

"Adjoint Modeling"

A Brief Introduction

Ichiro Fukumori

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California

10 February 2022

© 2022. California Institute of Technology. Government sponsorship acknowledged.

What is this about?

Adjoint models can efficiently compute sensitivity between input and output of models. While popularly known for their use in estimation and optimization (e.g., data assimilation), adjoint models are also increasingly employed to investigate the workings of complex systems like the ocean. What could collectively be called "adjoint modeling", these emerging applications afford insight that can be difficult to attain otherwise.

Here we describe briefly what an adjoint is, how it is useful, and how it is derived.

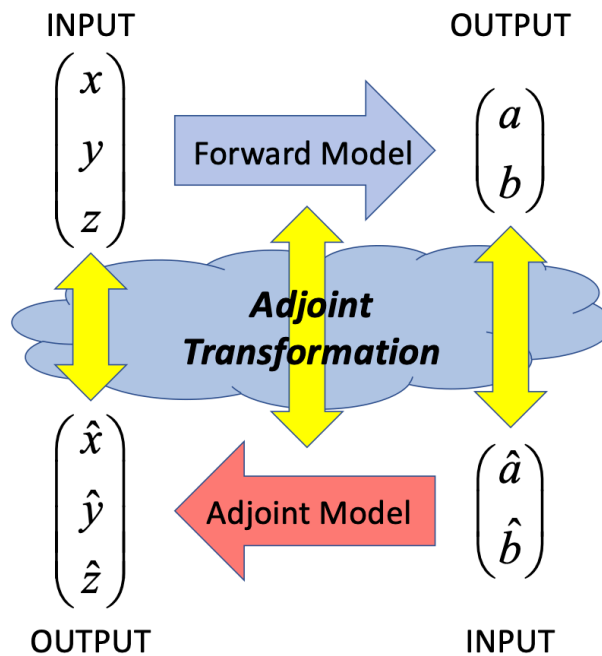
Adjoint is not data assimilation.

What is an adjoint?

Adjoint¹ is a transformation used in studying mathematical relationships. An "adjoint model" refers to an adjoint-transformed version of a model. Here, "model" is simply a set of rules, or a computer program, that takes input and calculates output. To distinguish it from its adjoint, the original pre-transformed model is often referred to as the "forward model" in contrast to its adjoint that computes things "backward", switching roles of what is input and what is output² (see diagram below). Adjoint models are particularly useful for computing how forward model output depends on its input, i.e., *sensitivity* (a.k.a. derivative, gradient).

¹ The adjoint here is technically called the Hermitian Adjoint.

² Adjoint models are distinct from inverse models that are also "backward." Inverse models solve input from output, whereas adjoint models evaluate dependencies between the two.



A schematic illustration of a forward model and its adjoint. Note the input and output being opposite between the two models.

For illustration, suppose we have a simple “forward model” written as

$$\begin{aligned} a &= x - 2y + 3z \\ b &= 4x - 5z \end{aligned} \tag{1}$$

Here, the model's input is the variables (x, y, z) used in the computation on the right-hand-side and its output is the variables (a, b) on the left-hand-side assigned (=) to the results of this computation. Now suppose we have a quantity of interest J (often called “cost function”, “objective function”, or “target function”) defined in terms of the forward model output,

$$J = 6a - 7b \tag{2}$$

and that we are interested in the sensitivity of this quantity to the model input (x, y, z) . In other words, how much would J change, if $x, y,$ or z were to change by one?

One way of evaluating this sensitivity is to use the forward model, by changing its input $x, y,$ and z one by one and evaluating how J would change each time. For example, J 's sensitivity to x can be computed by setting $x=1$ and $y=z=0$ as input for the forward model (Eq 1) and using its output to evaluate J (Eq 2), which result is -22 . J 's sensitivity to y and z can be obtained similarly, each requiring a separate evaluation of the forward model. Thus, this way of deriving the sensitivity, the answer of which is $(-22, -12, 53)$, requires three separate evaluations of the forward model, one for each input variable.

The “adjoint model”, in comparison, allows us to evaluate this sensitivity in one step. The adjoint of the forward model (Eq 1) is³

$$\begin{aligned}\hat{x} &= \hat{a} + 4\hat{b} \\ \hat{y} &= -2\hat{a} \\ \hat{z} &= 3\hat{a} - 5\hat{b}\end{aligned}\tag{3}$$

where carets (^) denote variables of the adjoint model (“adjoint variables”). Here, the adjoint model’s input used in the computation on the right-hand-side is (\hat{a}, \hat{b}) and its output assigned on the left-hand-side is $(\hat{x}, \hat{y}, \hat{z})$, which are opposite to what they are in the forward model. That is, the adjoint model conducts its computation “backward” compared to the “forward model” as noted above.

