The Model Manifold: Four Subspaces in Regression

Companion notes for the "Model Manifold Subspaces" slideshow. The video unfolds in five movements, mirroring the four fundamental subspaces of a regression matrix and what each one means for a model with redundant parameters. Light click markers help you follow along either while presenting or while reading.

Movement 1: Setting up the model map

We fit three noisy data points at $t = 0.3, 1.0, 1.7$ — observed values $(1.58,\ 0.54,\ 2.91)$ — with a fancier model than a straight line:

$$y = \theta_1 x + \theta_2 x^2 + \theta_3\, x(x+1)$$

built from three basis functions, $\phi_1(x) = x$, $\phi_2(x) = x^2$, $\phi_3(x) = x(x+1)$. Evaluating each basis function at each t_i builds the familiar matrix-vector equation:

$$\begin{bmatrix} 0.30 & 0.09 & 0.39 \\ 1.00 & 1.00 & 2.00 \\ 1.70 & 2.89 & 4.59 \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \end{bmatrix} = \begin{bmatrix} 1.58 \\ 0.54 \\ 2.91 \end{bmatrix}$$

With the model matrix in hand, we again fill parameter space with a rainbow cloud of $(\theta_1, \theta_2, \theta_3)$ guesses and map them over to data space, with each guess's prediction curve drawn faintly in the y-vs-t view. Even though we now have three free parameters — one more than data points — the cloud stubbornly refuses to fill the full 3-D data space. Something is redundant.

Movement 2: The four subspaces

Shrinking the y-vs-t view away and growing the parameter-space and data-space panels to fill the screen (rotating each as new content appears, the same technique used in the earlier spaces-and-transformations video) lays the answer bare. The rainbow cloud in data space collapses onto a 2-D **range** plane — not the full 3-D space — spanned by the model matrix's own first two columns. Correspondingly, parameter space has a genuine **null space**: a whole line of parameter combinations, direction $(1, 1, -1)$, that all predict exactly the same thing.

Turning to A^T: the **row space** is a plane in parameter space spanned by $(1, 0, 1)$ and $(0, 1, 1)$, and the **left null space** is a line in data space, direction roughly $(0.85, -0.51, 0.15)$ — orthogonal to the range plane, just as the null space is orthogonal to the row space.

Movement 3: The null space is what you *can't* learn

Bringing the y-vs-t view back with the data and best-fit line, we write the null-space direction explicitly as a bracketed vector, $(1, 1, -1)$, and add $\theta_{\text{fit}} + \alpha(1, 1, -1)$ with a single slider on α. Sliding α back and forth visibly drags the three parameters all over the place — yet the best-fit line and the predicted point in data space never move. Moving along the null space changes the parameters without changing a single prediction: there is simply no way to learn *which* point along that line is "correct" from this data.

Movement 4: The row space is what you *can* learn — and the left null space is what the model misses

Repeating the exercise with the row-space directions $(1, 0, 1)$ and $(0, 1, 1)$ (sliders α, β) tells the opposite story: moving around in the row space *does* move the predicted point and the best-fit line. These are the identifiable directions — the combinations of parameters the data can actually pin down.

Finally, sliding the *observed* data point itself along the left null-space direction (slider γ) shows it always projecting onto the same point on the model manifold — the best-fit line never changes. The left null space is exactly the part of the data the model has no way to explain: not a parameter ambiguity, but genuine model mismatch.

Movement 5: Why — the redundant basis function

The culprit is hiding in the model itself. Since $x(x + 1) = x^2 + x$, the model

$$y = \theta_1 x + \theta_2 x^2 + \theta_3\, x(x + 1)$$

is "double-dipping" — it can be rewritten as

$$y = (\theta_1 + \theta_3)\, x + (\theta_2 + \theta_3)\, x^2$$

We can never learn $\theta_1, \theta_2, \theta_3$ individually — only the two combinations $\theta_1 + \theta_3$ and $\theta_2 + \theta_3$. Those are exactly the row-space directions $(1, 0, 1)$ and $(0, 1, 1)$: the *identifiable* combinations. The one combination we can never pin down is $\theta_1 + \theta_2 - \theta_3$ — the null-space direction.

Putting the four subspaces side by side:

Subspace	Meaning here
$R(A)$	the model manifold — every prediction the model can reach
$N(A)$	unidentifiable parameter combinations
$R(A^T)$	identifiable parameter combinations
$N(A^T)$	model mismatch — what's true but structurally uncapturable

As both panels slowly spin through a full turn, these four labels close out the video: the same four-subspace structure from a simple 2-D map, now shown to be exactly what governs which parts of a regression model can and can't be learned from data.