2.5} concentrations and their impact on model performance evaluation in the lower Fraser valley, *Atmos. Environ.*, 40(26),4973-4985, doi:10.1016/j.atmosenv.2005.10.069, 2006.
- Karydis, V. A., Capps, S. L., Russell, A. G., and Nenes, A.: Adjoint sensitivity of global cloud droplet number to aerosol and dynamical parameters, *Atmos. Chem. Phys.*, 12, 9041-9055, doi:10.5194/acp-12-9041-2012, 2012.
- Koo, B., Dunker, A. M., and Yarwood, G.: Implementing the decoupled direct method for sensitivity analysis in a particulate matter air quality model, *Environ. Sci Technol.*, 41(8), 2847-2854, doi:10.1021/es0619962, 2007.
- Koo, B., Wilson, G. M., Morris, R. E., Dunker, A. M., and Yarwood, G.: Comparison of source apportionment and sensitivity analysis in a particulate matter air quality model, *Environ. Sci. Technol.*, 43(17), 6669-6675, doi:10.1021/es9008129, 2009.
- Kopplitz, S. N., Mickley, L. J., Marlier, M. E., Buonocore, J. J., Kim, P. S., Liu, T., Sulprizio, M. P., DeFries, R. S., Jacob, D. J., Schwartz, J., and Pongsiri, M.: Public health impacts of the severe haze in Equatorial Asia in September–October 2015: demonstration of a new framework for informing fire management strategies to reduce downwind smoke exposure, *Environ. Res. Lett.*, 11(9), 094023, doi:10.1088/1748-9326/11/9/094023,2016.
- Lacey, F. G., Henze, D. K., Lee, C. J., van Donkelaar, A., and Martin, R. V., Transient climate and ambient health impacts due to national solid fuel cookstove emissions, *Proc. Natl. Acad. Sci.*, 114(6), 1269-1274, doi:10.1073/pnas.1612430114, 2017.

Deleted: Krewski, D., Jerrett, M., Burnett, R. T., Ma, R., Hughes, E., Shi, Y., Turner, M. C., Pope III, C. A., Thurston, G., Calle, E. E. and Thun, M. J.: Extended follow-up and spatial analysis of the American Cancer Society study linking particulate air pollution and mortality (No. 140). Boston, MA: Health Effects Institute, 2009.

