



Technical Note

Efficient gradient computation for dynamical models



B. Sengupta*, K.J. Friston, W.D. Penny

Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, 12 Queen Square, London WC1N 3BG, UK

ARTICLE INFO

Article history:

Accepted 12 April 2014

Available online 23 April 2014

Keywords:

Augmented Lagrangian

Adjoint methods

Dynamical systems

Dynamic causal modelling

Model fitting

ABSTRACT

Data assimilation is a fundamental issue that arises across many scales in neuroscience – ranging from the study of single neurons using single electrode recordings to the interaction of thousands of neurons using fMRI. Data assimilation involves inverting a generative model that can not only explain observed data but also generate predictions. Typically, the model is inverted or fitted using conventional tools of (convex) optimization that invariably extremise some functional – norms, minimum descriptive length, variational free energy, etc. Generally, optimisation rests on evaluating the local gradients of the functional to be optimized. In this paper, we compare three different gradient estimation techniques that could be used for extremising any functional in time – (i) finite differences, (ii) forward sensitivities and a method based on (iii) the adjoint of the dynamical system. We demonstrate that the first-order gradients of a dynamical system, linear or non-linear, can be computed most efficiently using the adjoint method. This is particularly true for systems where the number of parameters is greater than the number of states. For such systems, integrating several sensitivity equations – as required with forward sensitivities – proves to be most expensive, while finite-difference approximations have an intermediate efficiency. In the context of neuroimaging, adjoint based inversion of dynamical causal models (DCMs) can, in principle, enable the study of models with large numbers of nodes and parameters.

© 2014 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/3.0/>).

Introduction

An important goal of systems neuroscience is to integrate empirical data from various neuroimaging modalities with biologically informed models that describe the underlying generative processes. Here, the data to be explained are for example M/EEG and fMRI recordings made while subjects perform various experimental tasks, and the underlying neurodynamic processes are framed in terms of differential equations describing activity in neural masses, mean fields, or neural fields (David et al., 2006; Deco et al., 2008; Friston et al., 2003).

Considerable insight can be gained from studying the emergent properties of such neurodynamic processes. These can then be qualitatively compared with empirical data, allowing consilience among multiple levels of description (Gazzaniga, 2010; Hopfield and Brody, 2001; Wilson, 1999). An alternative approach is to directly fit neurodynamical models to neuroimaging data using standard model fitting procedures from statistics and machine learning (Bishop, 2006; Press et al., 1992). Differences in the generative processes induced by experimental manipulations can then be associated with changes in underlying brain connectivity. One example of such an approach is Dynamic Causal

Modelling (DCM) (Friston et al., 2003) which fits differential equation models to neuroimaging data using a variational Bayesian scheme (Friston et al., 2007).

More generally, in the statistics and machine learning literature various methods have been employed to fit differential equations to data, from maximum likelihood approaches (Ramsay et al., 2007) to Bayesian sampling algorithms (Calderhead and Girolami, 2009; Vyshemirsky and Girolami, 2008). The majority of these convex optimisation approaches involve computing the gradient; the change in the cost function produced by a change in model parameters. This gradient is then combined with information from line searches (e.g., Wolfe's conditions) or methods involving a Newton, quasi-Newton (low-rank) or Fisher information based curvature estimators to update model parameters (Bishop, 1995; Nocedal and Wright, 2006; Press et al., 1992). The main computational bottleneck in these algorithms is the computation of the gradient (or the curvature) of the parametric cost function. This motivates the search for efficient methods to evaluate gradients.

This paper compares three different methods for computing gradients, and studies the conditions under which each is preferred. The first is the Finite Difference (FD) method, which is the simplest and most general method – and is currently used in DCM. The second is the Forward Sensitivity (FS; also known as tangent linear) method, which has previously been proposed in the context of modeling fMRI time series (Deneux and Faugeras, 2006). The third is the Adjoint Method

* Corresponding author.

E-mail addresses: b.sengupta@ucl.ac.uk (B. Sengupta), k.friston@ucl.ac.uk (K.J. Friston), w.penny@ucl.ac.uk (W.D. Penny).

(AM) which has previously been used in the context of dynamical systems theory (Wang, 2013), weather forecasting (Errico, 1997), image registration (Clark, 2011) and single-neuron biophysics (Stemmler et al., 2012).

The paper is structured as follows – the methods section describes each approach including a mathematical derivation of the adjoint method. Examples of the FS and AM updates are then provided for the case of simple Euler integration. The results section reports numerical simulations that disclose the scaling characteristics of each method. Simulations are provided for linear dynamical and weakly-coupled oscillator systems. We conclude with a discussion of the relative merits of each method.

Methods

We consider dynamical systems of the form

$$\begin{aligned} \dot{x} &= f(x, p) \\ j(x, p) &= -\frac{1}{2} \|y - g(x, p)\|^2 \end{aligned} \quad (1)$$

where x is a state variable, the dot notation denotes a time derivative $\frac{dx}{dt}$ t is time, $f(\cdot)$ is the flow equation (dynamics), and p are model parameters. The model produces a prediction via an observation function $g(x, p)$ and an instantaneous cost function $j(x, p)$ measures the squared difference from data points y . The total cost is then given by the integral up to time point T

$$J(p) = \int_0^T j(x, p) dt. \quad (2)$$

We consider three methods for computing the gradient $\frac{dJ}{dp}$.

