

Profiling the robustness, efficiency and limits of the forward-adjoint method for 3D mantle convection modelling

M.G. Price¹ and J.H. Davies¹

¹ School of Earth and Ocean Sciences, Cardiff University, Cardiff CF10 3AT UK. Email: PriceMG@cardiff.ac.uk

8 November 2017

SUMMARY

Knowledge of Earth's past mantle structure is inherently unknown. This lack of knowledge presents problems in many areas of Earth science, including in mantle circulation modelling (MCM). As a mathematical model of mantle convection, MCM's require boundary and initial conditions. While boundary conditions are readily available from sources such as plate reconstructions for the upper surface, and as free slip at the core-mantle boundary (CMB), the initial condition is not known. MCM's have historically 'created' an initial condition using long 'spin up' processes using the oldest available plate reconstruction period available. Whilst these do yield good results when models are run to present day, it is difficult to infer with confidence results from early in a model's history. Techniques to overcome this problem are now being studied in geodynamics, such as by assimilating the known internal structure (e.g. from seismic tomography) of Earth at present day backwards in time. One such method is to use an iterative process known as the forward-adjoint method, which, while an efficient means of solving this inverse problem still strains all but the most cutting edge computational systems. In this study we endeavour to profile the effectiveness of this method using synthetic test cases as our known data source. We conclude that savings in terms of computational expense for forward-adjoint models can be achieved by streamlining the time-stepping of the calculation, as well as determining the most efficient method of updating initial conditions in the iterative scheme. Furthermore we observe that in the models presented, there exists an upper limit on the time interval over which solutions will practically converge, although this limit is likely to be linked to Rayleigh number.

Key words: Mantle processes, Inverse theory, Structure of the Earth, Numerical modelling.

1 INTRODUCTION

The adjoint method has been shown to be an excellent method of assimilating information in order to minimise the misfit between observed and predicted fields in a number of numerical modelling disciplines from meteorology (Courtier & Talagrand 1987) to seismology (Tarantola 1984) to oceanography (Menemenlis & Wunsch 1997). It has also been shown that this technique can be extended to address the problem of having an incomplete picture of Earth's historical mantle structure (Bunge et al. 2003).

The adjoint method for mantle circulation models (MCMs) has previously been utilised using both synthetic (Bunge et al. 2003; Ghelichkhan & Bunge 2016) and tomography derived (Horbach et al. 2014) true state forward-adjoint mantle models. Further studies using a simplified adjoint calculation have also been conducted (Ismail-Zadeh et al. 2004; Liu et al. 2008; Spasojevic et al. 2009), which use a subset of the full adjoint equations. In all of these studies, the interval over which the forward-adjoint iterations were ran spanned from 40 to 100 Ma until present day. A possible reason for shorter timescale calculations could be due to hardware limits, as a high resolution forward-adjoint calculation requires in

excess of 10 TB of hard disk storage for a single iteration running over a 40 Myr time interval, and takes several days to weeks of runtime (Horbach et al. 2014).

For an adjoint system using the full set of adjoint equations very little literature exists defining how reliable results are further back in time. Vynnytska & Bunge (2014) attempt to quantify this question of convergence using a two-dimensional forward adjoint model which did use the full system of adjoint equations. In this study the authors determined that knowledge of the surface boundary condition is crucial in order for the solution to converge to a unique initial condition, a result that they attribute to the uniqueness theorem of Serrin (1959). This states that two incompressible Stokes flows are equivalent given they have the same initial and boundary conditions. Their models investigated adjoint models running up to one-third of the transit time of mantle material, where they observed consistent convergence for various time intervals up to this limit. Here we will look to expand on the method of Vynnytska & Bunge (2014), by running a range of adjoint calculations and comparing them to a synthetic known mantle state.

Equally, the forward-adjoint formulation has a number of parameters that have not been tested for their influence on the conver-

gence of the final solution. One source that has been investigated is the effect that the choice in first guess initial condition has on the solution convergence. [Horbach et al. \(2014\)](#) found that regardless of this first guess, from a one-dimensional temperature profile to unlikely temperature structures derived from tomography, all forward-adjoint models converge to the same global minimum.

Very low resolution investigations have also been made with a compressible forward-adjoint model by [Ghelichkhan & Bunge \(2016\)](#). In their study they compare a compressible and incompressible adjoint model to a compressible ‘true’ state reference. They found that due to the formulation of the compressible adjoint equations, there are only small differences between the incompressible and compressible adjoint solutions to the reference temperature field.

With the results of [Vynnytska & Bunge \(2014\)](#); [Horbach et al. \(2014\)](#) and [Ghelichkhan & Bunge \(2016\)](#) already found, we look to extend the investigation of the forward-adjoint method to cover more of the parameters that may control the convergence of solutions. In doing so we hope to understand better how the method can be altered to yield more efficient compute times and storage requirements without compromising the convergence.

We organise this paper as follows: firstly we present the method used to solve the forward-adjoint system, including the equations, solution scheme and model setup. Then by computing synthetic initial and final states from a mantle circulation model to use as a benchmark for our forward-adjoint models, we investigate the effects on convergence of altering components of the forward-adjoint model. We will first investigate the role of the forward time-step lengths on convergence; following this we will look at altering δ , the fraction used to update our initial state; finally with the availability of plate reconstructions that extend over 200 Myr we look at varying the time period over which the calculation is run to provide an upper bound on the time interval that can be used with this method.

2 METHOD

2.1 Numerical methodology

The mantle circulation models used in this study are carried out using an adapted version of the three-dimensional finite element code TERRA ([Baumgardner 1985](#); [Bunge et al. 1997](#)), which solves the equations governing conservation of mass, momentum and energy (eqs. (1)) for an incompressible fluid

$$\nabla \cdot \vec{u} = 0 \quad (1a)$$

$$\nabla \cdot \left\{ \eta (\nabla \vec{u} + (\nabla \vec{u})^T) \right\} - \nabla P + \alpha \rho (T_{av} - T) \vec{g} = 0 \quad (1b)$$

$$\frac{\partial T}{\partial t} + \vec{u} \cdot \nabla T - \kappa \nabla^2 T - H = 0, \quad (1c)$$

and the adjoint equations governing conservation of mass, momentum and energy (eqs. (2)) for an incompressible fluid

$$\nabla \cdot \vec{\phi} = 0 \quad (2a)$$

$$\nabla \cdot \left(\eta \{ \nabla \vec{\phi} + (\nabla \vec{\phi})^T \} \right) - \nabla \psi + \tau \nabla T = 0 \quad (2b)$$

$$\frac{\partial \tau}{\partial t} + \vec{u} \cdot \nabla \tau + \kappa \nabla^2 \tau + \alpha \rho \vec{g} \cdot \vec{\phi} = \partial_T \chi(T). \quad (2c)$$

Here η is the dynamic viscosity, α is the coefficient of thermal expansion, ρ the reference density, T_{av} the radial average temperature profile, \vec{g} gravitational acceleration, κ is the thermal diffusivity and H is radiogenic heat production. $\partial_T \chi(T) = (T_i - T_m) \delta(t - t_1)$

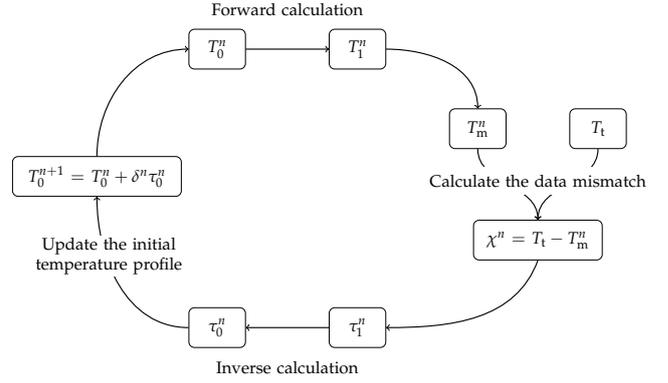


Figure 1. The forward-adjoint workflow

is the gradient of the misfit function χ that relates the model temperature T_m to the observed true temperature T_t . $\vec{\phi}$, ψ and τ are the adjoint complements to the forward terms \vec{u} , P and T .

These two sets of equations (eqs. (1) and (2)) are run together in an iterative loop with the temperature profile at t_0 updated using a conjugate gradient method ([Fletcher & Reeves 1964](#)),

$$T_0^{n+1} = T_0^n + \delta^n \tau_0^n. \quad (3)$$

Here δ is a damping factor (with $\delta \leq 1$), the value of which can be varied over the calculation. Figure 1 outlines the full iteration loop, and is cycled over until the misfit at t_1 falls below a pre determined level, or the total number of iterations is exceeded.

The convention for indexing of variables (\cdot_m^n) in this paper is as follows; superscripts denote the iteration number of the forward-adjoint calculation with $n \in \mathbb{N}_1$ and N denoting the final, optimal iteration; subscripts meanwhile denote the initial $m = 0$ or final $m = 1$ point in time of the model relative to the forward calculation. That is to say $m = 0$ refers to some point in the past from present day, regardless of whether discussing forward or adjoint variables. In addition to this, subscripts may also take the letter m or t ; where m refers to the *model* data source obtained from a forward calculation; and t refers to the external *true* data source which is being used as the target field.

The overall similarity between the forward (eqs. (1)) and adjoint (eqs. (2)) equations means that implementing the adjoint equations in a mantle convection code is fairly straightforward. Of course the presence of terms from the solution of the forward equations in the adjoint equations means that the two systems are intimately linked. This means that we are required to run a forward calculation, saving the velocity and temperature fields as we go for use in the following adjoint calculation. As this method looks to minimise the errors on the temperature field we look to perform successive forward-adjoint calculations to achieve this.