- Lee, C. J., Martin, R. V., Henze, D. K., Brauer, M., Cohen, A., van Donkelaar, A.: Response of global particulate-matter-related mortality to changes in local precursor emissions, *Environ. Sci. Technol.*, 49(7), 4335–4344, doi:10.1021/acs.est.5b00873, 2015.
- McRae, G. J., Goodin, W. R., Seinfeld, J. H.: Numerical solution of the atmospheric diffusion equation for chemically reacting flows, *J. Comput. Phys.*, 45(1), 1-42, doi: 10.1016/0021-9991(82)90101-2, 1982.
- Martien, P. T., and Harley, R. A.: Adjoint sensitivity analysis for a three-dimensional photochemical model: application to Southern California, *Environ. Sci. Technol.*, 40(13), 4200-4210, doi:10.1021/es051026z, 2006.
- Mesbah, S. M., Hakami, A., and Schott, S.: Improving NO_x cap-and-trade system with adjoint-based emission exchange rates, *Environ. Sci. Technol.*, 46(21), 11905-11912, doi:10.1021/es302406y, 2012.
- 10 Napelenok, S. L., Cohan, D. S., Hu, Y., and Russell, A. G.: Decoupled direct 3D sensitivity analysis for particulate matter (DDM-3D/PM), *Atmos. Env.*, 40(32), 6112-6121, doi:10.1016/j.atmosenv.2006.05.039, 2006.
- Navon, I.M.: Practical and theoretical aspects of adjoint parameter estimation and identifiability in meteorology and oceanography, *Dynam. Atmos. Ocean*, 27(1-4), 55-79, doi:10.1016/S0377-0265(97)00032-8, 1997.
- Nenes, A., Pandis, S. N., and Pilinis, C.: ISORROPIA: a new thermodynamic equilibrium model for multiphase 15 multicomponent inorganic aerosols, *Aquat. Geochem.*, 4(123), doi:10.1023/A:1009604003981, 1998.
- Pappin, A. J., and Hakami, A.: Attainment vs exposure: ozone metric responses to source-specific NO_x controls using adjoint sensitivity analysis, *Environ. Sci. Technol.*, 47(23), 13519-13527, doi:10.1021/es4024145, 2013.
- Pappin, A. J., Mesbah, S. M., Hakami, A. and Schott, S.: Diminishing returns or compounding benefits of air pollution control? The case of NO_x and ozone, *Environ. Sci. Technol.*, 49(16), 9548-9556, doi:10.1021/acs.est.5b00950, 2015.
- 20 Pappin, A. J., Hakami, A., Blagden, P., Nasari, M., Szyszkowicz, M., and Burnett, R.T.: Health benefits of reducing NO_x emissions in the presence of epidemiological and atmospheric nonlinearities, *Environ. Res. Lett.*, 11(6), 064015, doi:10.1088/1748-9326/11/6/064015, 2016.
- Park, S.-Y., Kim, D.-H., Lee, S.-H., and Lee, H. W.: Variational data assimilation for the optimized ozone initial state and the short-time forecasting, *Atmos. Chem. Phys.*, 16(5), 3631-3649, doi:10.5194/acp-16-3631-2016, 2016.
- 25 Pascanu, R., Mikolov, T., Bengio, Y.: On the difficulty of training Recurrent Neural Networks, arXiv:1211.5063, 2013.
- Pinault, L. L., Weichenthal, S., Crouse, D. L., Brauer, M., Erickson, A., van Donkelaar, A., Martin, R. V., Hystad, P., Chen, H., Finès, P. and Brook, J. R.: Associations between fine particulate matter and mortality in the 2001 Canadian Census Health and Environment Cohort, *Environ. Res.*, 159, 406-415, doi: 10.1016/j.envres.2017.08.037, 2017.
- Pleim, J. E., and Chang, J. S.: A non-local closure model for vertical mixing in the convective boundary layer, *Atmos. Environ.*, 30 26A(6), 965-981, doi:10.1016/0960-1686(92)90028-J, 1992.
- Qi, L., Li, Q., Henze, D. K., Tseng, H.-L. and He, C.: Sources of springtime surface black carbon in the Arctic: an adjoint analysis for April 2008, *Atmos. Chem. Phys.*, 17(15), 9697, doi:10.5194/acp-17-9697-2017, 2017.
- Resler, J., Eben, K., Jurus, P., and Liczki, J.: Inverse modeling of emissions and their time profiles, *Atmos. Pollut. Res.*, 1(4), 288-295, doi:10.5094/APR.2010.036, 2010.