Finite difference method

The (one-sided) finite difference approximation to the gradient is then

$$\frac{dJ}{dp_i} = \frac{J(p + \delta_i) - J(p)}{\delta_i} \quad (3)$$

where δ_i denotes a small change (generally, $\sqrt{\epsilon}$ where ϵ is the machine epsilon) to the i th parameter. The error in the computation of this gradient is of order δ_i . The computation of $\frac{dJ}{dp}$ requires $P + 1$ runs of the integration process, one for each model parameter. It is also possible to use central differences

$$\frac{dJ}{dp_i} = \frac{J(p + \delta_i) - J(p - \delta_i)}{2\delta_i} \quad (4)$$

which has an error of order δ_i^2 but requires $2P + 1$ runs of the integration process. Variations on the vanilla FD approach are discussed in (Press et al., 1992; Richtmeyer and Morton, 1967).

Forward Sensitivity method

The original dynamical model (Eq. (1)) can be implicitly differentiated w.r.t parameters to give

$$\frac{d\dot{x}}{dp} = \frac{\partial f}{\partial x} \frac{dx}{dp} + \frac{\partial f}{\partial p}. \quad (5)$$

If the state variables are of dimension D and the parameters of dimension P then the quantity $\frac{dx}{dp}$ is a $D \times P$ matrix, which can be vectorized to form a new flow function. This forms a new dynamical system of dimension $D \times P$ that can then be integrated using any numerical method to produce $\frac{dx}{dp}$ as a function of time. The Forward

Sensitivity approach has been known since the publication of Gronwall's theorem (Gronwall, 1919). The cost gradient is then given by accumulating the sensitivity derivative $\frac{dx}{dp}$ over time according to:

$$\begin{aligned} \frac{dJ}{dp} &= \int_0^T \frac{dj}{dp} dt \\ \frac{dj}{dp} &= \frac{\partial j}{\partial x} \frac{dx}{dp} + \frac{\partial j}{\partial p} \\ &= \frac{\partial j}{\partial g} \frac{\partial g}{\partial x} \frac{dx}{dp} + \frac{\partial j}{\partial g} \frac{\partial g}{\partial p}. \end{aligned} \quad (6)$$

Euler example

This section illustrates the FS approach first-order Euler integration of the dynamics

$$x_n = x_{n-1} + \tau f(x_{n-1}, p) \quad (7)$$

at discrete times $t(n)$. The FS method is based on differentiating this equation to give

$$\frac{dx_n}{dp} = \frac{dx_{n-1}}{dp} + \tau \left[\frac{\partial f}{\partial x_{n-1}} \frac{dx_{n-1}}{dp} + \frac{\partial f}{\partial p} \right]. \quad (8)$$

This method is illustrated in Fig. 1 where the solid path indicates a trajectory of points x_n for a dynamical system with parameters p and the dotted path indicates the trajectory \underline{x}_n for the same dynamical system but with parameters $\underline{p} = p + \delta_i$. The dotted path can be obtained from the solid path via the total derivative $\frac{dx_n}{dp}$ in the direction of the perturbation, δ_i . The FS method provides a method for computing this derivative. Under a first order Euler approach for integrating the dynamics, this is implemented using the above recursion.

Because the perturbed path (dotted in Fig. 1) can be reached from the original trajectory via the total derivative $\frac{dx_n}{dp}$, there is no need to separately integrate the system with parameters \underline{p} . Geometrically, the points \underline{x}_n in Fig. 1 can be reached via the solid and dashed lines (rather than the dotted lines).




Fig. 1. Forward Sensitivity The solid path indicates a trajectory of points x_n , with $n = 1 \dots 5$, for a dynamical system with parameters p . The dotted path indicates the trajectory \underline{x}_n for the same dynamical system but with parameters $\underline{p} = p + \delta_i$. The dotted path can be reached from the solid path via the total derivative $\frac{dx_n}{dp}$. The Forward Sensitivity approach provides a method for computing this derivative.

We rewrite the recursion equation as

$$S_x(n) = S_x(n-1) + \tau [F_x(n-1)S_x(n-1) + F_p(n-1)] \quad (9)$$

where

$$\begin{aligned} F_p(n) &= \left. \frac{\partial f}{\partial p} \right|_{x_n} \\ F_x(n) &= \left. \frac{\partial f}{\partial x} \right|_{x_n} \\ S_x(n) &= \left. \frac{dx}{dp} \right|_{x_n}. \end{aligned} \quad (10)$$

S_x is a $[D \times P]$ matrix, F_x is $[D \times D]$ and F_p is $[D \times P]$. We then have

$$\begin{aligned} \frac{dJ}{dp} &= \sum_{n=1}^N \frac{\partial j}{\partial x_n} \frac{dx_n}{dp} \\ &= \sum_{n=1}^N j_x(n) S_x(n) \end{aligned} \quad (11)$$

and $j_x(n)$ is the derivative of $j(x, p)$ with respect to x , evaluated at x_n . This method requires the derivatives F_x and F_p . These will be specific to the dynamical model in question and, in this paper, are computed analytically. We provide the Euler example here as a simple illustration of the method. The numerical simulations in this paper use a more accurate integration method (see below).