Adjoint variables can be regarded as sensitivity. Specifically, by setting the adjoint model's input (\hat{a}, \hat{b}) to be the sensitivity of J to (a, b) (Eq 2),

$$\begin{aligned}\hat{a} &= 6 \\ \hat{b} &= -7\end{aligned}\tag{4}$$

the adjoint model's output $(\hat{x}, \hat{y}, \hat{z})$ gives the sensitivity of J to (x, y, z) . The result, $(\hat{x}, \hat{y}, \hat{z}) = (-22, -12, 53)$, is the same as the previous calculation. However, the present computation requires only one evaluation of the adjoint model (Eq 3) instead of the three that were required by the forward model (Eq 1).

For this example, owing to its simplicity, the difference is small between using the forward model and using the adjoint model in the amount of computation required to evaluate the sensitivity. But imagine the model being more complex such as the ocean general circulation model used by the Consortium for "Estimating the Circulation and Climate of the Ocean" (ECCO; <https://ecco-group.org>; Wunsch et al., 2009). In ECCO, a single forward or adjoint model evaluation requires several days of computation on a state-of-the-art supercomputer with millions of input and output variables. The sensitivity calculation using the forward model will

³ In many circumstances, as it is here, the adjoint corresponds to a transpose operation. This can be seen more clearly in matrix notation. Namely, using matrices, Equations (1) and (3) can be written, respectively, as,

$$\begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} 1 & -2 & 3 \\ 4 & 0 & -5 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix}\tag{1'}$$

$$\begin{pmatrix} \hat{x} \\ \hat{y} \\ \hat{z} \end{pmatrix} = \begin{pmatrix} 1 & 4 \\ -2 & 0 \\ 3 & -5 \end{pmatrix} \begin{pmatrix} \hat{a} \\ \hat{b} \end{pmatrix}\tag{3'}$$

Note the coefficient matrix on the right-hand-side of the adjoint model (Eq 3’) simply being the transpose of that in the forward model (Eq 1’), i.e., a reversal of rows and columns of the matrix. Sensitivity is independent of inhomogeneous terms and are, therefore, excluded from the example. For nonlinear systems, adjoints can be defined with respect to linearized forward models (i.e., tangent linear approximation).

require one forward model evaluation per input variable which altogether will take millions of days to complete (that's a couple thousand years!), but the equivalent calculation can be performed with only one evaluation of the adjoint model taking only a few days' time. Thus, it is impractical to use the forward model to evaluate the sensitivity of a model like ECCO's, whereas it is a routine calculation using its adjoint. This computational efficiency is what makes an adjoint model indispensable.

Adjoint models efficiently compute sensitivity.

How is an adjoint model useful?

Sensitivity (derivative) is foundational in calculus used in every branch of science and engineering that employs math. Just as how calculus has advanced these disciplines, adjoint-derived sensitivities provide new insight into problems.

Adjoint models were first introduced in oceanography in the context of data assimilation (*Thacker and Long, 1988*). Based on sensitivity of model-data differences, adjoint models are used to fit corresponding forward models to observations. Such approach ("adjoint method") is also at the heart of ECCO's Central Solution and most of its other estimation products (e.g., *Forget et al., 2015*). Values of an adjoint's sensitivity also tell us what observations are most effective for monitoring different quantities of interest and are useful in designing observing systems (*Köhl and Stammer, 2004*).

Importantly, adjoints also provide an effective means to investigate the workings of complex systems like the ocean. As a tangible example, consider a change in passive tracer, like dye in water, which can tell us how a tracer-tagged water spreads with time (fate of the water). In contrast, the sensitivity of such quantity to tracer concentrations in the past ("adjoint passive tracer"; *Fukumori et al., 2004*), tells us where that tracer-tagged water came from (origin of the water). Quantitative measures of origin and fate elucidate pathways of ocean circulation and provide insight into processes controlling its evolution.

Adjoint models are also employed in directly studying causation ("adjoint gradient decomposition"; *Fukumori et al., 2015*); in this, sensitivity is used to quantify effects of different elements driving the ocean, thus providing a tool to assess their relative contribution (e.g., wind vs heat flux). In contrast to commonly-used correlation, which quantifies similarities, adjoints reveal causation based on theoretical relationships underlying the models (first principles).

Adjoint models quantify causation that correlation doesn't.

How do you derive an adjoint?

