Lecture 15: Generalized matrix inversion and the singular value decomposition

Recap

We've been using the singular value decomposition (SVD) to extract the leading modes of variability in a data set, using an empirical orthogonal function analysis. We represented our matrix $\mathbf{G}$ as

$$\mathbf{G} = \mathbf{USV}^T, \tag{1}$$

and we noted that the leading order variability in the data could be represented by the leading column vecotors of the orthogonal matrices $\mathbf{U}$ and $\mathbf{V}$, scaled by the singular values. Today we're going to return to one of the primary uses of the SVD, which is to invert a matrix.

To do this, we'll examine a classic oceanographic problem, determining reference velocities for to reference geostrophic velocity determined from the thermal wind equations.

Determination of the geostrophic reference velocity using a control volume

Consider the time-honored oceanographic problem of determining a reference velocity for a geostrophic velocity calculation. Suppose hydrographic stations were made in the form of sections that enclose a volume of water (Figure 1). Define a coordinate system where x is horizontal distance along the section enclosing the control volume, and z is vertical. The section can then be unfolded and visualized in the x–z plane (Figure 2). Using the thermal wind equation, the geostrophic velocity is

$$v(x, z) = v_0(x) - \frac{g}{f \rho_0} \int_{z_0}^{z} \frac{\partial \rho}{\partial x}(x, z') \, dz' \tag{2}$$

Here v_0 is the reference velocity normal to the section at depth z_0, ρ is density, g is gravitational acceleration, f is the Coriolis parameter, and ρ_0 is reference density. In a classic oceanographic scenario, we measure density as a function of x and z, but we do not know the reference veocity v_0, so our challenge will be to find the best possible estimate of v_0. We also want to impose a set of constraints on the system:

1. The system is in geostrophic balance.

2. Mass is conserved in the closed domain.

3. Water properties are conserved in the closed domain. For example, there is no net heat gain within the box.

4. The system is in steady state.

5. The level of no motion (or alternatively "known motion") is relatively siple, and reference velocities do not oscillate wildly between adjacent stations.

When this question first arose, decades ago, oceanographers questioned whether all of these constraints could be met simultaneously. Work guided by Carl Wunsch and his students and collaborators formulated a framework for finding an optimal solution based on these constraints.

We start by discretize for horizontal location m, and vertical location n. Velocity is then

$$v_{nm} = v_{0m} + v'_{nm} \tag{3}$$

where v_{0m} is the unknown depth-independent reference velocity, and v'_{nm} is the known depth-dependent relative velocity from thermal wind. Conservation of a property C_{nm} in a vertical range Δz_{nm} implies

$$\sum_{m=1}^{M} C_{nm}(v_{0m} + v'_{nm})\Delta z_{nm}\Delta x_m = 0 \tag{4}$$

We suppose there are N such conservation statements for different properties and layers. Then the components of the data vector, model parameter vector, and data kernel matrix are

$$
\begin{aligned}
d_n &= \sum_{m=1}^{M} C_{nm}v'_{nm}\Delta z_{nm}\Delta x_m & (5)\\
m_m &= v_{0m} & (6)\\
G_{nm} &= C_{nm}\Delta z_{nm}\Delta x_{nm} & (7)
\end{aligned}
$$

We can write conservation statements for mass ($C_{nm} = 1$), for temperature or heat, for salt, for O_2, for nutrients and for potential vorticity. We can also require conservation in an arbitrarily large number of potential density layers. Thus you might suppose that you could set this problem up to be formally overdetermined so that you could solve for the reference velocities. Unfortunately, this is easier said than done, for several reasons. First, if you force the problem to have more equations than unknowns, the rows of G might turn out to be highly correlated, perhaps because the different variables are highly correlated (e.g. variations in temperature mirror variations in salinity), or perhaps because adjacent density layers are highly correlated. This will lead to an ill-conditioned matrix that will not invert. You could run into further problems because you might conclude that properties are not completely conserved within density layers and that you need to account for an unknown background diffusivity (or perhaps a vertical advection term) that connects adjacent layers.

In essence, there are usually more stations in a section (and more unknown reference velocities v_o) than sensible conservation statements. We'll start by treating this as an **underdetermined** problem. It might be argued that all oceanographic inverse problems are underdetermined, as we never have enough data to determine the complete state of the ocean.
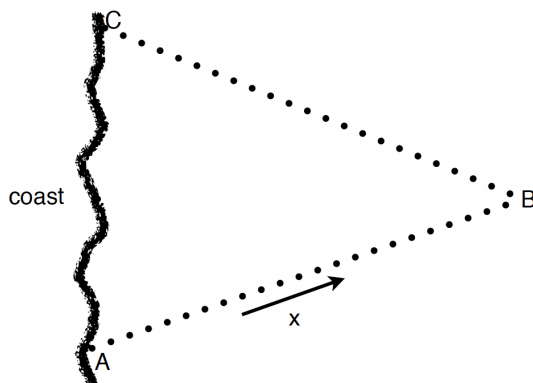

Figure 1: Hydrographic sections extending from the coast and enclosing a volume of water.