- Sandu, A., Daescu, D. N., Carmichael, G. R., and Chai, T.: Adjoint sensitivity analysis of regional air quality models, *J. Comput. Phys.*, 204(1), 222-252, doi:10.1016/j.jcp.2004.10.011, 2005.
- Seinfeld, J. H., and Pandis, S. N.: Atmospheric chemistry and physics: from air pollution to climate change, John Wiley, Hoboken, NJ, 2006.
- 5 Sirkes, Z., and Tziperman, E.: Finite difference of adjoint or adjoint of finite difference?, *Mon. Weather Rev.*, 125(12), 3373-3378, doi:10.1175/1520-0493(1997)125<3373:FDOAOA>2.0.CO;2, 1997.
- Squire, W., and Trapp, G.: Using complex variables to estimate derivatives of real functions, *SIAM Rev.*, 40(1), 110-112, doi:10.1137/S003614459631241X, 1998.
- Tai, A. P., Mickley, L. J., and Jacob, D. J.: Correlations between fine particulate matter (PM_{2.5}) and meteorological variables
10 in the United States: implications for the sensitivity of PM_{2.5} to climate change, *Atmos. Environ.*, 44(32), 3976-3984, doi:10.1016/j.atmosenv.2010.06.060, 2010.
- Tai, A. P., Mickley, L. J., Jacob, D. J., Leibensperger, E. M., Zhang, L., Fisher, J. A., and Pye, H. O. T.: Meteorological modes of variability for fine particulate matter (PM 2.5) air quality in the United States: implications for PM 2.5 sensitivity to climate change, *Atmos. Chem. Phys.*, 12(6), 3131-3145, doi:10.5194/acp-12-3131-2012, 2012.
- 15 Talagrand, O., and Courtier, P.: Variational assimilation of meteorological observations with the adjoint vorticity equation. I: Theory, *Q. J. R. Meteor. Soc.*, 113(478), 1311-1328, doi:10.1002/qj.49711347812, 1987.
- Thuburn, J., and Haine, T. W.: Adjoint of nonoscillatory advection schemes, *J. Comput. Phys.*, 171(2), 616-631, doi:10.1006/jcph.2001.6799, 2001.
- Turner, M. C., Jerrett, M., Pope III, C. A., Krewski, D., Gapstur, S. M., Diver, W. R., Beckerman, B. S., Marshall, J. D., Su,
20 J., Crouse, D. L., and Burnett, R. T.: Long-term ozone exposure and mortality in a large prospective study, *Am. J. Respir. Crit. Care Med.*, 193(10), 1134-1142, doi: 10.1164/rccm.201508-1633OC, 2016.
- Turner, M. D., Henze, D. K., Hakami, A., Zhao, S., Resler, J., Carmichael, G. R., Stanier, C. O., Baek, J., Sandu, A., Russell, A. G., Nenes, A.: Differences between magnitudes and health impacts of bc emissions across the united states using 12 km scale seasonal source apportionment, *Environ. Sci. Technol.*, 49(7), 4362-4371, doi: 10.1021/es505968b, 2015a.
- 25 Turner, M. D., Henze, D. K., Capps, S. L., Hakami, A., Zhao, S., Resler, J., Carmichael, G. R., Stanier, C. O., Baek, J., Sandu, A. and Russell, A. G.: Premature deaths attributed to source-specific BC emissions in six urban US regions, *Environ. Res. Lett.*, 10(11), 114014, doi:10.1088/1748-9326/10/11/114014, 2015b.
- Vukićević, T., Steyskal, M., and Hecht, M.: Properties of advection algorithms in the context of variational data assimilation, *Mon. Weather Rev.*, 129(5), 1221-1231, doi:10.1175/1520-0493(2001)129<1221:POAAIT>2.0.CO;2, 2001.
- 30 Wang, K. Y., Lary, D. J., Shallcross, D. E., Hall, S. M., and Pyle, J. A.: A review on the use of the adjoint method in four-dimensional atmospheric-chemistry data assimilation, *Q. J. R. Meteorol. Soc.*, 127(576), 2181-2204, doi:10.1002/qj.49712757616, 2001.
- Wang, Q., Moin, P., and Iaccarino, G.: Minimal repetition dynamic checkpointing algorithm for unsteady adjoint calculation. *SIAM J. Sci. Comput.*, 31(4), pp.2549-2567, doi: 10.1137/080727890, 2009.

West, J. J., Cohen, A., Dentener, F., Brunekreef, B., Zhu, T., Armstrong, B., Bell, M. L., Brauer, M., Carmichael, G., Costa, D. L., Dockery, D. W., Kleeman, M., Krzyzanowski, M., Künzli, N., Lioussé, C., Lung, S.-C. C., Martin, R. V., Pöschl, U., Pope III, C. A., Roberts, J. M., Russell, A. G., and Wiedinmyer C.: What we breathe impacts our health: improving understanding of the link between air pollution and health, *Environ. Sci Technol.*, 50, 4895-4904, doi:10.1021/acs.est.5b03827,

5 2016.