2.2 Model configuration

Similar to other studies investigating the forward-adjoint model we limit ourselves to a minimal number of assumptions in our setup with incompressible models, viscosity varying only radially and plate reconstructions used as a boundary constraint.

All test cases were modelled as an incompressible fluid on a mesh with 10 million finite elements, which corresponds to an average grid spacing of 50 km. The viscosity profile employed in our models varies as a function of depth (r) only according to the

Table 1. Model parameters

Parameter		Value	Units
Internal heating rate	H	4×10^{-12}	W kg^{-1}
Reference viscosity	η_0	3×10^{22}	Pa s
Density	ρ	4500	Kg m^{-3}
Thermal expansivity	α	2.5×10^{-5}	K^{-1}
Thermal conductivity	k	4	$\text{W m}^{-1} \text{K}^{-1}$
Thermal diffusivity	κ	1×10^{-6}	$\text{m}^2 \text{s}^{-1}$
Specific heat capacity	C_V	1000	$\text{J kg}^{-1} \text{K}^{-1}$
Basally heated Ra number	Ra_b	$\approx 2.4 \times 10^6$	–
Internally heated Ra number	Ra_H	$\approx 3.3 \times 10^7$	–

function

$$\eta(r) = \begin{cases} 100\eta_0 & r \leq 100 \text{ km} \\ \eta_0 & 100 \text{ km} < r \leq 660 \text{ km} \\ 30\eta_0 & r > 660 \text{ km} . \end{cases}$$

This provides our models with a stiff surface layer, to mimic the lithosphere, a much less viscous upper mantle, and a more viscous lower mantle. Such radial profiles as this are typically used as first order approximations in MCM's (e.g. Bunge et al. 2002; Bower et al. 2013; Nerlich et al. 2016). Other key model parameters are outlined in table 1.

The boundary conditions for the forward model are as follows. At the CMB we define the boundary on the velocity field to be free slip (due to the low viscosity of the liquid outer core). Meanwhile for the surface velocity boundary, to mimic the movements of the plates, we prescribe velocities using the plate motion reconstruction of Seton et al. (2012), with the velocities scaled to match the convective vigour of the underlying convection model (Bunge et al. 2002). The temperature field boundary conditions are set at 300 and 3000 K for the surface and CMB respectively. These surface velocities are assimilated in 1 Myr intervals over the course of our forward model calculations which ensures a smooth transition from one plate stage to the next. Due to the assumptions taken in deriving this set of adjoint equations (see Horbach et al. 2014, for a full derivation), all the boundary conditions in the adjoint calculation are assumed to be perfect (i.e. known with no error), hence their values are zero. For all our test cases, T_0^1 our forward-adjoint models starting initial condition is a simple 1D temperature structure .

The Rayleigh number, the non-dimensional value which quantifies the convective vigour of the system that is basally heated is defined as

$$\text{Ra}_b = \frac{\alpha \rho \Delta T D^3 g}{\kappa \eta} ,$$

where D is the mantle radius and ΔT the total temperature contrast across the depth of the mantle. As our models also contain internal heating we define a second Rayleigh number for an internally heated volume which is defined as

$$\text{Ra}_H = \frac{\alpha \rho^2 H D^5 g}{k \kappa \eta} .$$

Using the values used in this study this gives our models a Rayleigh number of $\text{Ra}_b \approx 2.4 \times 10^6$ and $\text{Ra}_H \approx 3.3 \times 10^7$, meaning we expect the vigour of convection in our models to be slightly below what is expected for Earth ($\text{Ra}_H \approx 10^8$). Because of this we increase the timescales of our models to allow a similar amount of

convection to occur in the models presented here. All times quoted in this study are after accounting for this scaling.

For our study into the limits, robustness and efficiencies of the forward-adjoint method we use a synthetically generated temperature field for our true data source, similar to the method used in Bunge et al. (2003). To generate this synthetic data source we begin with a mantle convection calculation (free-slip as the surface boundary condition) and run this until it has reached a quasi-steady thermal state, such that the surface heat flux is fairly constant. Once this state has been reached the convecting mantle volume is conditioned with the oldest available plate motion stage (200 Ma for the Seton et al. (2012) model) to introduce the tectonic features of the stage. This is achieved by continuing the calculation for ≈ 50 Myr with the oldest plate stage, introducing the temperature and density anomalies in the upper to mid mantle reflective of this plate stage. Following this setup we then run our reference model through to present day, recording the mantle structure at 10 Myr intervals. With this information we can then also compare the predicted initial stage mantle structure from our adjoint calculation to the evolving mantle structure in the synthetic case.

The main metric we use to compare different cases is the L^2 norm of the volume average of the mismatch in temperature T at t_0 and t_1 . This gives us a measure of the distance between the modelled T_m , and true T_t , temperature fields, with lower values reflecting a smaller distance between the two fields.

We conclude the methodology by outlining the motivation and setup for each each of the different cases we considered in this work.

2.3 Altering the forward time-step mechanism

Our first investigation concerns the time-stepping mechanism of the forward calculation. We see from eqs. (2) that the adjoint equations require knowledge of the forward iterations temperature and velocity values. This coupling of the two equations results in the adjoint calculation having to take the exact same time-steps back in time as the freely evolving forward calculations time-steps. The storage of the \vec{u} and T arrays at each time-step requires, depending on resolution over 3 GB of temporary storage for each time-step in high resolution models (grid spacing ≤ 25 km). As calculations grow into the thousands of time-steps, this can begin to push practical limits of storage available on high performance computing (HPC) systems. To this end, we hope that by constraining the forward time-steps evolution we can achieve practical gains in terms of storage space and runtime without compromising the convergence of our solutions.

In its current implementation in TERRA, time-stepping is performed using a second order Runge-Kutta scheme. The length of each time-step is adjusted as the calculation evolves, such that only one iteration of the code's multigrid algorithm is required to maintain the residual error in the solving of the momentum equation below a specified level. Furthermore, the length of each time-step is bounded from above in order to ensure the Courant-Friedrichs-Lewy (CFL) condition is not violated. The net result is a time-step which will increase $\times 1.1$ if all tolerances are reached or $\times 0.8$ if one or more tolerances are not met.

In order to investigate the effect that changing the time-step evolution has on the convergence of the forward-adjoint we shall conduct two synthetic tests which cover 10 Myr and 40 Myr of Earth history respectively. We consider three methods of time-step evolution:

• **Free** - This is the classic time-step as found in TERRA and described above (this is the default used elsewhere in the paper).

• **Upwards** - This limits the time-step by taking the maximum value out of the previous and current time-step values; this has the net effect of allowing the time-step to increase when the tolerances are met but never decrease in value.

• **Forced** - In this implementation the time-step is hard coded to the CFL limit apart from the first ten time-steps of each 1 Myr plate stage in order to allow the code to handle the change of plate boundaries.

2.4 Varying the choice of δ for the T_0^{n+1} update

At the end of a forward-adjoint iteration we update the initial condition guess using eq. (3), $T_0^{n+1} = T_0^n + \delta^n \tau^n$, where δ is the damping factor that controls the influence of the adjoint temperature in the update process. A well chosen function for determining δ could have a significant effect on the rate of convergence of our solutions, resulting in fewer forward-adjoint iterations required to reach satisfactory results. This of course would have the additional benefit of less total storage required, as well as shorter run times; both valuable when considering calculations run on large HPC systems.

In order to understand how the choice of δ affects the convergence of our forward-adjoint solution, we will conduct a series of synthetic forward-adjoint models running over a time period of both 10 Myr and 40 Myr. By running our tests using synthetic data we are able to directly compare the accuracy of our predicted initial condition for the different cases. For both of these time spans we consider the following for our choices for δ :

- $\delta = 1$ - A value where we apply the whole minimised adjoint field.
- $\delta = 0.8$ - A second fixed value addition of the adjoint field (this is the default used elsewhere in the paper).
- $\delta = 1 - 0.02n$ - A simple decreasing function where n corresponds to the iteration number.
- $\delta = \frac{1}{n + n_0}$ - A decreasing function for δ as used in Liu & Gurnis (2008), where we choose $n_0 = 0$. Again n corresponds to the iteration number.
- $\delta = \min \left\{ \frac{1}{1 + n}, \frac{\|\chi^n\|}{\|\tau_0^n\|} \right\}$ - Here δ is determined based on the gradient method as in Ismail-Zadeh et al. (2004). n once more corresponds to the iteration number, χ^n the misfit between the true and model temperatures for the current iteration, and τ_0^n the adjoint temperature at t_0 .

2.5 Reducing the interval of stored \vec{u} and T fields

Due to the Stokes flow nature of the equations governing mantle convection, the \vec{u} and T fields evolve slowly. Based on this assumption, there is the potential to output these fields not at each time-step but instead at some other interval. Then based on the assumption that the fields evolve sufficiently slowly, a simple interpolation between saved data can be used to determine the intermediate \vec{u} and T fields. The potential cost savings in terms of hard disk space for this technique would be significant

To investigate the feasibility of interpolating between saved information, we shall compare calculations run over time intervals of 10 Myr and 40 Myr which store either:

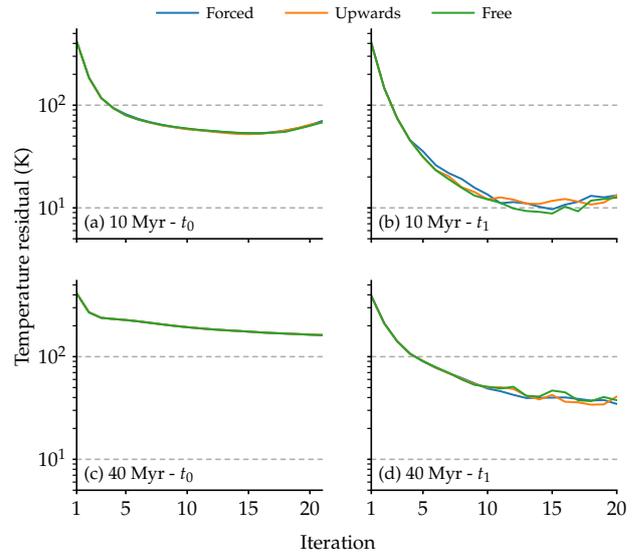


Figure 2. Logarithmic plots of the L^2 norms for the residual temperatures from 10 (top) and 40 (bottom) Myr synthetic forward-adjoint calculation with varying methods of time-step evolution. Residual temperatures at t_0 (left) and t_1 (right), are shown.