Adjoint method

Errico (1997) and Giles and Pierce (2000) provide introductions to the adjoint method. A derivation of the adjoint method for dynamical models is provided rigorously in Cao et al. (2003) and in Hindmarsh and Serban (2002). Here, we provide an informal derivation, starting with the cost function

$$J(p) = \int_0^T j(x, p) dt. \quad (12)$$

The constraints implied by the dynamics allow us to write the Lagrangian

$$\mathcal{L}(p) = \int_0^T j(x, p) dt + \int_0^T \lambda^T [\dot{x} - f(x, p)] dt. \quad (13)$$

Once the system has been integrated (solved for x) we have $\dot{x} = f(x, p)$. Hence the second term in the Lagrangian disappears and we have

$$\frac{dJ}{dp} = \frac{d\mathcal{L}}{dp}. \quad (14)$$

This is the gradient we wish to compute. So far it may seem that we have made no progress but it turns out that $\frac{d\mathcal{L}}{dp}$ can be computed efficiently.

Before proceeding further, we summarize the main ideas behind the adjoint method. The central concept is that the Lagrange vector λ^T constrains the dynamical system to variations around the forward path x_n . The Lagrange vectors are of the same dimension as x and form a time series. Algebraically, the contribution of the total derivative $\frac{d\mathcal{L}}{dp}$ to the gradient $\frac{dJ}{dp}$ is made zero, by setting λ^T appropriately. This means that the sensitivity derivative need not be calculated, resulting in a large computational saving. Instead, the gradient $\frac{dJ}{dp}$ can be expressed as a function of λ^T . We will now go through this in a bit more detail:

The proof proceeds by differentiating Eq. (13) to give the gradient

$$\frac{dJ}{dp} = \int_0^T \left(\frac{\partial j}{\partial x} \frac{dx}{dp} + \frac{\partial j}{\partial p} \right) dt + \int_0^T \lambda^T \left(\frac{d\dot{x}}{dp} - \frac{\partial f}{\partial x} \frac{dx}{dp} - \frac{\partial f}{\partial p} \right) dt. \quad (15)$$

The term involving the change in total derivative, $\frac{d\dot{x}}{dp}$ can be rewritten using integration by parts

$$\int_0^T \lambda^T \frac{d\dot{x}}{dp} dt = \left[\lambda^T \frac{dx}{dp} \right]_0^T - \int_0^T \frac{d\lambda^T}{dt} \frac{dx}{dp} dt. \quad (16)$$

Substituting this into the previous expression and rearranging to group together terms involving the sensitivity derivative $\frac{dx}{dp}$ give

$$\begin{aligned} \frac{dJ}{dp} &= \int_0^T \frac{dx}{dp} \left(\frac{\partial j}{\partial x} - \lambda^T \frac{\partial f}{\partial x} - \frac{d\lambda^T}{dt} \right) dt \\ &\quad + \int_0^T \left(\frac{\partial j}{\partial p} - \lambda^T \frac{\partial f}{\partial p} \right) dt + \left[\lambda^T \frac{dx}{dp} \right]_0^T. \end{aligned} \quad (17)$$

The adjoint vector λ^T can be used to eliminate the first term involving the sensitivity derivative. This term is zero when:

$$\frac{d\lambda^T}{dt} = \frac{\partial j}{\partial x} - \lambda^T \frac{\partial f}{\partial x}. \quad (18)$$

This is known as the adjoint equation and is used to compute λ^T . The gradient is then given by

$$\frac{dJ}{dp} = \int_0^T \left(\frac{\partial j}{\partial p} - \lambda^T \frac{\partial f}{\partial p} \right) dt + \left[\lambda^T \frac{dx}{dp} \right]_0^T. \quad (19)$$

As our goal has been to avoid computation of the sensitivity derivative, $\frac{dx}{dp}$, we can eliminate the last term above by integrating the adjoint equations backward in time, starting with $\lambda_T = 0$. The starting value for the adjoint equation is arbitrary and it can be proven that if $\lambda|_{t=t_{f,a}}$ and $\lambda|_{t=t_{f,b}}$ are two different starting values for the adjoint equation with solutions λ_a and λ_b respectively, then $\left. \frac{dJ}{dp} \right|_{\lambda_a} = \left. \frac{dJ}{dp} \right|_{\lambda_b}$. Therefore, there exist infinitely many starting conditions for the adjoint equation that yields the same parametric gradient.

If the initial conditions do not depend on the parameters, as we assume for our numerical examples, then we have $\frac{dx}{dp} = 0$ at $t = 0$ and the gradient reduces to

$$\frac{dJ}{dp} = \int_0^T \left(\frac{\partial j}{\partial p} - \lambda^T \frac{\partial f}{\partial p} \right) dt. \quad (20)$$

This equality can now be used to compute the parametric gradients, given the backwards solution of the adjoint equation.

There are no restrictions on the functional form to make the adjoint method viable — if one can pose the optimization problem via a Lagrangian, then the adjoint method could be used for any dynamical system (ordinary-, delay-, random- and partial-differential equation). The one and only constraint is the stability of the adjoint equation for the underlying dynamical system. Thus, static or dynamical systems that are routinely used in neuroimaging are amenable to an adjoint formulation under some loss-function including stochastic DCMs that have an analytical model for the noise (of class C^ω). Table 1 highlights the key differences between all of the methods and the crucial steps required in each of them.