Yarwood, G., Rao, S., Yocke, M. A., and Whitten, G. Z.: Updates to the carbon bond chemical mechanism: CB05, Final report to the US EPA, RT-0400675, 2005.

Zhang, L., Liu, L., Zhao, Y., Gong, S., Zhang, X., Henze, D. K., Capps, S. L., Fu, T.-M., Zhang, Q., and Wang, Y.: Source attribution of particulate matter pollution over North China with the adjoint method, *Environ. Res. Lett.*, 10(8), 084011,

10 doi:10.1088/1748-9326/10/8/084011, 2015.

Zhao, S., Pappin, A. J., Mesbah, S. M., Zhang, J. Y. J., MacDonald, N. L., and Hakami, A.: Adjoint estimation of ozone climate penalties, *Geophys. Res. Lett.*, 40(20), 5559-5563, doi:10.1002/2013GL057623, 2013.

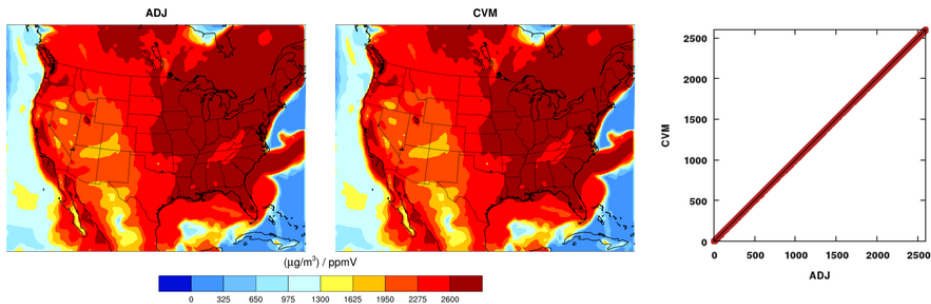
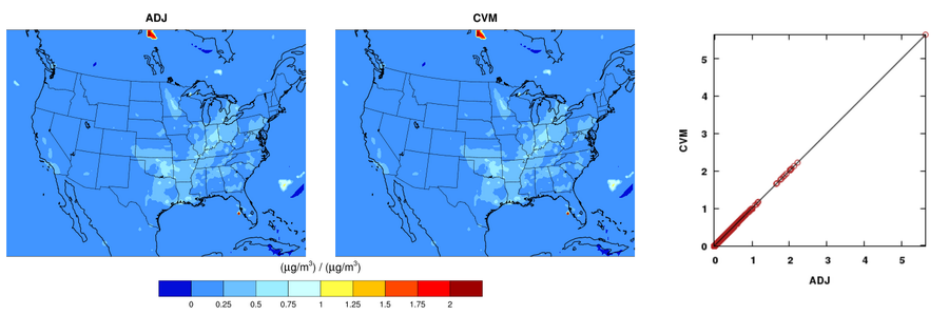


Figure 1: Evaluation of the SOA (Secondary Organic Aerosol) process with the ADJ (Adjoint) against the CVM (Complex Variable Method) sensitivities of the final concentrations of an accumulation-mode aerosol species AALKJ ($\mu\text{g}/\text{m}^3$) with respect to the initial concentrations of a semi-volatile species SV_ALK (ppmV) from a one-day test run. The perturbation size for the CVM is 1.E-24.



5
Figure 2: Evaluation of all aerosol sub-processes but thermodynamics with the ADJ (Adjoint) against the CVM (Complex Variable Method) sensitivities of the final concentrations of an accumulation-mode aerosol species ASO4J ($\mu\text{g}/\text{m}^3$) with respect to the initial concentrations of another Aitken-mode aerosol species ASO4I ($\mu\text{g}/\text{m}^3$) from a one-day test run. The perturbation size for the CVM is 1.E-12.