- \vec{u} and T fields at each time-step from the forward calculation (this is the default used elsewhere in the paper).
- \vec{u} and T fields only for every tenth time-step, with the interim time-step's fields determined by interpolation using our knowledge of the length of each forward time-step.

2.6 Varying the time interval of the forward-adjoint calculation

For our final investigation we shall look at how the time interval we run the forward-adjoint model over affects the residuals between the modelled and true temperature fields. As previously discussed, previous studies that looked at 3D forward-adjoint simulations ran over a time interval of 40 to 100 Myr. Both Liu et al. (2008); Bunge et al. (2003) ran global adjoint models that extended back to 100 Myr, with the latter using the full adjoint system of equations. Bunge et al. (2003) ran their calculation for 100 forward-adjoint iterations and saw a reduction of 90% at t_1 and 50% at t_0 . For our study we do not attempt to run for as many iterations due to the limits of the computational resource available for this work.

With the availability of plate reconstructions that extend back to the early Jurassic (Seton et al. 2012), we can therefore conduct forward-adjoint calculations that over a time period of twice that of other studies. In doing so we hope to determine at what point, if any, information is not significantly propagated back in time using an adjoint calculation. We will run several cases covering a range of intervals; 10, 20, 40, 80, 120, 160 and 200 Myr.

3 RESULTS

3.1 Altering the forward time-step mechanism for the forward adjoint calculation

Figure 2 show the results for the different methods of controlling the forward time-step for both the 10 Myr and 40 Myr calculations.

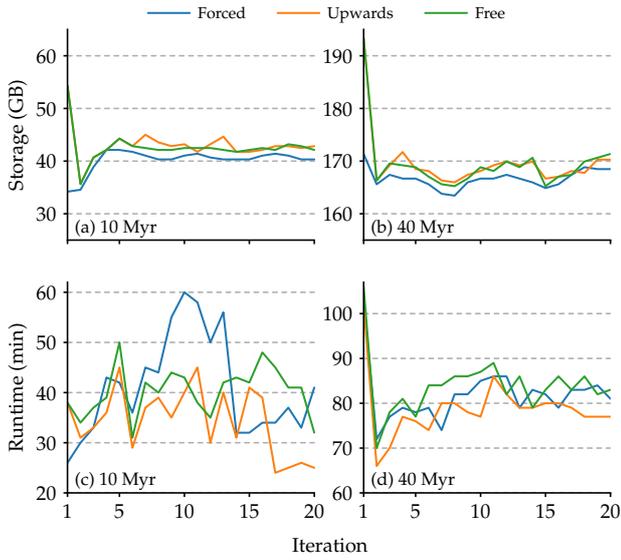


Figure 3. Graphs displaying the total hard disk storage required for saving \bar{u} and T from the forward calculation for each iteration for the (a) 10 Myr and (b) 40 Myr calculations. (c) and (d) depict the total computation time required for each iteration for the 10 and 40 Myr calculations respectively.

For the 10 Myr (Fig. 2a and b) calculation the t_0 profiles all display the same convergence to a minimum around 55 K globally before a small uplift in the norms is observed. At t_1 the choice of time-step method does have some influence on the convergence, however these differences are $\mathcal{O} \sim 1$ K.

For the 40 Myr calculation we again see that the convergence of the different solutions very closely match each other both at t_0 and at t_1 (Fig. 2c and d). Similar to the 10 Myr calculation we see in Fig. 2d that in the later iterations there is some variation in the L^2 norms $\mathcal{O} \sim 10$ K, although there is no noticeable favoured implementation. We also note that the minimum norm values at both end points are higher than those observed in the 10 Myr calculation. Finally we do not observe the small divergence of the L^2 norms in the later iterations at t_0 .

As the aim of testing these different time-step controls is to maximise savings in storage and runtime without compromising our solutions we also look at how these are affected by the time-step choices in Fig. 3.

At 0.36 GB per time-step at this resolution we see in Fig. 3a and b, that over the 10 and 40 Myr calculations we require a fairly uniform amount of storage for each iteration of roughly 40 and 170 GB respectively. In both length calculations, the initial iteration shows a much higher storage requirement for the first iteration for the free and upwards methods. We can interpret this as an effect of our choice of initial condition being a 1D temperature profile. As there is no information of the internal structure a freely evolving model will progress slowly as thermal structures are introduced into the model, which will limit the time-step growth as the model copes with these structures. The forced time-step method on the other hand ignores this and progresses at full pace, hence the smaller storage requirement in the first iteration. By the start of the second iteration some thermal structures have been introduced into all the test cases and so the calculation requires a similar amount of storage across all cases.

Turning to the runtime for our cases we find that the upwards time-step performs best on average over both time intervals (Fig. 3c

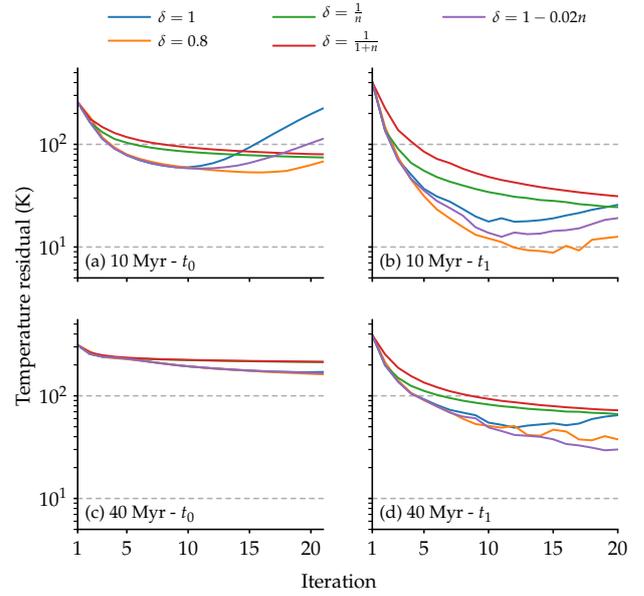


Figure 4. Logarithmic plots of the L^2 norms for the residual temperatures from 10 (top) and 40 (bottom) Myr synthetic forward-adjoint calculation for different δ implementations. Residual temperatures at t_0 (left) and t_1 (right), are shown.

and d). Perhaps surprisingly, the forced time-step does not outperform the others, and in the 10 Myr case we also see it performs noticeably worse for a period. We attribute this to the case ignoring any difficult to solve regions in the volume, therefore requiring extra multigrid cycles in order to satisfy the convergence tolerances of the code.

3.2 Varying the choice of δ for the T_0^{n+1} update

The temperature residuals for all cases for various choices in δ are presented in Fig. 4. Regardless of calculation time period, the choices in δ can be split into two separate groups based on their convergence.

The first group contains the linear choices of δ . Focusing on the 10 Myr results first, the three cases can be characterised by a very similar convergence rate at present day (Fig. 4b), with each reaching a minimum norm value ~ 10 -20 K for the misfit between the modelled and observed temperature field. While at t_1 there are some differences between the three choices of δ even early on, with $\delta = 0.8$ showing the best convergence, at t_0 all three cases show very similar convergence until iteration 10. Beyond this point $\delta = 1$ begins to quickly diverge, and our linear function shows similar, but less pronounced behaviour. All three cases obtain a minimum of ~ 50 -60 K. If we separate the calculated norms to cover only the upper and lower mantle respectively, we find that the divergence of our misfit at t_0 for the linear functions is only observed in the upper mantle. For the lower mantle norm, the linear functions are only marginally worse than the non-linear choices up until the last few iterations. Poorer upper mantle norms have been observed previously (Horbach et al. 2014), and we attribute the diverging norm values to a build up of over correction of features located at the near surface, introduced by these choices of δ . The difference between the upper and lower mantle is not observed at t_1 however, as over

the 10 Myr of forward model there is time to remove some of this over correction.

The second group contains the non-linear choices for δ . Both of these choices show a monotonically decreasing norm residual across both end points, a feature not shared by the first group. Whilst this may suggest a more stable convergence rate, we do see however that it is taking 20+ iterations for this group to begin falling below the minimum levels obtained by the first group. Across both the upper and lower mantle at t_1 they only just start to outperform the three choices from the linear group, namely $\delta = 1$. At t_0 we do not see the same divergence in the upper mantle due to over correction, but the minimum values after 20 iterations still fail to match the early norm levels of the linear δ group. We also see that $\delta = 1/n$ is consistently below the more complex non-linear function. From further investigation of this we found that the gradient determined by $\|\chi^n\|/\|\tau_0^n\|$ was larger than $\frac{1}{1+n}$ at all but the last two iterations. Because of this, the choice in δ was taken as the fraction $\frac{1}{1+n} < \frac{1}{n}$ and so explains the poorer convergence rate compared to the other non-linear choice.

For the 40 Myr calculations (Fig. 4c and d) the differences are less pronounced, but results still fall into the same two groups. We observe that at t_1 for four out of the five choices of δ , the temperature residuals continue to decrease across the 20 iterations obtaining a minimum value ranging from 30–70 K at the twentieth iteration. The exception is $\delta = 1$ which begins to show a diverging final stage misfit from the twelfth iteration onwards. The minimum value obtained by this choice is worse than the other linear functions at this point, showing no advantage to this choice in δ . In contrast to the 10 Myr calculations, the best observed choice at t_1 is $\delta = 1 - 0.02n$, which shows a much more stable convergence compared to the next nearest which shows an oscillating norm value. This oscillation for the 0.8 choice is actually from contributions in the upper mantle.