Table 1

Comparison of the different gradient computation methods. The flow eqn. is either linear or non-linear, with P parameters and N state variables.

	Finite differences	Forward sensitivities	Adjoint
Suitability	Arbitrary	$N \gg P$	$P \gg N$
Cost	$(1 + P)$ flow eqns.	P non-linear sensitivity eqns. + 1 flow eqn.	1 linear adjoint eqn. + 1 flow eqn.
Key steps	1. Integrate flow eqn. 2. Parametrically perturb flow P times	1. Integrate the coupled flow and sensitivity eqns.	1. Integrate flow eqn. 2. Integrate adjoint eqn.

Euler example

The specification of the adjoint method starts from the specification of the Lagrangian. For us, this has the particular form

$$\mathcal{L} = -\frac{1}{2} \sum_n \|y_n - g(x_n)\|^2 + \sum_n \lambda_n [x_n - x_{n-1} - \tau f(x_{n-1}, p)] \quad (21)$$

where the first term is the original cost function, the second term enforces the constraint embodied in the Euler integration of the state dynamics, and λ_n is a $[1 \times D]$ vector of Lagrange multipliers. Because \mathcal{L} is a scalar, and the state x_n is a column vector, the Lagrange multipliers must be a row vector. It is in this sense that they are adjoint (or transposed) to the states. The derivative of \mathcal{L} with respect to the states is then given by

$$\frac{d\mathcal{L}}{dx_n} = g_x(n)[y_n - g(x_n)] + \lambda_n - \lambda_{n+1} - \tau \lambda_{n+1} F_x(n) \quad (22)$$

where $g_x(n)$ is the derivative of $g(x, p)$ with respect to x , evaluated at x_n . Setting Eq. (22) to zero (i.e., solving for the states) gives

$$\lambda_n = \lambda_{n+1} [I + \tau F_x(n)] - g_x(n)[y_n - g(x_n)]. \quad (23)$$

This is a backward recursion, known as the adjoint equation, that starts with $\lambda_N = 0$. After solving the adjoint equations we can enter λ_n into Eq. (20), giving

$$\frac{dj}{dp} = \tau \sum_{n=1}^N [j_p(n) - \lambda_n F_p(n)] \quad (24)$$

where $j_p(n)$ is the derivative of $j(x, p)$ with respect to p , evaluated at x_n . If the observation function does not depend on model parameters then the first term disappears. A first order Euler Adjoint method has been used previously in the context of image registration (Clark, 2011). However, we provide the Euler example here as an illustration of the method. The numerical simulations in this paper use a more accurate integration method (see below).

Stability

It is known that if flows are prescribed as ODEs, then their adjoint solutions are also stable (Cao et al., 2003). Under these conditions, the numerical stability of the adjoint system is guaranteed when the adjoint equation is integrated backwards in time, in the sense that the flow is reversed. Consider a linear autonomous system, $f = Ax + B$ where $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^n$ and both are invariant in time. Being linear in states with pre-defined initial conditions, such a system can be analytically integrated to yield a solution as a sum of its n matrix exponentials with unique eigenvectors and their respective eigenvalues. Such a system is asymptotically stable when the eigenvalues have a negative real part i.e., $\text{Re}(A(A)) < 0$. For such a linear autonomous system the eigenvalues of the adjoint equation, $\dot{\lambda} = -\left(\frac{df}{dx}\right)^T \cdot \lambda = -A^T \lambda$ have a positive real part, proving to be asymptotically unstable. If one were now to reverse the flow i.e., $\dot{\lambda}_* = \left(\frac{df}{dx}\right)^T \cdot \lambda_* = A^T \lambda_*$, the eigenvalues then have a negative real part and the dynamics is asymptotically stable.

One can derive similar results for non-linear equations using a perturbation expansion, suggesting the condition of asymptotic and uniform stability is guaranteed when the adjoint equations are integrated backwards.

Integration

For the numerical results in this paper, we used MATLAB's ode15s function which implements a variable order method for integrating stiff differential equations (Shampine and Reichelt, 1997). Two important parameters governing the operation of this algorithm are the absolute, a , and relative, r , error tolerances. The estimated error in each integration step is constrained to be less than $\max(r|x_n|, a)$.

The absolute and relative tolerances were set to 10^{-7} for each of the gradient computation methods although results were also obtained with different sets of tolerances, taking on values a_{FD} , a_{FS} , a_{AM} and r_{FD} , r_{FS} , r_{AM} for the Finite Difference, Forward Sensitivity and Adjoint Methods respectively. When tolerances were set differently, these values were tuned for each problem (linear/nonlinear) so as to achieve good agreement among the methods.

Results

Custom MATLAB scripts were written to implement each of the gradient computation methods.