Figure 3: Evaluation of all aerosol sub-processes with the ADJ (Adjoint) against the CVM (Complex Variable Method) sensitivities of the final concentrations of an accumulation-mode aerosol species ASO4J ($\mu\text{g}/\text{m}^3$) with respect to the initial concentrations of the Aitken-mode aerosol species ASO4I ($\mu\text{g}/\text{m}^3$) from a one-day test run: left, original results; right, ISORROPIA branches set consistent between the CVM and the ADJ. The perturbation size for the CVM is 1.E-12.




Figure 4: Evaluation of all aerosol sub-processes with the ADJ (Adjoint) against the CVM (Complex Variable Method) sensitivities of the final concentrations of an accumulation-mode aerosol species ASO4J ($\mu\text{g}/\text{m}^3$) with respect to the initial concentrations of NH_3 (ppmV) from a one-day test run. The perturbation size is 1.E-12.




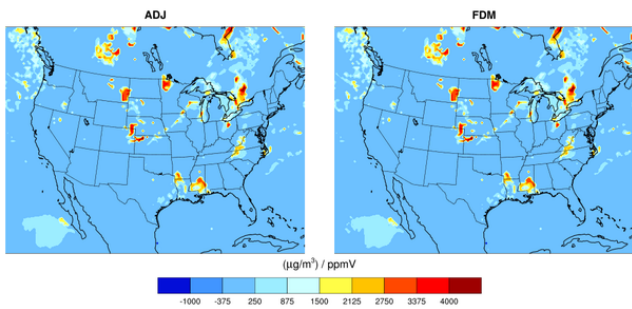



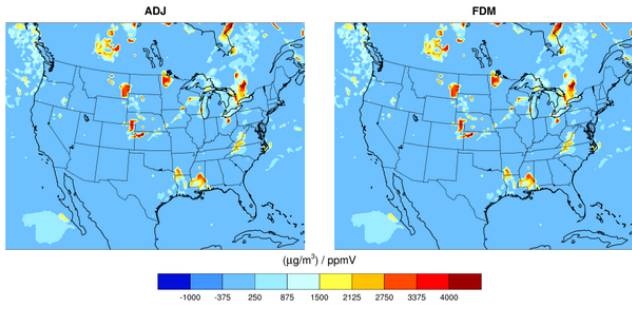


Figure 5: Evaluation of sub-grid cloud mixing with the ADJ (Adjoint) against the FDM (Finite Difference Method) sensitivities of the final concentrations of an accumulation-mode aerosol species ASO4J ($\mu\text{g}/\text{m}^3$) with respect to the initial ASO4J ($\mu\text{g}/\text{m}^3$) from a one-day test run. The perturbation size for the FDM is $0.01\mu\text{g}/\text{m}^3$.



5
Figure 6: Evaluation of cloud dynamics and aqueous chemistry with the ADJ (Adjoint) against the FDM (Finite Difference Method) sensitivities of the final concentrations of an accumulation-mode aerosol species ASO4J ($\mu\text{g}/\text{m}^3$) with respect to the initial ASO4J ($\mu\text{g}/\text{m}^3$) from a one-day test run. The perturbation size for the FDM is $0.01\mu\text{g}/\text{m}^3$.