The picture at t_0 is much less diverse, with the values at each iteration for the differing choices of δ falling into the two previously mentioned groups of linear and non-linear (Fig. 4c). All the values show a monotonically decreasing residual, though the non-linear choices are decreasing by only a few K over the last ten iterations. The minimum values obtained range from 160–220 K.

As both the linear and non-linear choices in δ show different ideal properties (namely faster convergence, and monotonic convergence respectively), for our remaining investigations we shall continue to use both $\delta = 0.8$ and $\delta = \frac{1}{n}$ with hopes of finding further distinction between the two.

3.3 Reducing the interval of stored \vec{u} and T fields

The temperature residuals resulting in the forward-adjoint calculations where we reduce the number of saved \vec{u} and T fields are shown in Fig. 5.

Looking first at the present day residuals we see that compared to the non-interpolated calculations the 10 Myr (Fig. 5b) interpolated calculations both follow the same reduction in temperature residual. Over 40 Myr (Fig. 5d) however, whilst we again observe reducing residuals, the rate of decrease is much slower than when we save information at every time-step. Also at 40 Myr we see that the interpolated calculations, whilst still converging, actually perform worse than both non-interpolated cases.

The results at t_0 however, are more startling. Over the 10 Myr calculation (Fig. 5a), regardless of the choice in δ , after the first few iterations the interpolated calculations follow a very similar trend with their residual stagnating around 130 K. Looking at the

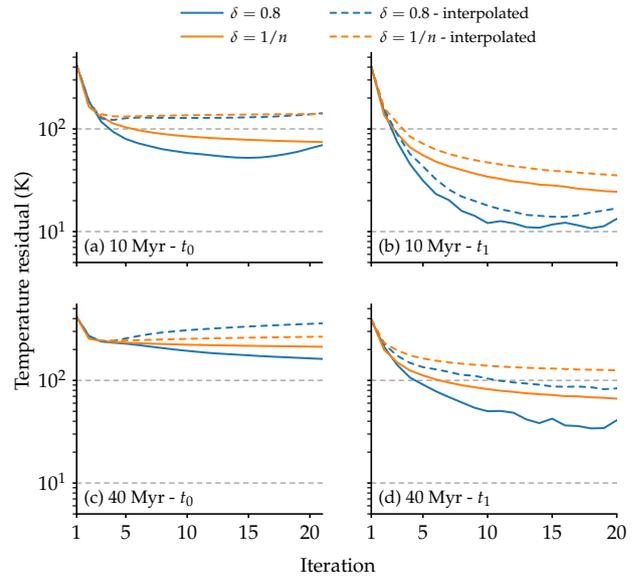


Figure 5. Logarithmic plots of the L^2 norms for the residual temperatures from 10 (top) and 40 (bottom) Myr synthetic forward-adjoint calculation where either the \vec{u} and T field is saved at every time-step from the forward calculation or where only every tenth time-step is saved and the intermediate time-steps are interpolated. Residual temperatures at t_0 (left) and t_1 (right), are shown.

40 Myr results in Fig. 5c, the $\delta = 0.8$ interpolated case performs significantly worse, showing a growing divergence in the residual. The other interpolated case also performs worse than the non-interpolated cases, with the residual stagnating around 330 K.

We attribute the distinct behaviour in the residuals at t_0 down to the missing information from the forward calculation which we have to interpolate to reconstruct the interim fields, which introduces errors into the final solution obtained at t_0 . In contrast, the residuals at t_1 are less obviously worse as the forward calculation is using the same boundary conditions regardless which limits the growth of these errors (Colli et al. 2015).

3.4 Varying the time interval of the forward-adjoint calculation

The temperature residuals at both t_0 and t_1 for these models are presented in Fig. 6 for $\delta = 0.8$ and $\delta = \frac{1}{n}$. For both choices in δ the results at t_1 are very similar with the trends we observe applicable to both cases, whereas there are pronounced differences at t_0 which we note.

The first feature we observe from the t_1 residual norms is that following the first forward calculation there are a range of values for the different length calculations. This is unsurprisingly arranged in order, with the shortest test case showing the greatest difference from the reference state, to the full 200 Myr model which has a residual norm nearly half that of the worst. We attribute this difference to the vastly different length assimilation times of our models, as having a long period of surface boundary conditions has previously been shown to provide a good match to present day observations (Bunge et al. 2002). Looking beyond the first iteration we see that by the end of the second forward iteration that the shorter length models are already displaying the best match to the observed mantle field, and by the fifth iteration the models have re-

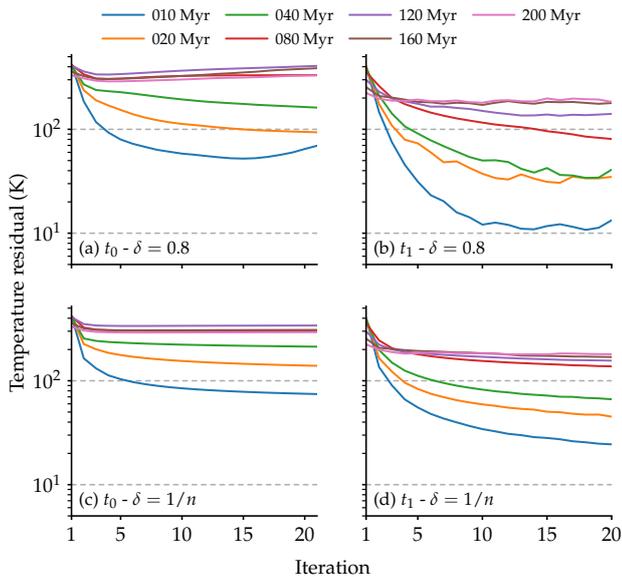


Figure 6. Logarithmic plots of the L^2 norms for the residual temperatures, at t_0 (left) and t_1 (right), of forward-adjoint models run over differing time intervals as a function of forward-adjoint iterations when $\delta = 0.8$ (top) and $\delta = \frac{1}{n}$ (bottom).

ordered completely in terms of best match. Over the course of the twenty iterations there are very small reductions in the residuals of the longest length calculations, whereas the short time scale models show much better matches.

It can be seen that for the linear δ models all the time intervals reach a similar or lower minimum compared to the $\frac{1}{n}$ cases. We also observe that for all the linear models ≤ 80 Myr there is continuing convergence while those ≥ 120 Myr appear to reach their minimum and oscillate about a final value. When $\delta = \frac{1}{n}$ all the cases apart from 200 Myr are monotonically converging at t_1 .

Looking at the temperature residuals observed at t_0 between the predicted initial condition and the reference we find a vastly contrasting picture between the choices in δ compared to at t_1 . First focussing on the similarities between the two choices in δ , we see at iteration 1 with our 1D radial temperature field that the longest calculations, 200 to 160 Myr, have a smaller misfit compared to the other cases. This could stem from the total length of the plate reconstruction model available to us together with how the reference model was calculated. We recall that the reference model is created by conditioning a plain mantle volume with the oldest plate stage for up to 50 Myr before the sequential assimilation of the full plate reconstruction. The reduced misfit values for the 200 to 160 Myr for T_0^1 suggests that the standard method of setup for a forward model, as used for the reference state, provides a mantle volume that is still fairly heterogeneous.

For both δ cases at t_0 , unlike at t_1 , the residuals are not arranged from shortest to longest assimilation periods. While 10, 20 and 40 Myr are arranged in order, the next best fitting model is the 200 and 80 Myr respectively with the 120 Myr calculation showing the worst misfit values across the whole calculation.

Looking at the differences between the choices in δ over the various time spans, for the linear choice all models > 40 Myr achieve their minimum misfit following three forward-adjoint iterations before their norms steadily diverge (Fig. 6a). In contrast, the effects of using $\delta = \frac{1}{n}$ observed previously continue (Fig. 6c), with

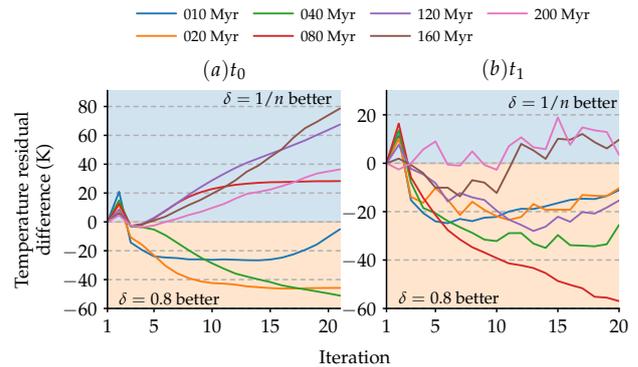


Figure 7. The difference between the whole mantle L^2 norms for the various length cases with $\delta = 0.8$ and $\delta = \frac{1}{n}$ at (a) t_0 and (b) t_1 . Positive values (shaded blue) denote where $\delta = \frac{1}{n}$ produces a closer match to the 'true' data, whilst negative values (shaded orange) are where $\delta = 0.8$ produces the closer match.

the residual values converging monotonically, albeit at a slower rate.

In order to directly compare the results from the different choices in δ , in Fig. 7 we show the difference at each iteration between the residuals for the different δ choices. At t_1 , $\delta = 0.8$ yields a better present day match for all but the longest length calculations. The reasoning behind this could be that, as more information is added to our starting 1D temperature structure initial condition at each iteration when $\delta = 0.8$, the shorter length forward calculations do not run for sufficiently long enough to introduce significant features into the lower mantle. The long length calculations however, (200 and 160 Myr in particular) run for long enough to negate this limitation. The 10, 20 and 40 Myr calculations begin to trend back towards $\delta = \frac{1}{n}$ at later iterations due to the previously observed instability of the linear choice in δ over multiple iterations.