Derivation of adjoint models is conceptionally straightforward but technically challenging owing to the size and complexities of computer programs defining most state-of-the-art forward models. Remarkably, however, there are software that can automatically transform computer programs into their adjoint (“automatic differentiation”, “algorithmic differentiation”, “AD tools”; *International Conference on Automatic Differentiation, 2004, 2012*). For example, the ocean circulation model used by ECCO, MITgcm (*Marshall et al., 1997*), was purposefully created to use such automatic differentiation tool to generate its adjoint (*Marotzke et al., 1999*). Today, MITgcm remains one of the few state-of-the-art ocean general circulation models that has an adjoint readily available (<https://github.com/MITgcm/MITgcm.git>). A rigorous process is in place to routinely test the “adjointability” of new features introduced in the model, thus assuring availability of the adjoint for the latest version of MITgcm for use in ocean state estimation and other applications.

AD tools help generate adjoint models.

Concluding remark

Despite its foundational nature, adjoint modeling is still far uncommon today compared to forward modeling. This difference is due in part to most adjoint models being developed for purposes of estimation and optimization, which relatively speaking are specialized fields of investigation, as opposed to general modeling applications. As the community becomes more familiar with adjoint models, their availability and scope of utilization are bound to broaden synergistically. Adjoint modeling is still at a nascent stage and is ripe for innovation.

Adjoint modeling is on the cusp of transformation.

Acknowledgement

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (NASA).

References and recommended reading

- Errico, R. M., 1997: What is an adjoint model? *Bull. Am. Meteorol. Soc.*, **78**, 2577–2591, [https://doi.org/10.1175/1520-0477\(1997\)078<2577:wiaam>2.0.co;2](https://doi.org/10.1175/1520-0477(1997)078<2577:wiaam>2.0.co;2).
- Forget, G., J. M. Campin, P. Heimbach, C. N. Hill, R. M. Ponte, and C. Wunsch, 2015: ECCO version 4: an integrated framework for non-linear inverse modeling and global ocean state

- estimation. *Geosci. Model Dev. Discuss.*, **8**, 3653–3743, <https://doi.org/10.5194/gmdd-8-3653-2015>.
- Fukumori, I., T. Lee, B. Cheng, and D. Menemenlis, 2004: The origin, pathway, and destination of Niño-3 water estimated by a simulated passive tracer and its adjoint. *J. Phys. Oceanogr.*, **34**, 582–604, <https://doi.org/10.1175/2515.1>.
- Fukumori, I., O. Wang, W. Llovel, I. Fenty, and G. Forget, 2015: A near-uniform fluctuation of ocean bottom pressure and sea level across the deep ocean basins of the Arctic Ocean and the Nordic Seas. *Prog. Oceanogr.*, **134**, 152–172, <https://doi.org/10.1016/j.pocean.2015.01.013>.
- Giering, R., and T. Kaminski, 1998: Recipes for adjoint code construction. *Acm Trans. Math. Softw.*, **24**, 437–474, <https://doi.org/10.1145/293686.293695>.
- International Conference on Automatic Differentiation (4th : 2004 : Chicago, Ill.), *Automatic Differentiation : Applications, Theory and Implementations*. Berlin: Springer; 2006. doi:10.1007/3-540-28438-9
- International Conference on Automatic Differentiation (6th : 2012 : Fort Collins, CO.), *Recent Advances in Algorithmic Differentiation*. Berlin: Springer; 2012. doi:10.1007/978-3-642-30023-3
- Köhl, A., and D. Stammer, 2004: Optimal observations for variational data assimilation. *J. Phys. Oceanogr.*, **34**, 529–542, <https://doi.org/10.1175/2513.1>.
- Marotzke, J., R. Giering, K. Q. Zhang, D. Stammer, C. Hill, and T. Lee, 1999: Construction of the adjoint MIT ocean general circulation model and application to Atlantic heat transport sensitivity. *J. Geophys. Res.*, **104**, 29529–29547, <https://doi.org/10.1029/1999jc900236>.
- Marshall, J., A. Adcroft, C. Hill, L. Perelman, and C. Heisey, 1997: A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers. *J. Geophys. Res.*, **102**, 5753–5766, <https://doi.org/10.1029/96jc02775>.
- Thacker, W. C., and R. B. Long, 1988: Fitting dynamics to data. *J. Geophys. Res. Ocean.*, **93**, 1227–1240, <https://doi.org/10.1029/JC093iC02p01227>.
- Wunsch, C., 2006: *Discrete Inverse and State Estimation Problems: With Geophysical Fluid Applications*. Cambridge University Press, 371 pp.
- Wunsch, C., P. Heimbach, R. M. Ponte, and I. Fukumori, 2009: The global general circulation of the ocean estimated by the ECCO-Consortium. *Oceanography*, **22**, 88–103, <https://doi.org/10.5670/oceanog.2009.41>.