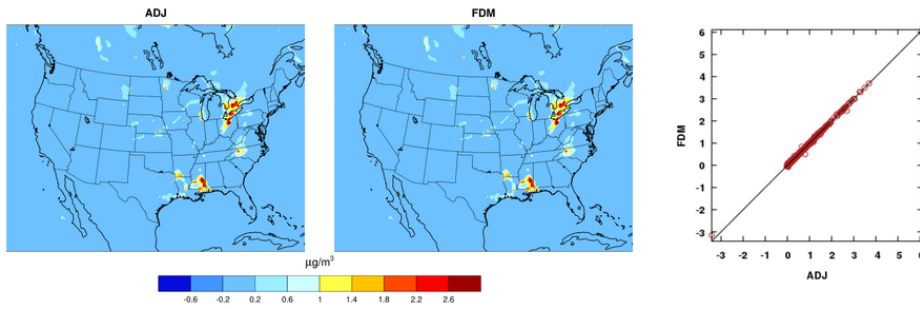




Figure 7: Evaluation of cloud dynamics and aqueous chemistry with the ADJ (Adjoint) against the FDM (Finite Difference Method) sensitivities of the final concentrations of an accumulation-mode aerosol species ASO4J ($\mu\text{g}/\text{m}^3$) with respect to the initial SO₂ (ppmV) from a one-day test run. The perturbation size for the FDM is 0.01/0.001/0.0001 ppb for the top three figures. [\(from top to bottom\)](#). The bottom figure is of sensitivities semi-normalized by the initial conditions of SO₂ and the perturbation size for the FDM is 0.001 ppb.

5



5 Figure 8: Evaluation of gas-phase chemistry with the ADJ (Adjoint) against the FDM (Finite Difference Method) sensitivities of the final O_3 concentrations (ppmV) with respect to the initial NO_2 concentrations (ppmV) from a one-day test run. The perturbation sizes are 0.1, 0.01 and 0.001 ppb for the FDM in the plots from the top to the bottom.



5 Figure 9: The Jacobians of absolute sensitivities from the ADJ of gas-phase chemistry at a grid cell from a one-step run. The X and Y axes represent all species involved in gas-phase chemistry.




Figure 10: Evaluation of gas-phase chemistry with the Jacobians of absolute sensitivities from the ADJ (Adjoint) and the CVM (Complex Variable Method) at a grid cell from a one-step run. A circle represents a pair of the absolute sensitivities from the ADJ (shown in Figure 9) and the CVM.

5






Figure 11: Evaluation of horizontal advection in the x direction with the DADJ (discrete adjoint) against the CVM (Complex Variable Method) sensitivities of the final concentrations of an accumulation-mode aerosol species ASO4J ($\mu\text{g}/\text{m}^3$) with respect to the initial ASO4J ($\mu\text{g}/\text{m}^3$) from a one-day test run. For the tilc plots, the x - and y -axes represent the horizontal y direction and the vertical layers, respectively. The perturbation size for the CVM is 1.E-12.

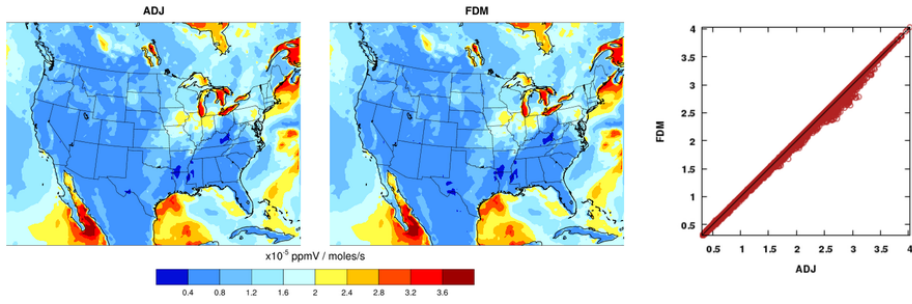
Deleted: op

10



5 **Figure 12:** Evaluation of advection with the CADJ (continuous adjoint) against the CVM (Complex Variable Method) sensitivities of the final concentrations of an accumulation-mode aerosol species ASO4J ($\mu\text{g}/\text{m}^3$) with respect to the initial ASO4J ($\mu\text{g}/\text{m}^3$) from a one-day test run. The perturbation size for the CVM is $1.E-12$.