At t_0 the effect of the diverging residuals is more readily apparent, as all the longer length calculations produce a residual that is lower when we chose a non-linear δ . The effects of a significant reduction in residual in the early iterations means that the shorter calculations (≤ 40 Myr) when $\delta = 0.8$ perform better than the non-linear cases. However, the effects of over-correcting the T_0 temperature can be observed as the 10 Myr case shows signs of yielding a better result using the non-linear δ by the final iteration.

To highlight the evolution of T_0^n , in Fig. 8 and 9 we present snapshots of some of the adjoint derived initial conditions at a selection of iterations together with the reference model which we hope to match with for $\delta = 0.8$ and $\frac{1}{n}$ respectively.

The excellent convergence of the misfit observed previously for the 10 Myr calculation is apparent here with minimal visual difference by the fifth iteration for both cases. The reconstructed initial condition displays the same subducting features together with similar lithospheric thickening. Further iterations show very little change in the features despite further decrease in the residuals.

Similarly in both cases, the selected snapshots at 40 Myr also show a good match from the fifth iteration with an improving picture in the subsequent images. The region located at 90° east shows a less robust match though we conclude that the broad structure is correct.

Beginning from 80 Myr onwards we begin to see differences in the two choices in δ . For the case where $\delta = 0.8$, at 80 Myr,

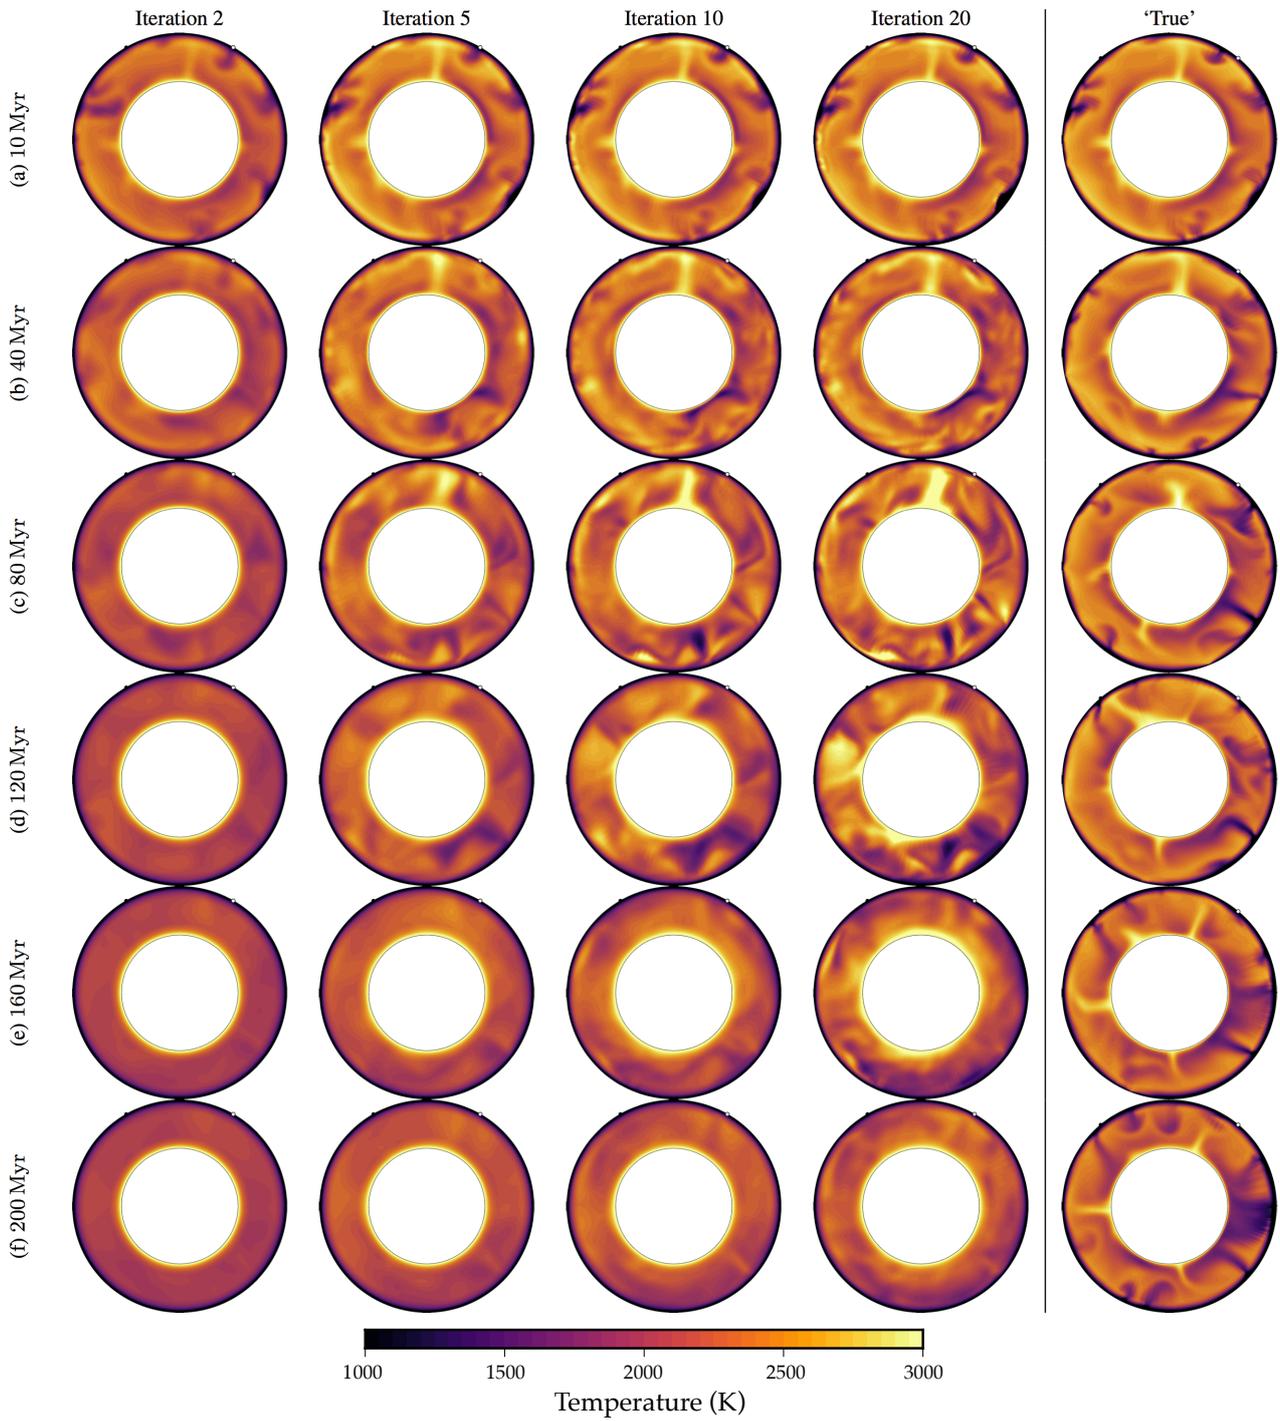


Figure 8. Equatorial cross section plots of our 3D models through the equator at t_0 for various time intervals forward-adjoint calculations when $\delta = 0.8$. A range of iterations are shown highlighting the evolution of T_0^p , together with the expected ‘true’ mantle structure taken from a separate MCM calculation.

many of the features generated in our predicted temperature field differ to our ‘true’ state.

Looking towards the longer time spanning cases, the power of features introduced in the first few iterations are much lower and broader in comparison to the shorter cases. This is particularly pronounced between 120 and 160 Myr where the former shows more defined features compared to the latter. Considering this together with the knowledge that the 120 Myr case showed the worst match to the reference state, we conclude that by having less pronounced

features in the longer cases results in less dramatic shifts from the observed state and hence the observed anomalous ordering in the t_0 residuals.

The longer cases for the non-linear δ (Fig. 9) begin to produce the broad scale, low power features we observed in the linear case, after many more iteration We observe that the 80 Myr case produces less extreme features but which still differ from our anticipated ‘true’ state. Meanwhile the models running > 80 Myr all produce visually similar mantle structures for their initial condi-