Linear models

First we consider the linear models

$$\dot{x} = Ax \quad (25)$$

where x is a D -dimensional state vector with initial value $x_0 = 1$, and A is a $D \times D$ connectivity matrix. Readers familiar with DCM for fMRI will recognize A as the endogenous or average connectivity matrix. A model with D states therefore has $P = D^2$ parameters (Fig. 2(A)). The system is integrated from time 0 to T . To ensure stability, we constructed A using the linear expansion

$$A = \sum_{d=1}^D q_d v_d v_d^T \quad (26)$$

where $v_d \sim \mathcal{N}(v_d; 0, 1)$ are standard D -dimensional Gaussian random vectors, which are serially orthogonalized. The scalars q_d are negative real numbers so that the corresponding eigenstates are exponentially decaying modes. The values of q_d were set so that the corresponding time constants were between $T/5$ and T . Fig. 2(B) shows the time series for five such eigenstates.

For each model dimension considered (see below) we generated a state trajectory using known model parameters generated as described above. We then created an observable data time series $y_n = g(x_n)$ with the observation function $g(x_n) = x_n$, that is, all of the dynamical states are observed.

We then created 'perturbed' parameters by adding Gaussian noise with a standard deviation of 10% of the original parameters. The cost

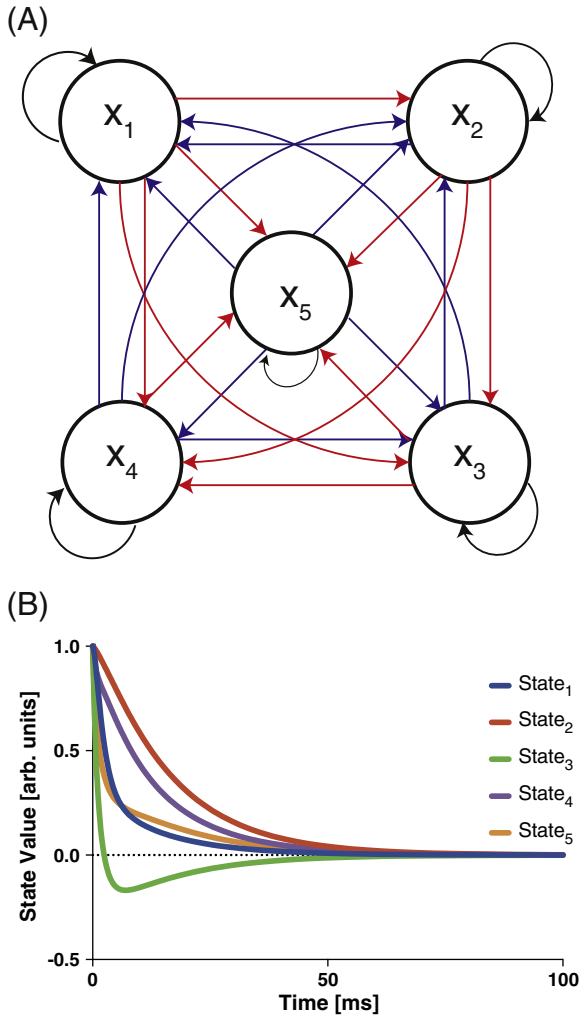


Fig. 2. Linear System (A) The 5-dimensional state-space model and (B) the linear evolution of its eigenstates.

function was defined as

$$J = -\frac{1}{2} \sum_n [y_n - g(x_n)]^2. \quad (27)$$

To summarize, the ‘data points’ y_n were created using the original parameters and the ‘model predictions’, $g(x_n)$ used the perturbed parameters. Gradients were then estimated at this perturbed point.

The systems were integrated using the tolerances of FD and FS fixed at 10^{-7} . Although, the tolerance of AM was adjusted so as to achieve the best fit to the FD based gradient estimate, for the efficiency-scaling simulations we fixed it at a lower value of 10^{-3} . This is illustrated in Fig. 3(A) that shows the estimated gradients for a $D = 5$ dimensional linear system. Setting the tolerance of the AM method to 10^{-3} did not affect the mean-squared deviation of the gradients obtained between the FD and the AM methods (data not shown).

We then compared the three gradient computation methods. Fig. 3(B) plots the computation time as a function of state dimension. For a 28-node system with 784 model parameters the computation time for the adjoint method is 77 times less than for the finite difference method.

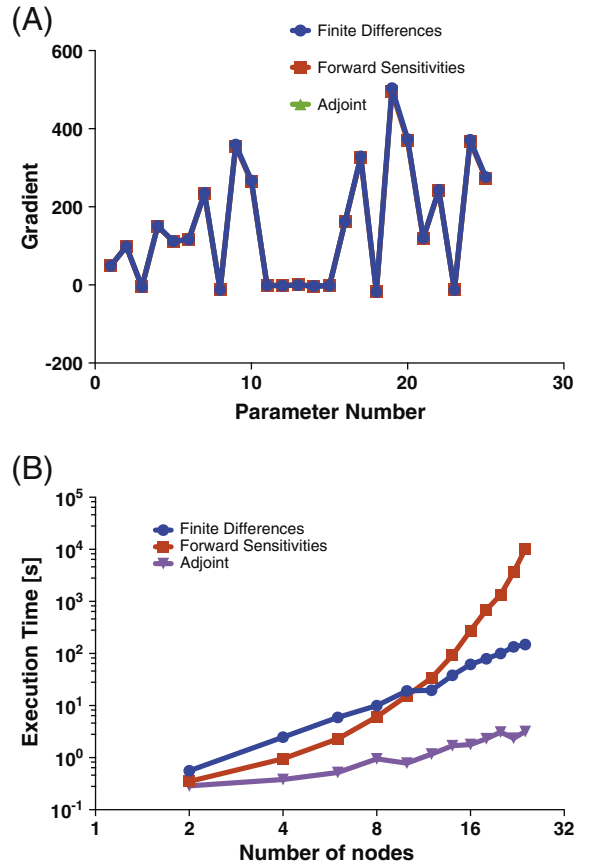


Fig. 3. Computational efficiency for linear systems (A) Comparison of the parametric gradient obtained by the three methods. (B) Scaling of run-time as a function of the number of nodes. The absolute and relative tolerances of FD and FS methods were set to 10^{-7} while the tolerances for the AM method were fixed to 10^{-3} . Simulation time was fixed at 400 ms.