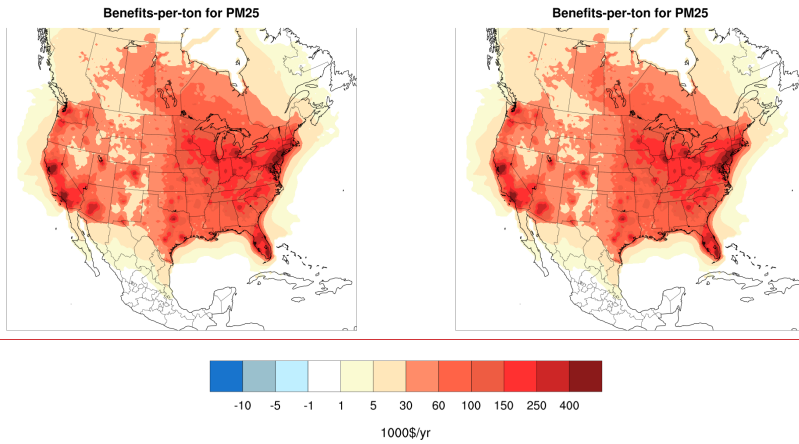
(b)

Figure 13: Evaluation of vertical diffusion with the ADJ (Adjoint) against the FDM (Finite Difference Method) from a one-day test run: (a) sensitivities of the final ozone (ppmV) with respect to the initial ozone (ppmV) with a 1-ppb perturbation size for the FDM; (b) sensitivities of the final nitrogen dioxide (ppmV) with respect to the initial nitrogen dioxide emissions (moles/s) with a perturbation size of 0.1 moles/s. The values in (b) are multiplied by 10^5 .


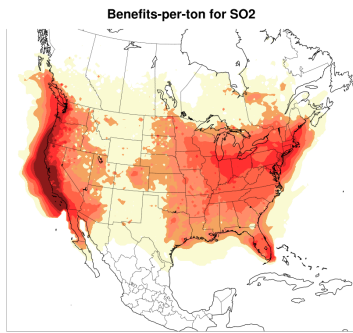
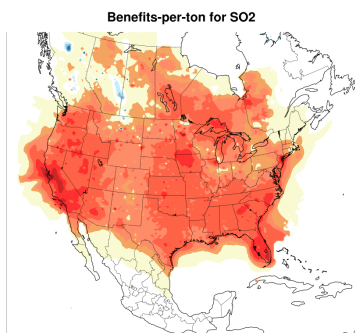

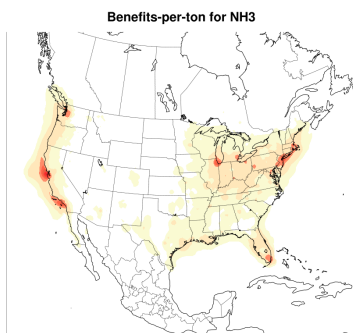
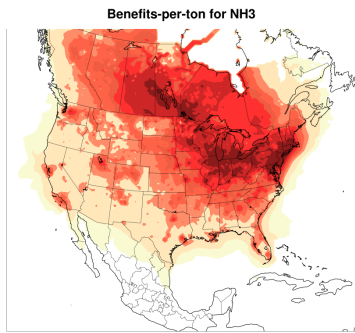
Formatted: Line spacing: single