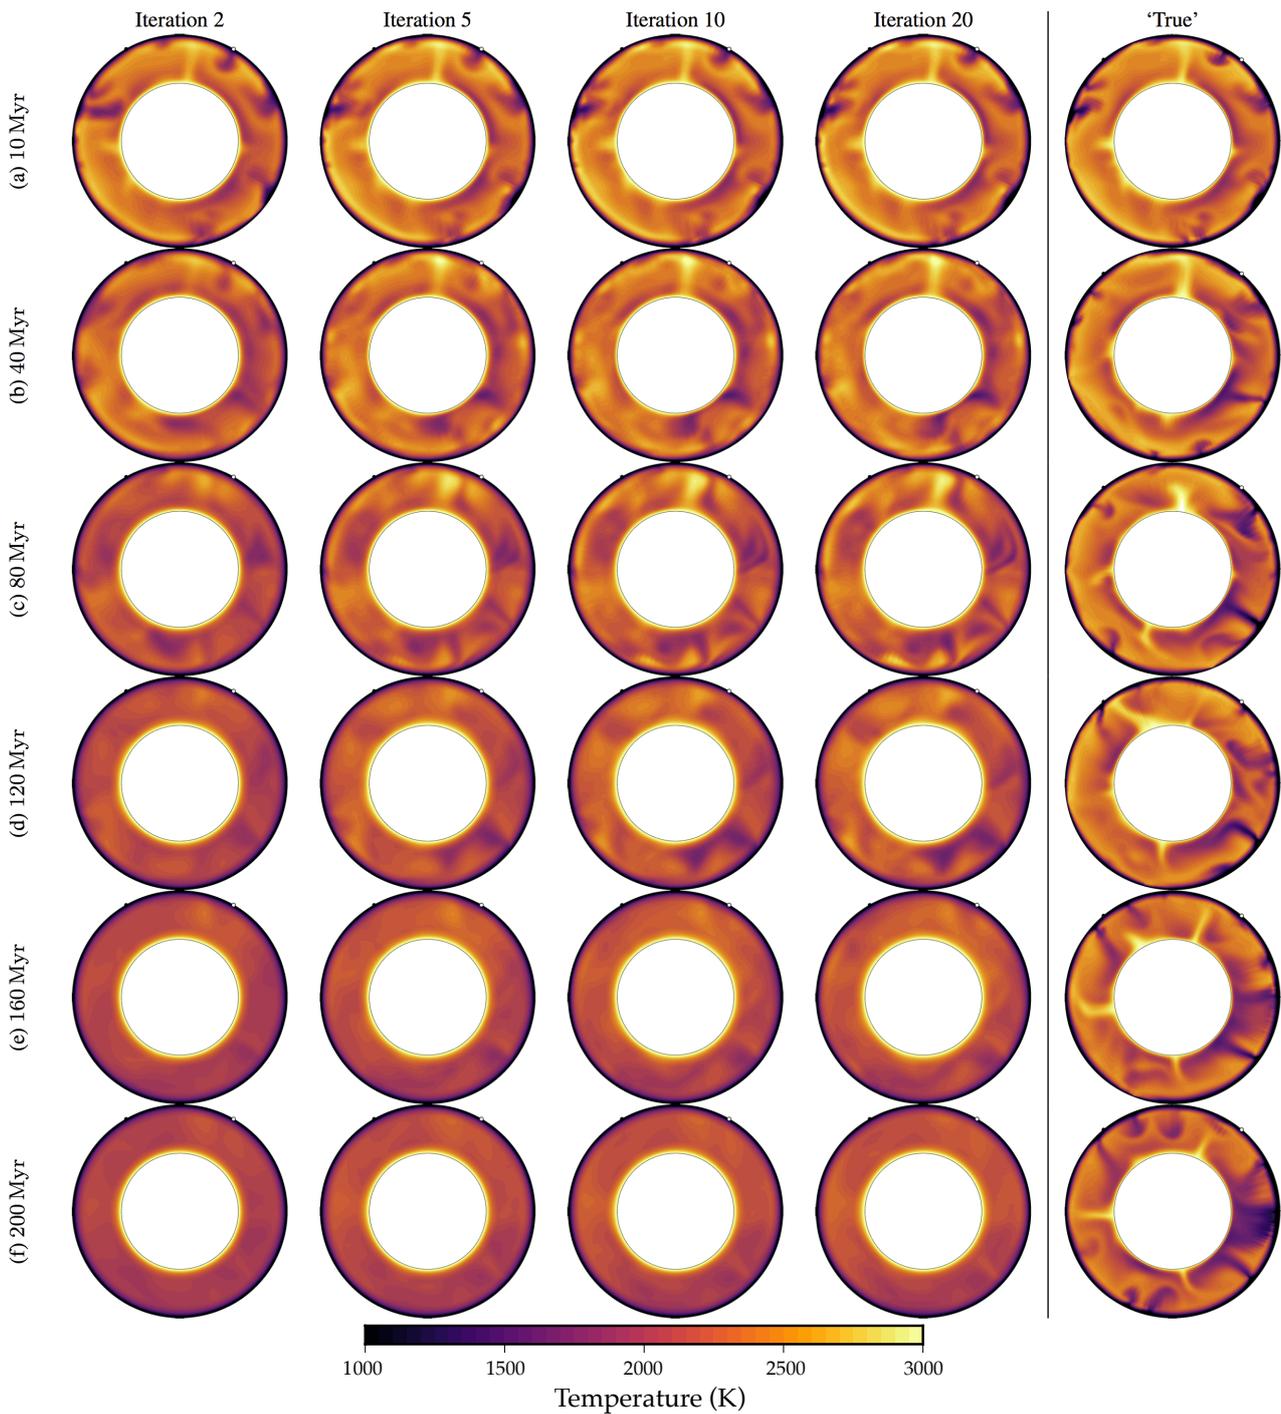


Figure 9. As Fig. 8 but with $\delta = \frac{1}{n}$.

tions following 20 forward-adjoint iterations, with their structure appearing to have evolved minimally over the 20 iterations beyond a broad scale realigning of the radial temperature profile by the fifth iteration.

Finally we recognise that in Colli et al. (2015), the authors note that the L^2 norms can over emphasise small scale differences. As previously stated, the longer scale models show a range of temperature residuals at t_1 , which runs counter to previous studies that suggest models driven by plate reconstructions should converge to a solution. In Fig. 10 we show the T_1^{20} field when $\delta = \frac{1}{n}$ for the

various time intervals next to the corresponding ‘true’ states. Whilst the norms at t_1 vary greatly (from 30 K to over 200 K, Fig. 6d) for our different cases, comparing the various cases side by side, we see that they all contain very similar structures. This does highlight the limitation in using the L^2 norm, but as we have seen from looking at the cross sections of T_0^n , it is still a valuable tool.

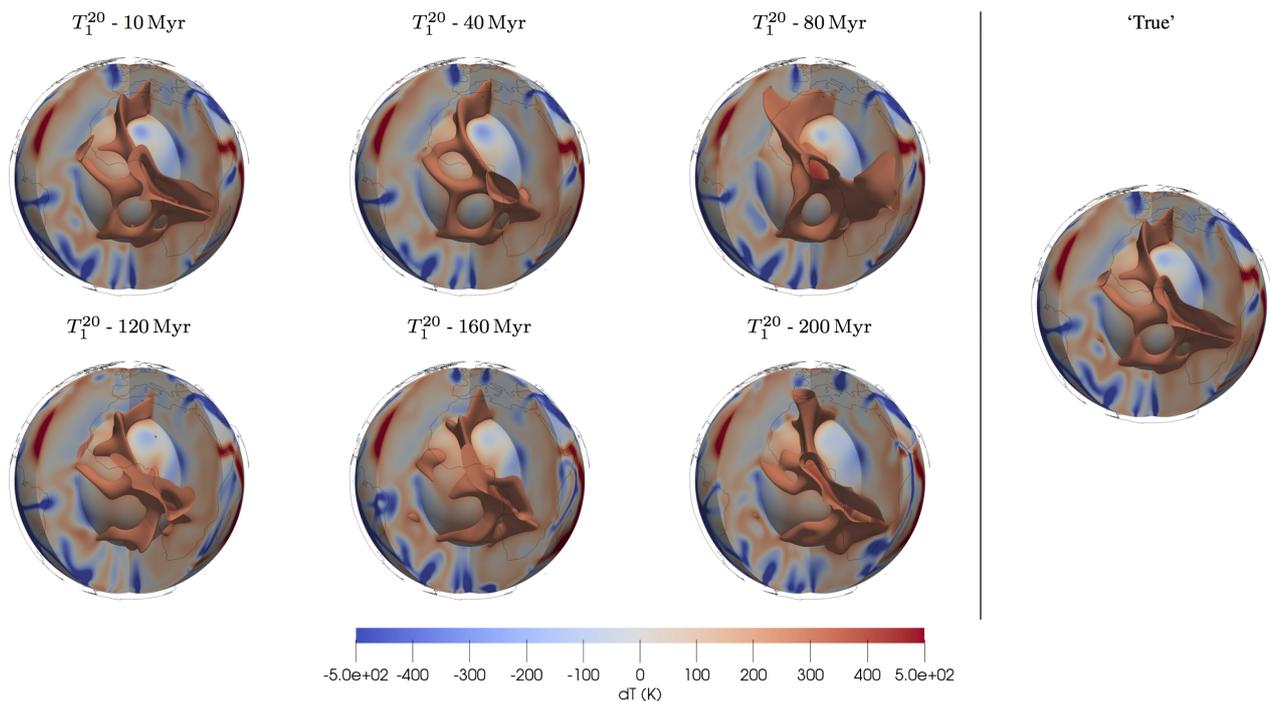


Figure 10. 3D views of the T_1^{20} fields various time intervals when $\delta = \frac{1}{n}$, together with the reference ‘true’ field. Inner and outer boundaries shown are at 2790 and 100 km depth respectively. Lateral temperature variations are shown, with the colour scale saturated at ± 500 K, together with an isosurface at $+300$ K between 300 and 2890 km depth. Present day coastlines (black) are also shown for reference.

4 DISCUSSION

By conducting synthetic tests of the forward-adjoint method we have been able to constrain a number of parameters that can help improve the efficiency of models.

While considering how best to conserve storage and runtime we have seen that all our choices for time-step method had no noticeable effect on the convergence rate. This allows us to view the three choices in mechanism equally without fear of compromising our solution. There is no noticeable best choice for storage savings, as all mechanisms require roughly the same storage requirements outside of the first iteration where forcing the time-step works best. The benefit of saving disk space for a single iteration using the forced time-step is likely minimal even at higher resolutions.

For the runtime required for each of our six test cases the upwards method per iteration, is on average quicker than both the forced and free methods for both calculation lengths. Over the shorter time interval of 10 Myr, the runtime for 20 iterations for the upward method was 689 minutes compared to 805 for the free method (the next fastest). Scaling these values up to the high resolution models that best match Earth’s convective vigour we would estimate to see this difference to amount to a difference of over 24 hours in runtime. Such a difference has great significance when the user has limited access to computing resources as is typical of the queueing systems on HPC systems.