Nonlinear models

Next, we consider weakly coupled oscillators of the form

$$\dot{x}_i(t) = f_i + \sum_{j=1, j \neq i}^D (\alpha_{ij} \sin[x_i(t) - x_j(t)] + \beta_{ij} \cos[x_i(t) - x_j(t)]) \quad (28)$$

where the model parameters comprise the parameters f , α and β . A model with D states therefore has $P = 2D^2 - D$ parameters. We used a cost function equal to the mean square deviation between observed and predicted state trajectories i.e., the norm of the prediction error (again, all states were observed).

The tolerance parameters of the integration process were set identical to those used for the linear models. Again, the adjoint equation being a linear first order ODE enables the use of lower tolerances (10^{-3}). This process was implemented for a $D = 5$ dimensional problem and Fig. 4(B) shows the estimated gradients.

We then compared the three gradient computation methods. Fig. 4(C) plots the computation time as a function of state dimension. For a 24-node system with 1128 model parameters the computation time for the adjoint method is 50 times less than for the finite difference method.

The efficiency of the AM formulation is due to two reasons – first, the adjoint equation is linear and second it is integrated only once to compute the gradient. Given that the AM equation is linear, the condition number is low, enabling any ODE integrator to integrate the adjoint equation with ease. Indeed, if the ODE integrator is subjected to

unnecessary high tolerances it spends more time integrating the adjoint equation. Thus, the advantage of the adjoint scheme reveals both the parsimonious integration scheme as well as the linearity of this equation that requires less-conservative tolerances.

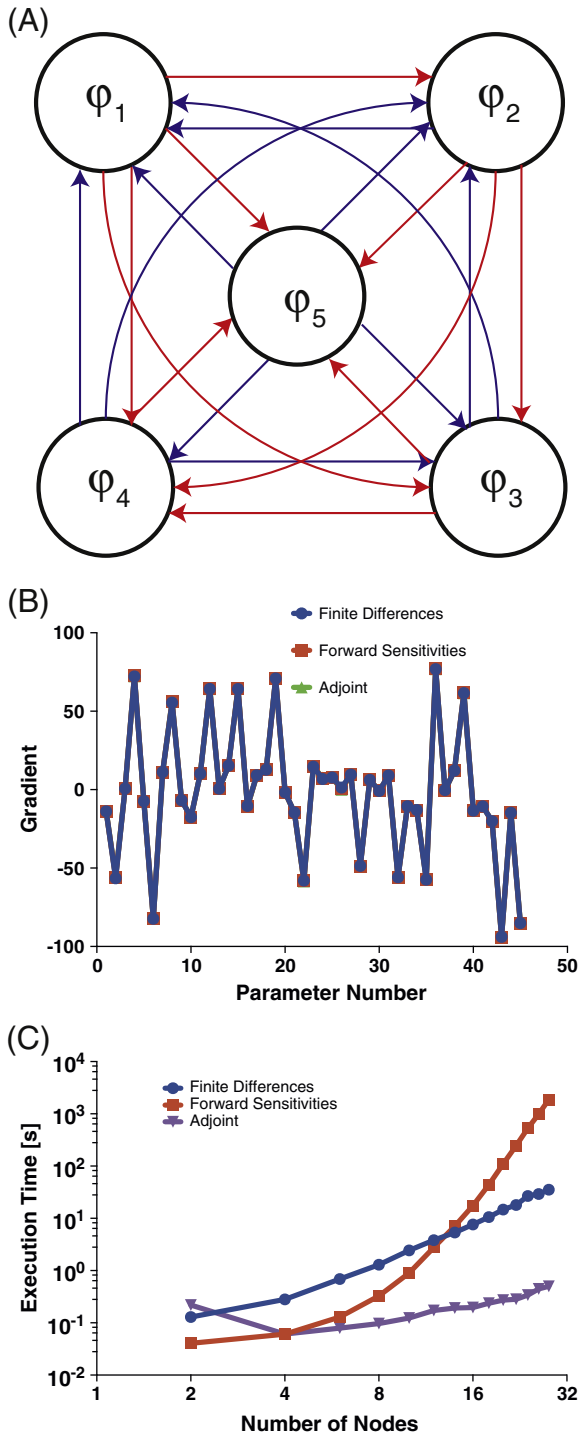


Fig. 4. Computational efficiency for non-linear systems (A) Comparison of the gradient obtained by the three methods. Here, the last five parameters quantify the intrinsic oscillator frequencies, and the first 40 parameters the sine and cosine interaction terms. (B) Scaling of run-time as a function of the number of nodes. The absolute and relative tolerances of FD and FS methods were set to 10^{-7} while the tolerances for the AM method were fixed to 10^{-3} . Simulation time was fixed at 100 ms.

