Reply to Reviewer #1

We thank the reviewer for his/her comments, which have helped to improve the paper and clarify the details of our model. Our point-by-point response is given below.

RC – Reviewer comments
AC – Author's comments

1 Overview

RC: The manuscript by Wilson et al., "Development of a variational flux inversion system (INVICAT v1.0) within the TOMCAT chemical transport model" presents a new 4D-Var tool for inverse modeling of atmospheric tracers, such as CO2 or CH4. The authors present a clear overview of the utilities of this tool, followed by details of how it was constructed, and concluding with validation tests of the adjoint model gradients as well as a simple toy inversion. Overall the manuscript is clear and well written. My only overall suggestion is that in some places the text can be thinned a bit. Below are some specific suggestions to help clarify this and a few other aspects. This manuscript will be suitable for publication after minor revision.

AC: Thank you for these comments. We have thinned the text as suggested. Details are in the specific comments below.

2 Specific comments

RC: The authors might consider a different / revised title. Currently it implies that INVICAT is contained within TOMCAT, but rather the TOMCAT model would be contained within INVICAT, which comprises TOMCAT, ATOMCAT, and M1QN3.

AC: Correct. We have changed the title of the paper to *"Development of a variational flux inversion system (INVICAT v1.0) using the TOMCAT chemical transport model".*

RC: 7120.20: similar that → similar to that

AC: Done.

RC: 7120.28: possible → is reasonable? Because it's always possible to change them quite a bit, it just might not be reasonable.

AC: We agree, and have changed the sentence accordingly.

RC: 7124.22: modelling explained → modelling are explained.

AC: Done.

RC: 7125:8: The reference here seems quite outdated. There are any number of more recent papers analyzing the performance of L-BFGS in the context of modern 4DVar problems; I would suggest including something along these lines.

AC: The reference here has been changed to: (Gilbert and Lemarechal (1989); Nocedal and Wright (2006)), with some expansion in detail about the minimisation program used.

RC: 7125.21: At this point in the manuscript I wonder what the difference is between M' and M, both of which seem to be the linearized forward model. Is it possible to unify the notation?

AC: At this particular point in the manuscript, we are discussing only the theory of inverse modelling, without reference to TOMCAT in particular. Therefore M and M' are representing the full model and the tangent linear model, respectively. However, you are correct that in practice, these two models are identical for TOMCAT. To clarify this, we have added a sentence earlier in Section 3 explaining that the matrix **M** represents the linearised forward model, which $\mathbf{M}^T$ is produced from. We also explain in this section, rather than later on in the paper, that in TOMCAT's case, the forward model and linearized model are identical.

RC: The authors discuss checkpointing a bit, but typically this is only necessary for nonlinear processes. Tracer transport is, on paper, linear. So can they explain what is necessary to checkpoint? Is this because of the nonlinearities in the advection algorithm?

AC: This was indeed unclear and we have attempted to clarify these details. Firstly, we have now explicitly used the term 'checkpointing' in Section 3.1, before Equation (9), to describe the process of saving model data for use with the adjoint model. Secondly, we have explained in the first paragraph of Section 5 that since the forward model is linear, we do not need to 'checkpoint' forward model concentrations in order to run the adjoint model. However, we do save variables that are used the in forward model, such as convective flux rates, for use in the adjoint model as this is quicker than recalculating them.

RC: 7126.10: Strictly speaking, it's not correct to say that the adjoint calculates the sensitivity of c. Rather, it calculates the sensitivity of scalar metrics of c, such as J. Calculating the sensitivity of c would be the complete Jacobian, which is computationally prohibitive. Same issue regarding the text on line 24 of page 7120.

AC: We have changed both sentences in order to reflect this fact. The old line 7126.10 now reads "…propagates variables backwards through time in order to give the sensitivity of a scalar metric of **c** to model input parameters." The old line 7120.24 has been similarly changed.

RC: 7126.20: The authors justify the discrete approach for taking the adjoint of their advection scheme though Sirkes and Tziperman (1997). However, the work of Gou and Sandu (Atmospheric Env., 2011) provides a more relevant example, and, more importantly, they show that the discrete adjoint can actually lead to worse performance than the continuous adjoint for 4D-Var applications with CTMs.
Granted, the outcome is likely dependent upon the actual advection scheme employed. Since the authors don't appear to be using either the scheme from Sirkes and Tziperman (leap frog) or that studied in Gou and Sandu (piece-wise parabolic), they may need to think more carefully about the justification for their particular model.