Therefore from our three options we conclude that the upwards time-step is the best of the mechanisms tested. We can see that choosing the freely evolving time-step as in a standard TERRA calculation has no noticeably benefits over the upwards mechanism in terms of storage, runtime or convergence of our forward-adjoint. While before hand one might have made an argument that due to the iterative nature of the forward-adjoint method we can force the time-step to save time and space, actually only small gains are made

in terms of storage ($\mathcal{O} \sim 25$ GB); although at the next increment of resolution we would extrapolate this to correspond to a 400 GB saving at the first iteration. The upwards time-step meanwhile does not compromise on convergences, is no worse or better than the other methods in terms of storage, and consistently shows the best runtimes of the methods tested.

In the same vain of determining ‘shortcuts’ to reduce storage requirements we investigated the benefit of only storing the required \vec{u} and T fields for every tenth time-step. Our results conclude that whilst previous studies which have applied the forward-adjoint method to real Earth data and observed a good convergence to the ‘true’ data at present day, the information obtained in the past does not fair so well. We find from our synthetic tests that the reconstructed t_0 temperatures in fact fail to converge to our anticipated ‘true’ state, whereas when we store all the fields and do not need to interpolate we do see convergence. As previously discussed, the length of each step is limited by the CFL limit, and so we do not anticipate that it is the length of our time-steps which result in the interpolated cases inability to recover past mantle structure.

For our choice in δ we have observed that the more sophisticated choices used in some studies that looked at forward-adjoint models (Liu et al. 2008; Ismail-Zadeh et al. 2004), while achieving a monotonically decreasing residual, do not achieve the best convergence of residuals in our tests. The simple fixed value < 1 as used in Bunge et al. (2003) shows excellent convergence at both t_0 and t_1 for both length calculations. We do however note that there is some oscillation for these fixed values which does not occur when using the simple function $1 - 0.02n$. The strictly monotonically decreasing nature of the non-linear functions is an undoubtably attractive feature in a numerical model, and as computational power increases these methods would be preferential. While we believe a choice of $\delta = 0.8$ or $\delta = 1 - 0.02n$ can be considered the pre-

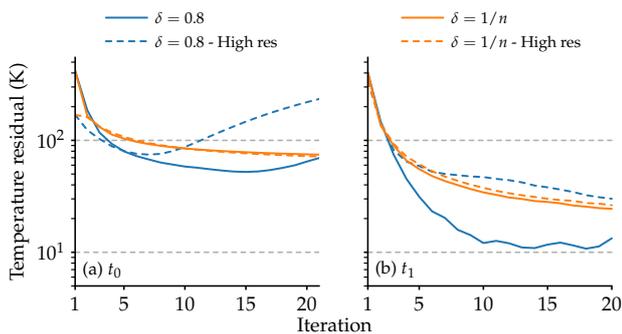


Figure 11. Logarithmic plots of the L^2 norms for the residual temperatures, at t_0 (a) and t_1 (b), of 10 Myr forward-adjoint models where $Ra \sim \mathcal{O}(10^7)$ (as elsewhere in this study) and a higher resolution case where $Ra \sim \mathcal{O}(10^8)$.

ferred factor to use for updating the initial condition temperature field, it is obvious from our results that it should be anticipated that given sufficient forward-adjoint iterations the solutions will readily diverge. Therefore we would advocate the use of a non-linear choice of δ . This would ideally be a modified version of that found in [Ismail-Zadeh et al. \(2004\)](#), such that a large percentage of the adjoint temperature field is applied to T_0^{n+1} in early iterations, but less pronounced further into the calculation.

By varying the length that the forward-adjoint models were run, we have also gained excellent insight into the limits of the method. We conclude that for models extending up to at least 40 Myr we can expect excellent convergence at t_1 and the features present at t_0 can be taken as accurate. For models between 80 to 120 Myr while we would expect to see a converging present day residual we do not observe a similar convergence at t_0 . Models longer than 120 Myr require many more forward-adjoint iterations to introduce meaningful mantle structures, with our snapshots showing minimal downwelling features compared to upwellings, with predominantly broad scale deep mantle features being the main reconstructed features.

Looking at the present day mantle structures we see that despite the differences in the norm values (for the 80 Myr, the $\delta = 0.8$ norm is nearly 60 K better; at 160 Myr, the $\delta = \frac{1}{n}$ norm is 10 K better), the overall structures are very similar. This is inline with the results of [Colli et al. \(2015\)](#), where they had found that a prescribed surface boundary condition limits the divergence of the solutions, but also raises the question of the optimal way of comparing similar mantle volumes.

We recall that the models presented in this study are of a lower convective vigour than that of Earth ($Ra \sim \mathcal{O}(10^7)$ here compared to $Ra \sim \mathcal{O}(10^8)$ in the mantle). Despite accounting for the lower convective vigour in terms of scaling our models to have the same transit times, the thermal diffusion time is longer for our lower Rayleigh number cases compared to Earth. Therefore, as thermal diffusion has more time to act in our scaled, lower Rayleigh number cases, we would anticipate our results here to be a lower bound for the time interval over which information may be recovered using the adjoint method for higher, more realistic, Rayleigh number calculations.

Therefore as a final test we ran two final 10 Myr cases at a higher grid resolution, which allows us to use a lower reference viscosity ($\eta_0 = 3 \times 10^{21}$) and therefore increase the Rayleigh number by an order of magnitude. The results of these two cases are pre-

sented in [Fig. 11](#). At high resolution the choice in δ is clear to see, with the linear choice rapidly diverging at t_0 after the first several iterations, allowing us to more robustly discard a simple linear choice for δ . The high resolution non-linear case meanwhile performs marginally better at t_0 compared to our previous case. This reinforces our assumption that our results present a lower bound on the time interval over which the adjoint can be reasonably expected to recover information. The story at t_1 is broadly similar with the high resolution calculations performing almost equal to the lower Rayleigh number calculations.

The difference between converging residuals at t_1 compared to at t_0 leads us to conclude that results from longer spanning models must be taken with caution as even here in our relatively simple synthetic tests discrepancies exist across the whole volume. We recall that in [Bunge et al. \(2003\)](#), their 100 Myr forward-adjoint models presented showed an excellent match to the true state after running for 100 iterations. While we cannot discount that a factor of five increase in iterations here could yield similar results, the practicality of solving this many iterations makes solving 100 iterations unrealistic. Furthermore, there is of course less certainty in the ‘true’ state we are comparing our adjoint results to at 200 to 120 Ma, and so this raises the question of how bad a match these models really are.

We also find, as other studies have previously (e.g. [Horbach et al. \(2014\)](#)), that calculated norms for the upper mantle are worse than the lower mantle. This is attributed to the t_0 predicted field having to assimilate the effects of crustal and subducting slab structures.

In spite of this, even if we can only reliably use deep mantle features generated from the adjoint method, this is not without its uses in Earth sciences. Deep Earth mantle structures have been shown to have a significant influence on surface dynamics ([Lithgow-Bertelloni & Richards 1998](#)) and the long term nature of deep Earth features such as the Large Low Shear Velocity Provinces (LLSVP), which have been the focus of several studies (e.g. [Zhang et al. 2010](#); [Bower et al. 2013](#); [Bull et al. 2014](#)).

We note that a simpler model utilising a ‘backwards in time’ calculation that only reverses the time dependent terms of the governing equations (eqs. (1)) and ignores the diffusive term performed calculations back from 75 to 126 Ma before drawing conclusions ([Conrad & Gurnis 2003](#)). In their study however, the authors observe that instabilities in the thermal boundary layers make it difficult for this method to accurately reconstruct the present day structures the model began with. Due to this, we conclude that the method used in our study is superior to these methods even when ran over shorter time frames due to the lack of instabilities in the boundary layers.

A third technique of recovering past mantle structures that is of note is the quasi-reversibility (QRV) method ([Ismail-Zadeh et al. 2007](#)). This method is similar to the adjoint in that backwards heat, motion and continuity equations are solved, but here the backwards heat equation contains an additional term that involves a small regularisation parameter together with a higher order temperature derivative. A best fit between observational and model fields is then sought by minimising the regularisation parameter. [Glišović & Forte \(2014\)](#) and [Glišović & Forte \(2017\)](#) utilise this technique in global models to investigate regional evolution dynamics over the previous 70 Myr. In [Glišović & Forte \(2014\)](#), the QRV method is tested on a known data source, with the accuracy of the QRV predictions observed to begin falling away beyond 45 Myr. The error levels up to 70 Ma are determined to be acceptable however, allowing the authors to investigate the evolution of Earth over the entire

Cenozoic. In Glišović & Forte (2017) the QRV method is applied to observational data to investigate the evolution of upwellings in the Indian ocean. With the time intervals over which these QRV studies can reliably be calculated over falling broadly in line with the results shown here, choosing either the adjoint or QRV method can be argued from both sides.

This investigation into adjoints run over different time intervals is more intriguing when considering previous work that looks at error growth in mantle models, which look to determine a ‘limit of predictability’ (Bello et al. 2014; Colli et al. 2015). Results from these studies suggest at their most conservative estimates, a limit of predictability of 95 Myr (Bello et al. 2014) without an assimilated surface boundary condition. With an assimilated surface it is suggested that solutions will evolve towards a single solution, which corroborates the theorem of Serrin (1959) which states that two incompressible viscous bodies are equivalent given their boundary conditions are the same. Ismail-Zadeh et al. (2007) meanwhile determine that the adjoint solution is limited by the characteristic thermal diffusion time, which for large features (> 100 km) suggests a limit over 300 Myr, and less than 5 Myr for fine features (< 10 km). This falls in line with what we have seen here, with the large scale features being predicted for all time scales (given sufficient iterations) and finer features not captured (such as in the upper mantle). We therefore conclude that despite being within the limit of predictability for this class of models, it is the thermal diffusion time which is limiting our adjoint predictions.