Discussion

Optimization theory attaches mathematical well-posedness to the issue of biological data-assimilation by formalizing the relationship between empirically measured data and a model generating those responses. In this paper, we introduced three different methods for numerical gradient estimation that forms an integral part of any convex optimization framework. Our comparison establishes that the adjoint method is computationally more efficient for numerical estimation of parametric gradients for state-space models – both linear and non-linear, as in the case of a dynamical causal model (DCM). As is apparent from the gradient equations, the adjoint method is efficient when the numbers of parameters are much greater than the number of states determining the cost function. The contrary is true for the Forward Sensitivity approach albeit for large state-space models, finite-difference based gradients prove to be beneficial. There are two remarks that can be made about the adjoint formulation. First, regardless of whether the flow is linear or non-linear the adjoint method requires the integration of a single linear equation – the computational efficiency is inherent in the structure of the equation. Second, the appearance of a transpose on the adjoint vector implies that the flow of information in the system of equations is reversed; it is in this sense that the adjoint equations are integrated backwards in time.

Although, adaptive error correction is invariably used in the integration of differential equations, the numerical simulations suggest that the tolerance used for integrating the flow and adjoint differential equations are vital in determining the accuracy of the parametric gradients, due to the presence of discretization error. In theory, plugging in the solution field to the flow equation should yield zero, but due to the existence of discretization error the residual is generally non-zero. The same is true for the adjoint equation. In fact, a theorem by [Becker and Rannacher \(2001\)](#) shows how discretization of the gradient depends on the average of the errors and the residuals accumulated in the integration of flow and the adjoint equations ([Bangert and Rannacher, 2003](#)). This is also the case for obtaining gradients via finite-differencing, where we find that the fidelity of error-free discretization of the flow equations is a prerequisite for guaranteeing parametric gradients that are a reliable estimate of the true gradient.

It is known that if the flows are prescribed as ODEs the numerical stability of the adjoint system is guaranteed when the adjoint equation is integrated backwards in time, in the sense that the flow is reversed. Our derivation of the adjoint method is mathematically informal so as to illustrate the basic working principle; rigorous mathematical proofs that accommodate higher order differential algebraic equations, time-dependent parameters or objective functionals that depend on initial conditions are available elsewhere ([Cao et al., 2003](#)).

For DCM inversions that allow problem specification in a pre-defined form it may be generally time-efficient to derive the gradient functions analytically rather than using automatic differentiation ([Bischof et al., 2012](#)). Automatic differentiation is particularly important for partial differential equations (PDEs) that have 3-dimensional representations, requiring automatization and therefore proving to be error resilient ([Sala et al., 2004](#)). For a PDE-constrained optimization problem the solution is governed by a fully coupled Karush–Kuhn–Tucker (KKT) system of equations. These can be computationally expensive for parabolic and hyperbolic PDEs, as well as displaying slow convergence of the defined objective functional (ill-conditioning). The adjoint formulation remedies this by decoupling the coupled PDEs and replacing them by iterative solves of a linear adjoint PDE equation. Additional success of adjoint-based gradient methods for PDE-constrained optimization relies on the fact that mesh independent convergence can be attained. Further speedup could also be obtained by using compiled implementation of forward and adjoint sensitivity methods available in the SUNDIALS time integration package ([Hindmarsh and Serban, 2002](#)). This code is written in C and may offer substantial speed advantages over MATLAB implementations.

For data assimilation, it is only rarely that we have precise information on the states or the parameters (Wiener, 1964). Is the adjoint method equally efficient when there is noise on the states and the parameters? One way to represent uncertainty in a mathematical model, whether static or dynamic is to formulate it as a polynomial chaos expansion (Wiener, 1938), one for each noisy state or parameter. This then enables the characteristic statistical quantities to be evaluated as some function of the expansion coefficients – the uncertainty now becomes parameterized. The estimation of the numerical gradient can then proceed akin to a deterministic dynamical model where the computational burden does not depend on the number of parameters (Aleksiev et al., 2010). Alternatively, adjoint methods can be gracefully combined with Markov Chain Monte Carlo (MCMC) sampling-based evaluation of the posterior densities (Ghate and Giles, 2005). In a forthcoming paper we address how second-order adjoint gradient estimates could be obtained in the context of Bayesian inversion of neural masses, mean fields, and neural field equations.

Constrained optimization problems arise in many scientific fields, from neuroscience to financial mathematics, therefore a fundamental need for efficient computational methodologies arises. Our work promotes such an endeavor especially for data-sets arising in neuroscience, for example the inversion of large-scale DCMs that have been routinely used to test hypotheses about different functional brain architectures.

