

2023-08-29

1 Bias-variance tradeoffs and regularization

At the end of the last class, we talked about generalization error for least squares in the presence of noise. In particular, we saw that if there is a true linear model $Ax = b$ and we compute an estimate $A_1\hat{x} = \hat{b}_1$ involving a subset of the data with noisy measurements ($\hat{b}_1 = b_1 + e$), then

$$\|Ax - A\hat{x}\| \leq \|AA_1^\dagger\| \|e\| \leq \|A\| \|A_1^\dagger\| \|e\|.$$

We cannot generally overcome the effects of the measurement error, but we can at least hope not to amplify them. The amplification factor $\|A\| \|A_1^\dagger\|$ could be large if $\|A\|$ is large, but a more common problem is that $\|A_1^\dagger\|$ is large. In numerical analysis, we might call this a problem with *ill-conditioning*; in statistics, we might refer to this as an issue with *high variance* in the estimator.

More generally, suppose we have $y = f(x) + \epsilon$ where ϵ is a noise term with mean zero and variance σ^2 , and we use data y to fit an estimator function $\hat{f}(x)$. Then for an unseen point x

$$\mathbb{E}[(y - \hat{f}(x))^2] = \mathbb{E}[\hat{f}(x) - f(x)]^2 + \text{Var}[\hat{f}(x)] + \sigma^2;$$

that is, the squared error involves the squared *bias* (caused by inability of the model to fit the function even with perfect data); the *variance* (associated with stability of the estimation procedure under perturbations from noise); and a term associated with measurement error.

If we have an ill-conditioned fitting problem (or a high variance estimator, if you prefer), what should we do? There are three options:

1. Get more data.
2. Reduce the noise.
3. Change the fitting problem to reduce the ill-conditioning.

Getting more data or reducing the noise is not always practical or economical, and in any event is an issue unrelated to numerical methods. Thus, we will turn now to the last approach: *regularizing* the problem to reduce the ill-conditioning, possibly at the expense of a small bias.

2 Factor selection and pivoted QR

In ill-conditioned problems, the columns of A are nearly linearly dependent; we can effectively predict some columns as linear combinations of other columns. The goal of the column pivoted QR algorithm is to find a set of columns that are “as linearly independent as possible.” This is not such a simple task, and so we settle for a greedy strategy: at each step, we select the column that is least well predicted (in the sense of residual norm) by columns already selected. This leads to the *pivoted QR factorization*

$$A\Pi = QR$$

where Π is a permutation and the diagonal entries of R appear in descending order (i.e. $r_{11} \geq r_{22} \geq \dots$). To decide on how many factors to keep in the factorization, we either automatically take the first k or we dynamically choose to take k factors where r_{kk} is greater than some tolerance and $r_{k+1,k+1}$ is not.

The pivoted QR approach has a few advantages. It yields *parsimonious* models that predict from a subset of the columns of A – that is, we need to measure fewer than n factors to produce an entry of b in a new column. It can also be computed relatively cheaply, even for large matrices that may be sparse.

3 Tikhonov

A second approach is to say that we want a model in which the coefficients are not too large. To accomplish this, we add a penalty term to the usual least squares problem:

$$\text{minimize } \|Ax - b\|^2 + \lambda^2\|x\|^2.$$

Equivalently, we can write

$$\text{minimize } \left\| \begin{bmatrix} A \\ \lambda I \end{bmatrix} x - \begin{bmatrix} b \\ 0 \end{bmatrix} \right\|^2,$$

which leads to the regularized version of the normal equations

$$(A^T A + \lambda^2 I)x = A^T b.$$

In some cases, we may want to regularize with a more general norm $\|x\|_M^2 = x^T M x$ where M is symmetric and positive definite, which leads to the regularized equations

$$(A^T A + \lambda^2 M)x = A^T b.$$

If we know of no particular problem structure in advance, the standard choice of $M = I$ is a good default.

It is useful to compare the usual least squares solution to the regularized solution via the SVD. If $A = U \Sigma V^T$ is the economy SVD, then

$$\begin{aligned} x_{LS} &= V \Sigma^{-1} U^T b \\ x_{Tik} &= V f(\Sigma)^{-1} U^T b \end{aligned}$$

where

$$f(\sigma)^{-1} = \frac{\sigma}{\sigma^2 + \lambda^2}.$$

This *filter* of the inverse singular values affects the larger singular values only slightly, but damps the effect of very small singular values.

4 Truncated SVD

The Tikhonov filter reduces the effect of small singular values on the solution, but it does not eliminate that effect. By contrast, the *truncated SVD* approach uses the filter

$$f(z) = \begin{cases} z, & z > \sigma_{\min} \\ \infty, & \text{otherwise.} \end{cases}$$

In other words, in the truncated SVD approach, we use

$$x = V_k \Sigma_k^{-1} U_k^T b$$

where U_k and V_k represent the leading k columns of U and V , respectively, while Σ_k is the diagonal matrix consisting of the k largest singular values.

5 ℓ^1 and the lasso

An alternative to Tikhonov regularization (based on a Euclidean norm of the coefficient vector) is an ℓ^1 regularized problem

$$\text{minimize } \|Ax - b\|^2 + \lambda \|x\|_1.$$

This is sometimes known as the “lasso” approach. The ℓ^1 regularized problem has the property that the solutions tend to become sparse as λ becomes larger. That is, the ℓ^1 regularization effectively imposes a factor selection process like that we saw in the pivoted QR approach. Unlike the pivoted QR approach, however, the ℓ^1 regularized solution cannot be computed by one of the standard factorizations of numerical linear algebra.

6 Tradeoffs and tactics

All four of the regularization approaches we have described are used in practice, and each has something to recommend it. The pivoted QR approach is relatively inexpensive, and it results in a model that depends on only a few factors. If taking the measurements to compute a prediction costs money — or even costs storage or bandwidth for the factor data! — such a model may be to our advantage. The Tikhonov approach is likewise inexpensive, and has a nice Bayesian interpretation (though we didn’t talk about it). The truncated SVD approach involves the best approximation rank k approximation to the original factor matrix, and can be interpreted as finding the k best factors that are linear combinations of the original measurements. The ℓ_1 approach again produces models with sparse coefficients; but unlike QR with column pivoting, the ℓ_1 regularized solutions incorporate information about the vector b along with the matrix A .

So which regularization approach should one use? In terms of prediction quality, all can provide a reasonable deterrent against ill-posedness and overfitting due to highly correlated factors. Also, all of the methods described have a parameter (the number of retained factors, or a penalty parameter λ) that governs the tradeoff between how well-conditioned the fitting problem will be and the increase in bias that naturally comes from looking at a smaller class of models. Choosing this tradeoff intelligently may be rather more important than the specific choice of regularization strategy. A detailed discussion of how to make this tradeoff is beyond the scope of the class; but we will see some of the computational tricks involved in implementing specific strategies for choosing regularization parameters before we are done.