5 CONCLUSION

We have investigated a number of parameters that can be varied in the full forward-adjoint system for mantle convection. A minimum number of modelling assumptions have been made in our investigation here and so our results are relatively straightforward to interpret. In order to optimise the runtime and storage requirements of the forward-adjoint calculation one should consider using a gentle forcing of the time-step mechanism, as doing so sees moderate improvements in terms of calculation speed whilst not compromising on the misfit convergences. Furthermore whilst simple linear functions for the choice of δ (the factor that controls the amplitude of the adjoint solution used to update the initial condition) can achieve a minimum residual value in relatively few iterations, extreme caution must be taken as their t_0 solutions have been shown to quickly diverge. More complex non-linear choices on the other hand have been shown to converge monotonically at t_0 for a variety of time intervals, albeit at a slower rate compared to linear choices. Therefore unless the T_0 structure is known already, the use of a linear factor for δ could unwittingly lead to incorrectly calculated T_0 structures.

The results from investigating varying the time interval that the forward-adjoint calculations are run over suggest that drawing conclusions of fine features should be done with caution for longer calculations. It was observed that recreated initial conditions over time intervals of greater than 40 Myr showed increasingly poorer recreation of the short scale features. We do believe though that the results found here can act as a lower bound on the time interval for models run at a higher Rayleigh number (i.e. $Ra \sim \mathcal{O}(10^8)$).

Despite the problems in recreating short scale features, initial conditions derived using the adjoint method may still be better than methods currently employed in the geodynamics community for the recreation of broad scale features. The effect using these adjoint derived initial conditions have on the final output compared to current methods is an open question that could readily be investigated

further. As an example, the effects on dynamic topography over recent Earth history is one avenue of investigation that can benefit from these adjoint derived initial conditions (Colli et al. 2017).

With these results, together with those of Horbach et al. (2014); Vynnytska & Bunge (2014); Ghelichkhan & Bunge (2016) which also investigate some of the parameters involved in running adjoint models, future adjoint model based investigations can be focused on applying the method to real world scenarios using the wealth of present day information available. As the examples presented here utilise a synthetic data source in order to ascertain this methods viability, future studies could also be focused on investigating how true data sources based on Earth observations are incorporated into forward-adjoint methods.

ACKNOWLEDGMENTS

Both authors contributed to writing the paper. MGP developed the code and tests; and carried out the simulations, data processing and visualisation. JHD conceived the project. MGP thanks Hans-Peter Bunge and André Horbach for discussions on the adjoint method. We would also like to thank two anonymous reviewer for their helpful comments, which greatly helped to improve this manuscript. This work was supported by Fujitsu. All computations in this study were carried out on the Welsh national supercomputer system HPC Wales. Figures were created using Matplotlib (Hunter 2007), the Generic Mapping Tools (GMT) software (Wessel et al. 2013) and Paraview.

REFERENCES

- Baumgardner, J. R., 1985. Three-dimensional treatment of convective flow in the earth’s mantle, *J. Stat. Phys.*, **39**(5-6), 501–511.
- Bello, L., Coltice, N., Rolf, T., & Tackley, P. J., 2014. On the predictability limit of convection models of the earth’s mantle, *Geochem. Geophys. Geosyst.*, **15**(6), 2319–2328.
- Bower, D. J., Gurnis, M., & Seton, M., 2013. Lower mantle structure from paleogeographically constrained dynamic earth models, *Geochem. Geophys. Geosyst.*, **14**(1), 44–63.
- Bull, A. L., Domeier, M., & Torsvik, T. H., 2014. The effect of plate motion history on the longevity of deep mantle heterogeneities, *Earth planet. Sci. Lett.*, **401**, 172–182.
- Bunge, H.-P., Richards, M. A., & Baumgardner, J. R., 1997. A sensitivity study of three-dimensional spherical mantle convection at 10^8 rayleigh number: Effects of depth-dependent viscosity, heating mode, and an endothermic phase change, *J. geophys. Res.*, **102**(B6), 11991–12007.
- Bunge, H.-P., Richards, M. A., & Baumgardner, J. R., 2002. Mantle circulation models with sequential data assimilation: Inferring present-day mantle structure from plate-motion histories, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **360**(1800), 2545–2567.
- Bunge, H.-P., Hagelberg, C. R., & Travis, B. J., 2003. Mantle circulation models with variational data assimilation: Inferring past mantle flow and structure from plate motion histories and seismic tomography, *Geophys. J. Int.*, **152**(2), 280–301.
- Colli, L., Bunge, H.-P., & Schuberth, B. S. A., 2015. On retrodictions of global mantle flow with assimilated surface velocities, *Geophys. Res. Lett.*, **42**(20), 8341–8348.
- Colli, L., Ghelichkhan, S., Bunge, H.-P., & Oeser, J., 2017. Retrodictions of mid paleogene mantle flow and dynamic topography in the atlantic region from compressible high resolution adjoint mantle convection models: Sensitivity to deep mantle viscosity and tomographic input model, *Gondwana Res.*

- Conrad, C. P. & Gurnis, M., 2003. Seismic tomography, surface uplift, and the breakup of gondwanaland: Integrating mantle convection backwards in time, *Geochem. Geophys. Geosyst.*, **4**(3), 1031.
- Courtier, P. & Talagrand, O., 1987. Variational assimilation of meteorological observations with the adjoint vorticity equation. II: Numerical results, *Q. J. Roy. Meteor. Soc.*, **113**(478), 1329–1347.
- Fletcher, R. & Reeves, C. M., 1964. Function minimization by conjugate gradients, *Comput. J.*, **7**(2), 149–154.
- Ghelichkhan, S. & Bunge, H.-P., 2016. The compressible adjoint equations in geodynamics: Derivation and numerical assessment, *GEM - International Journal on Geomathematics*, **7**(1), 1–30.
- Glišović, P. & Forte, A. M., 2014. Reconstructing the cenozoic evolution of the mantle: Implications for mantle plume dynamics under the pacific and indian plates, *Earth planet. Sci. Lett.*, **390**, 146–156.
- Glišović, P. & Forte, A. M., 2017. On the deep-mantle origin of the deccan traps, *Science*, **355**(6325), 613–616.
- Horbach, A., Bunge, H.-P., & Oeser, J., 2014. The adjoint method in geodynamics: Derivation from a general operator formulation and application to the initial condition problem in a high resolution mantle circulation model, *GEM - International Journal on Geomathematics*, **5**(2), 163–194.
- Hunter, J. D., 2007. Matplotlib: A 2d graphics environment, *Comput. Sci. Eng.*, **9**(3), 90–95.
- Ismail-Zadeh, A., Schubert, G., Tsepelev, I., & Korotkii, A., 2004. Inverse problem of thermal convection: numerical approach and application to mantle plume restoration, *Phys. Earth planet. Inter.*, **145**(1-4), 99–114.
- Ismail-Zadeh, A., Korotkii, A., Schubert, G., & Tsepelev, I., 2007. Quasi-reversibility method for data assimilation in models of mantle dynamics, *Geophys. J. Int.*, **170**(3), 1381–1398.
- Lithgow-Bertelloni, C. & Richards, M. A., 1998. The dynamics of cenozoic and mesozoic plate motions, *Rev. Geophys.*, **36**(1), 27–78.
- Liu, L. & Gurnis, M., 2008. Simultaneous inversion of mantle properties and initial conditions using an adjoint of mantle convection, *J. geophys. Res.*, **113**(B8), B08405.
- Liu, L., Spasojevic, S., & Gurnis, M., 2008. Reconstructing farallon plate subduction beneath north america back to the late cretaceous, *Science*, **322**(5903), 934–938.
- Menemenlis, D. & Wunsch, C., 1997. Linearization of an oceanic general circulation model for data assimilation and climate studies, *J. Atmos. Ocean. Tech.*, **14**(6), 1420–1443.
- Nerlich, R., Colli, L., Ghelichkhan, S., Schuberth, B., & Bunge, H.-P., 2016. Constraining central neo-tethys ocean reconstructions with mantle convection models, *Geophys. Res. Lett.*, **43**(18), 9595–9603.
- Serrin, J., 1959. Mathematical principles of classical fluid mechanics, in *Encyclopedia of Physics / Handbuch der Physik*, pp. 125–263, ed. Truesdell, C., Springer Science & Business Media.
- Seton, M., Müller, R., Zahirovic, S., Gaina, C., Torsvik, T., Shephard, G., Talsma, A., Gurnis, M., Turner, M., Maus, S., & Chandler, M., 2012. Global continental and ocean basin reconstructions since 200ma, *Earth-Sci. Rev.*, **113**(3-4), 212–270.
- Spasojevic, S., Liu, L., & Gurnis, M., 2009. Adjoint models of mantle convection with seismic, plate motion, and stratigraphic constraints: North America since the Late Cretaceous, *Geochem. Geophys. Geosyst.*, **10**(5), Q05W02.
- Tarantola, A., 1984. Inversion of seismic reflection data in the acoustic approximation, *Geophysics*, **49**(8), 1259–1266.
- Vynnytska, L. & Bunge, H.-P., 2014. Restoring past mantle convection structure through fluid dynamic inverse theory: Regularisation through surface velocity boundary conditions, *GEM - International Journal on Geomathematics*, **6**(1), 83–100.
- Wessel, P., Smith, W. H. F., Scharroo, R., Luis, J., & Wobbe, F., 2013. Generic mapping tools: Improved version released, *Eos, Transactions American Geophysical Union*, **94**(45), 409–410.
- Zhang, N., Zhong, S., Leng, W., & Li, Z.-X., 2010. A model for the evolution of the earth's mantle structure since the early paleozoic, *J. geophys. Res.*, **115**(B6), B06401.