WINTER

SUMMER



10



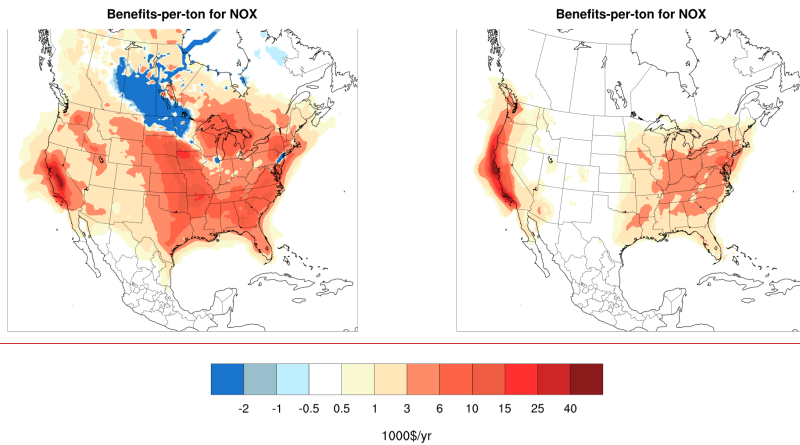
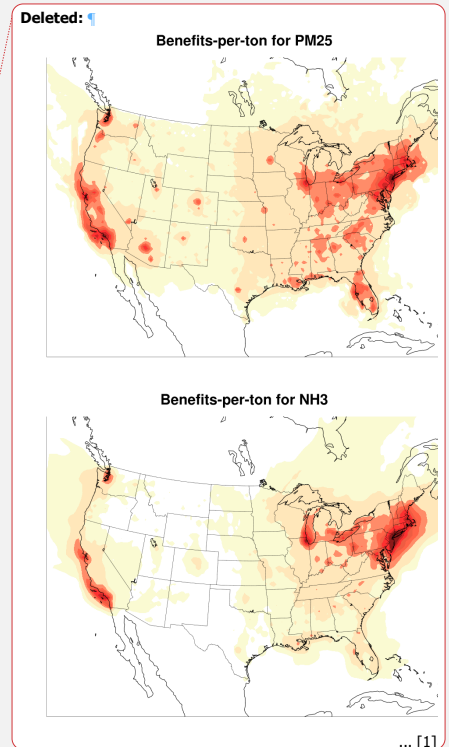


Figure 14: Application of the adjoint model for adjoint sensitivity analysis to estimate the benefits-per-ton (BPTs) related to long-term $PM_{2.5}$ exposure for primary $PM_{2.5}$ emissions in the US and its precursors NH_3 , SO_2 , and NO_x . The BPT is time-integrated and location-specific, i.e., each value in the figures represents the BPT for the specific emissions at the specific location. For example, a value of \$30,000/ton for SO_2 emissions at a location suggests that reducing emission of SO_2 at that location entails \$30,000 in valuated benefits across the U.S. The left and right columns show BPTs from three-month simulations for winter and summer, respectively.

Cells	#1	#2	#3	#4	#5	#6
CADJ	0.63	5.63	0.65	0.51	0.40	0.92
DADJ	0.45	5.86	0.47	0.93	0.71	0.88
CVM	0.42	5.96	0.47	0.95	0.77	0.85
FDM	0.17	6.70	21.21	-8.87	-23.20	0.54

Table 1: Evaluation of the full adjoint model with the ADJ (Adjoint) against the CVM (Complex Variable Method) sensitivities of the concentrations of an accumulation-mode aerosol species $ASO4J$ ($\mu g/m^3$) at Hour 24 with respect to the concentrations of a gas species SO_2 (ppmV) at Hour 23. The cells are arbitrarily picked. The perturbation size for the CVM is $1.E-12$ and the one for the FDM 10%. The relation of FDM and CVM sensitivities with CADJ and DADJ results has been discussed in Section 3.3.



<u>Process</u>	<u>Aerosols</u>	<u>Clouds</u>	<u>Chemistry</u>	<u>Vertical Diffusion</u>	<u>Horizontal Advection</u>	<u>Vertical Advection</u>
<u>File Size (GB)</u>	<u>24</u>	<u>24</u>	<u>17</u>	<u>1</u>	<u>0.2</u>	<u>0.2</u>

Table 2: Sizes of checkpoint files for the science processes in CMAQ for a single day. The computational domain has 148 columns, 112 rows, and 24 vertical layers. The synchronization time step of CMAQ is 12 minutes. Shown for horizontal and vertical advectons are checkpointing file sizes from the continuous version of adjoint.

5