AC: This is an interesting point, and this section has now been lengthened slightly in order to discuss the justification for use of the discrete adjoint within ATOMCAT. We have mentioned the results of Gou and Sandu (2011), and their relevance to our scheme. We feel that due to the linear nature of the advection scheme used in TOMCAT, and the fact that Gou and Sandu (2011) found little difference in performance between the two types of adjoint in 'real-world' inverse simulations, the use of the discrete adjoint is justified. However, further tests in the future (beyond the scope of this study) should be carried out in order to examine the performance of the continuous and discrete adjoints in comparison with each other. We have now included text clarifying this point in the paper.

RC: 7129.6: 2001) and → 2001), and

AC: Done.

RC: Regarding the Lagrange equality test, I have two concerns. First, this check can potentially overlook fortuitous cancelation of errors if the variables being tested are vectors, as is implied by the notation in the manuscript. A more stringent, albeit costly, test is to check individual variables. Second, in my experience a significant challenge in developing an adjoint is not the creation of adjoints of individual subroutines, but rather the "glueing" back all of the individual adjoint subroutines together in correct way. Thus I'm a bit concerned that the authors only test their subroutines for a single iteration (7134.5). Can they test over longer periods? Can they compare individual sensitivities to brute force (finite difference) sensitivities?

AC: First, we feel that the fact that the Lagrange equality test is just one of a number of tests that we have used in order to assess the accuracy of ATOMCAT means that using it in its current form is justified.

Second, we have extended number of iterations carried out during the equality test for the full model (rather than for individual subroutines), from one iteration (one hour) to 480 iterations (20 days). ATOMCAT still passes this test with an extremely high level of accuracy (equal to machine zero).

Third, we have included a new finite difference test for the adjoint and described the results in the paper. Again ATOMCAT passes this test of its accuracy.

RC: The reciprocity test is a nice way of checking the transport adjoint validity.

AC: Thank you!

RC: The inverse modeling test is perhaps overly simple, mainly in that it doesn't appear that any random error was ascribed to the observations. Most of the issues related to inverse modeling won't become apparent until noise is included in the observed data, so the results shown here will not likely be representative of the performance achieved in any real application.

AC: We appreciate this fact, and have therefore included a sentence describing this issue within the text. This test was not intended to be fully representative of a real application of the inverse model, but an attempt to examine the accuracy of the adjoint model and minimisation program.

RC: 7120.19: It might be good to explain what the differences are here.

AC: We have added a sentence here to clarify the differences between NWP weather prediction and our methodology.

RC: 7121.5-9: This is a bit strongly worded, suggest rephrase. I recognize that the availability of satellite data has driven the method to advance a lot in atmospheric chemistry the past decade, but the adjoint-based variational inverse method itself has been used decades earlier (e.g., nuclear reactor design in the 1940's) in other fields.

AC: We recognise that the wording here was fairly strong, and have rewritten the sentence to correct this as suggested.

RC: The material starting on the last line of 7123 through line 9 of 7124 feels redundant and unnecessary. The meaning of the cost function has already been explained. The utility of defining the cost function terms separately is not apparent.

AC: This section has been removed.

RC: 7124.25 – 7125.9: Since the details of the optimization are presented later, it doesn't seem necessary to describe the process qualitatively here.

AC: This section has been significantly shortened, although a brief description of the minimisation process remains here in order to provide context/reasoning for the development of the adjoint model.

RC: 7126.23: accurate in comparison → consistent with

AC: Done

RC: 7133.3: The sentence "These tests . . . " seems somewhat expendable.

AC: We have removed the first half of this sentence, but we have kept the text explaining the consequences of using an inaccurate adjoint model, which is not discussed elsewhere in the paper.

RC: 7136.12: If I'm not mistaken, this uses the BFGS algorithm, so that should be mentioned specifically with appropriate references.

AC: This has been added.

RC: 7139.16: Presenting this as an equation with no equality is a bit odd.

AC: Equality to RMSE_x added into the equation.

RC: Appendix A: given the vast body of literature using adjoint models for gradient based optimization, it doesn't seem that this section is necessary.

AC: Agreed. Appendix A has been removed.