Solution 1: Treat as an underdetermined problem

If we have more unknowns than equations, we can treat this as an underdetermined problem. We talked about this earlier, and we'll refresh details now. (See Dan Rudnick's notes entitled "reference_velocity.pdf" for a continuous version of this problem.) The overall size of the reference velocities are measured by the size of $\mathbf{m}$:

$$\mathbf{m}^T \mathbf{m} \tag{8}$$

In discrete form, we can define a cost function that minimizes the size of the reference velocities subject to the model requirements:

$$\mathcal{L} = \sum | = \infty^{\mathcal{M}} \mathbb{1}_|^\in \cdot \S\rangle - \in \sum_{\rangle = \infty}^{\mathcal{N}} \lambda_\rangle \left(\sum_{|=\infty}^{\mathcal{M}} \mathcal{G}_{\backslash \mathbb{1}} \mathbb{1}_| \cdot \S_| - \lceil_\rangle \right) \tag{9}$$

Minimizing (9) is a problem in variational calculus. Taking the variation with respect to m, and setting the result to zero yields

$$\delta \mathbf{L} = \sum_{|=\infty}^{\mathcal{M}} \left(\in \mathbb{1}_| - \in \sum_{\rangle=\infty}^{\mathcal{N}} \lambda_\rangle \mathcal{G}_{\rangle|} \right) \delta \mathbb{1}_| \lceil \S = \prime \tag{10}$$

Thus, to reach a minimum, the expression in parentheses must be zero:

$$m_j = \sum_{i=1}^{N} \lambda_i G_{ij} \tag{11}$$

So the smallest model is a linear combination of layer thicknesses. If $N = 1$, that is we conserve mass in only one layer, then the best model is simply a scaling of the layer thickness. Substitute (11) into the constraints (7) to find the Lagrange multipliers λ_i.

$$\sum_{j=1}^{M} \sum_{k=1}^{N} G_{ij} G_{ik} \lambda_j \, dx = d_i \tag{12}$$
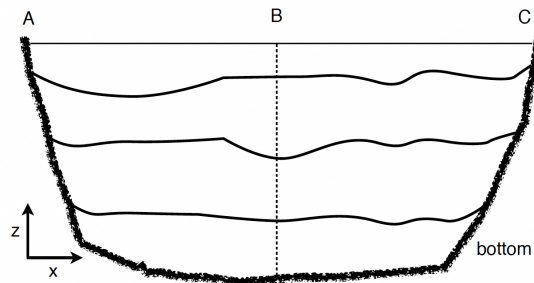

Figure 2: The section of Figure1, unfolded as a function of depth and distance.

Define the $N \times N$ matrix $\mathbf{H}$ whose elements are the inner products of the layer thicknesses

$$H_{ij} = \sum_{k=1}^{N} G_{ki} G_{kj} \tag{13}$$

or in matrix form
$$\mathbf{H} = \mathbf{GG}^T \tag{14}$$

Then (12) can be written as the matrix equation:

$$\mathbf{Ha} = \mathbf{d} \tag{15}$$

where $\mathbf{a}$ is the vector of Lagrange multipliers λ_i and $\mathbf{d}$ is the vector of data d_i. The Lagrange multipliers are found by solving (15)

$$\mathbf{a} = \mathbf{H}^{-1}\mathbf{d} \tag{16}$$

Combining (16) with (11) is the solution we seek.

$$\mathbf{m} = \mathbf{G}^T(\mathbf{GG}^T)^{-1}\mathbf{d}, \tag{17}$$

which is referred to as the Moore-Penrose inverse.

In setting up this problem, we choose layer depths consistent with observed potential temperature, or potential density. So, what would be the properties of a good set of layers in the sense that they improve the determination of the reference velocity? If the layer thicknesses are orthogonal, then $\mathbf{H}$ would be diagonal, and all N constraints would provide useful information. On the other hand if some of the layer thicknesses are similar, then $\mathbf{H}$ would have small eigenvalues, and the solution would be unstable with large values in the regions where there are small differences between layer thicknesses.