Acknowledgments

This work is supported by a Wellcome Trust/DBT Early Career (500182/Z/11/Z) fellowship to BS. KJF and WDP are supported by the Wellcome Trust (091593/Z/10/Z). Initial code-base used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

References

- Aleksiev, A.K., Navon, I.M., Zelentsov, M.E., 2010. The estimation of functional uncertainty using polynomial chaos and adjoint equations. *Int. J. Numer. Methods Fluids* 67, 3283–341.
- Bangerth, W., Rannacher, R., 2003. *ETH Zürich Lectures in Mathematics*, Chapter Adaptive Finite Element Methods for Differential Equations. Birkhäuser.
- Becker, R., Rannacher, R., 2001. An optimal control approach to a posteriori error estimation in finite element methods. *Acta Numerica* 10, 1102.
- Bischof, C., Martin Bücker, H., Vehreschild, A., Willkomm, J., 2012. Automatic Differentiation for MATLAB (ADiMat). *MATLAB-Day, MATLAB-Day*, Aachen, Germany, October 2012.
- Bishop, C.M., 1995. *Neural Networks for Pattern Recognition*. Oxford University Press, Oxford.
- Bishop, C.M., 2006. *Pattern Recognition and Machine Learning*. Springer, New York.
- Calderhead, B., Girolami, M., 2009. Estimating Bayes factors via thermodynamic integration and population MCMC. *Comput. Stat. Data Anal.* 53 (12), 4028–4045.
- Cao, Y., Li, S., Petzold, L., Serban, R., 2003. Adjoint sensitivity analysis for differential-algebraic equations: the adjoint DAE system and its numerical solution. *SIAM J. Sci. Comput.* 24, 1076–1089.
- Clark, A., 2011. *Geodesic Shooting for Anatomical Curve Registration on the Plane*. (PhD thesis) Department of Aeronautics, Imperial College, London.
- David, O., Kiebel, S., Harrison, L., Mattout, J., Kilner, J., Friston, K., 2006. Dynamic causal modeling of evoked responses in EEG and MEG. *NeuroImage* 30 (4), 1255–1272 (May 2006).
- Deco, G., Jirsa, V.K., Robinson, P.A., Breakspear, M., Friston, K., 2008. The dynamic brain: from spiking neurons to neural masses and cortical fields. *PLoS Comput. Biol.* 4 (8), e1000092.
- Deneux, T., Faugeras, O., 2006. Using nonlinear models in fMRI data analysis: model selection and activation detection. *NeuroImage* 32, 1669–1689.
- Errico, R., 1997. What is an adjoint model? *Bull. Am. Meteorol. Soc.* 78, 2577–2591.
- Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. *NeuroImage* 19 (4), 1273–1302.
- Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., Penny, W., 2007. Variational free energy and the Laplace approximation. *NeuroImage* 34 (1), 220–234 (Jan 2007).
- Gazzaniga, M., 2010. Neuroscience and the correct level of explanation for understanding mind. *Trends Cogn. Sci.* 14, 291–292.
- Ghate, D., Giles, M.B., 2005. Inexpensive Monte Carlo Uncertainty Analysis. Symposium on Applied Aerodynamics and Design of Aerospace Vehicles.
- Giles, M., Pierce, N., 2000. An introduction to the adjoint approach to design. *Flow Turbul. Combust.* 65, 393–415.
- Gronwall, T., 1919. Note on the derivatives with respect to a parameter of the solutions of a system of differential equations. *Ann. Math.* 20, 292–296.
- Hindmarsh, A., Serban, R., 2002. User Documentation for CVODES, and ODE Solver with Sensitivity Analysis Capabilities. Technical report. Centre for Applied Scientific Computing, Lawrence Livermore National Laboratory.
- Hopfield, J., Brody, C., 2001. What is a moment? Transient synchrony as a collective mechanism for spatiotemporal integration. *Proc. Natl. Acad. Sci.* 98 (3), 1282–1287.
- Nocedal, J., Wright, S.J., 2006. *Numerical Optimization*, 2nd edition. Springer, New York.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 1992. *Numerical Recipes in C*. Cambridge University Press, New York.
- Ramsay, J., Hooker, H., Campbell, D., Cao, J., 2007. Parameter estimation for differential equations: a generalized smoothing approach. *J. R. Stat. Soc. Ser. B* 69 (5), 741–796.
- Richtmeyer, D., Morton, K., 1967. *Difference Methods for Initial Value Problems*. Wiley, New York.
- Sala, Marzio, Heroux, Michael A., Day, David M., 2004. Trilinos Tutorial. Technical Report SAND2004-2189. Sandia National Laboratories.
- Shampine, L., Reichelt, M., 1997. The MATLAB ODE Suite. *SIAM J. Sci. Comput.* 18, 1–22.
- Stemmler, M., Sengupta, B., Laughlin, S.B., Niven, J.E., 2012. Energetically Optimal Action Potentials. *Advances in Neural Information Processing Systems*, pp. 1566–1574.
- Vysheirsky, V., Girolami, M., 2008. Bayesian ranking of biochemical system models. *Bioinformatics* 24 (6), 833–839.
- Wang, Q., 2013. Forward and adjoint sensitivity computation of chaotic dynamical systems. *J. Comput. Phys.* 235, 1–13.
- Wiener, Norbert, 1938. The homogeneous chaos. *Am. J. Math.* 60 (4), 897–936 (ISSN 00029327).
- Wiener, Norbert, 1964. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. The MIT Press.
- Wilson, H., 1999. *Spikes, Decisions and Actions: The Dynamical Foundations of Neuroscience*. Oxford University Press, New York.