Referring back to the formulation of this problem lecture, the matrix $\mathbf{H}$ is approximated by $\mathbf{GG}^T$. The singular values of $\mathbf{G}$ correspond to the square roots of the eigenvalues of $\mathbf{H}$. The simplest form of this problem considers only mass conservation within each layer. When we include tracers in this problem, assuming that heat and salt are conserved within layers, how does the solution improve? For a tracer c_i, the conservation statement is expressed in (5).

If C_{nm} is a constant, that is it does not vary along a section, then adding the tracer would add no information. To add information, the tracer should be linearly independent of the layer thicknesses. An "ideal" tracer would be one for which:

$$\sum_{i=1}^{N} C_{ij}G_{ij}G_{ik} = 0 \tag{18}$$

for all combinations of j and k.

In mathematical terms, this is the reason why oceanographers spend so much time looking at tracers.

The null space of our inverse problem consists of all vectors $\mathbf{q}$ for which

$$\sum_{i=1}^{N} q_i G_{ij} = 0 \tag{19}$$

for all i. Any solution of the form

$$\hat{\mathbf{m}} = \mathbf{m} + \mathbf{q} \tag{20}$$

would satisfy the constraints (5-7), but would be larger than $\mathbf{m}$ by the measure defined by (8).

When we solve this problem, we build the station spacing (the width of each cell, Δx_i) into the matrix $\mathbf{G}$. That could influence the relative the final solution, since our minimizationof $\mathbf{m}^T\mathbf{m}$ places equal weight on eqch velocity, while our solution to $\mathbf{Gm} - \mathbf{d}$ will assign less cost to narrower boxes. We could tune this by defining a weight matrix $\mathbf{W}_m$ to account for varying column

widths, so that we minimized the width-weighted reference velocity $\mathbf{m}' = \mathbf{W}_m^{1/2}\mathbf{m}$. In this case, the column-weighted solution becomes

$$\mathbf{m} = \mathbf{W}_m^{-1/2}(\mathbf{W}_m^{-1/2}\mathbf{G}^T)(\mathbf{G}\mathbf{W}_m^{-1}\mathbf{G}^T)^{-1}\mathbf{d}. \tag{21}$$

Solution 2: Row weight and column weight

Once we introduce to this problem the possibility of column weighting by data uncertainty as well as row weighting, you might conclude that we should be able to manage the overdetermined version of this problem, using the column weighting to prefernt the matrix from being singular. This could potentially help us out. In this case, we'd look for a solution of the form:

$$\mathbf{m} = (\mathbf{G}^T\mathbf{W}_e^{-1}\mathbf{G} + \lambda\mathbf{W}_m^{-1})^{-1}\mathbf{G}^T\mathbf{W}_e^{-1}\mathbf{d} \tag{22}$$

In this case adding $\mathbf{W}_m^{-1}$ to the diagonal of the matrix that we invert will stabilize the inversion. Low uncertainties or small covariances imply large weights, which will stabilize the inversion (and force it to deviate less from a prior guess of zero).

Solution 3: Solve using SVD

Our final strategy for this inversion problem is to compute the singular value decomposition. In this case

$$\mathbf{G} = \mathbf{U}\mathbf{S}\mathbf{V}^T, \tag{23}$$

and the solution becomes

$$\mathbf{m} = \sum_{l=1}^{K} \alpha_l\mathbf{v}_l = \sum_{l=1}^{K} \frac{\mathbf{u}_l^T\mathbf{d}}{s_l}\mathbf{v}_l, \tag{24}$$

where α_l are coefficients for the lth singular vector $\mathbf{v}_l$, taken from the lth column of the orthonormal matrix $\mathbf{V}$. Likewise $\mathbf{u}_l$ represents the lth singular vector taken as the lth column of the orthonormal matrix $\mathbf{U}$. The singular value s_l is the lth value from the diagonal of $\mathbf{S}$.

This solution is sensitive to the number of singular values that we decide to use. If the matrix $\mathbf{G}$ is full rank, we would have $K = M$, where M is the number of columns in $\mathbf{G}$. But assuming the matrix is singular or near singular, we should truncate the solution to avoid using the zero or near-zero singualr values. In matrix form the equation becomes

$$\mathbf{m} = \mathbf{V}_K\mathbf{S}_K^{-1}\mathbf{U}_K^T\mathbf{d} \tag{25}$$

By choosing K we control the size of the solution and the size of the null space. Smaller values of K imply smaller $\mathbf{m}$—smaller reference velocities. And the $K + 1$ to M or $K + 1$ to N define the null space.

Synopsis: How to infer reference velocities

Here's a quick synopsis of the approach to inverting to find reference velocieies.

1. Use the measured data to set up a system of equations of the form $\mathbf{G}\mathbf{m} = \mathbf{d}$. Guess the unknown initial reference velocity. You will solve for anomalies relative to this initial guess.

2. When you set up the system of equations, sort the data by isopycnal (or isotherm) and assume conservation of mass, heat, salt, and other properties within each layer.

3. Column weight the matrix $\mathbf{G}$ by height and width. Row weight based on data uncertainties.

4. Solve for reference velocities. Since the problem is intrinsically underdetermined, constrain the solution to be as small as possible.

Resolution matrices

A sensible question given an inverse solution is how well the data and the model parameters are resolved. This question is addressed through the resolution matrices.

Data resolution matrix

Assume that an estimate of the model parameter vector $\mathbf{m}$ is

$$\hat{\mathbf{m}} = \mathbf{G}^{-g}\mathbf{d} \tag{26}$$

So, how well does our estimated data $\hat{\mathbf{d}}$ match our measured data $\mathbf{d}$? Answer this question as follows:

$$\hat{\mathbf{d}} = \mathbf{G}\hat{\mathbf{m}} = \mathbf{G}\mathbf{G}^{-g}\mathbf{d} \tag{27}$$

We define an $N \times N$ matrix

$$\mathbf{N} = \mathbf{G}\mathbf{G}^{-g}, \tag{28}$$

which is called the data resolution matrix. If $\mathbf{N} = \mathbf{I}$ then the predicted and measured data are equal. If the data possess a natural ordering, the rows of $\mathbf{N}$ are an approximation to a delta function (Figure 3). That is, the rows tell how well the data are resolved. An identity matrix infers perfect resolution. Off diagonal elements imply smoothing.
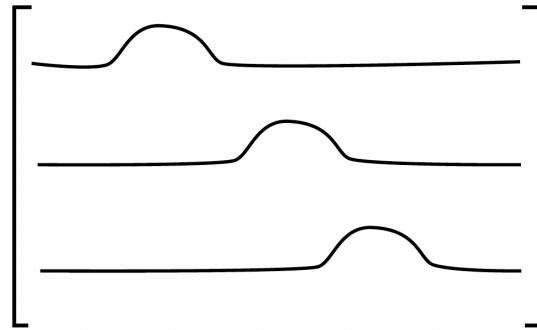

Figure 3: The data resolution matrix represented graphically. The rows are the approximation to the delta function implied by the inverse solution.

Consider the pseudoinverse

$$\mathbf{G}^{+} = \mathbf{V}\mathbf{S}^{+}\mathbf{U}^{T} \tag{29}$$

The data resolution matrix in this case is

$$
\begin{aligned}
\mathbf{N} &= \mathbf{G}\mathbf{G}^{+} & (30)\\
&= \mathbf{U}\mathbf{S}\mathbf{V}^{T}\mathbf{V}\mathbf{S}^{+}\mathbf{U} & (31)\\
&= \mathbf{U}\mathbf{S}\mathbf{S}^{+}\mathbf{U}^{T} & (32)\\
&= \mathbf{U}\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}\mathbf{U}^{T} & (33)\\
&= \mathbf{U}_{K}\mathbf{U}_{K}^{T} & (34)
\end{aligned}
$$

where $\mathbf{U}_K$ is the $N \times K$ matrix made of the first K columns of $\mathbf{U}$, and K is the rank of $\mathbf{G}$.

Model resolution matrix

How well does our estimate of the model parameters $\hat{\mathbf{m}}$ match the true model parameters $\mathbf{m}_0$? Address this question as follows

$$\hat{\mathbf{m}} = \mathbf{G}^{-g}\mathbf{d} = \mathbf{G}^{-g}\mathbf{G}\mathbf{m}_0 \tag{35}$$

The $M \times M$ matrix

$$\mathbf{R} = \mathbf{G}^{-g}\mathbf{G} \tag{36}$$

is called the model resolution matrix. The rows of $\mathbf{R}$ are approximations to the delta function, in direct analogy with the data resolution matrix. For the pseudoinverse

$$\mathbf{R} = \mathbf{G}^{+}\mathbf{G} = \mathbf{V}_K\mathbf{V}_K^{T} \tag{37}$$

where $\mathbf{V}_K$ is the $M \times K$ matrix made of the first K columns of $\mathbf{V}$, and K is the rank of $\mathbf{G}$.

Having good resolution of data and model parameters is a reasonable goal for an inverse. In fact, one can optimize resolution, by minimizing some measure of the difference between $\mathbf{N}$, $\mathbf{R}$, and the identity matrix. If a suitable measure is chosen, the solutions are identical to some we have already derived by minimizing L_2 norms of misfit and model size.

Bibliography

Wunsch, C. 1996. *The Ocean Circulation Inverse Problem*, Cambridge University Press, New York, 442 pp., https://doi.org/10.1017/CBO9780511629570.