

Regression and Gradients

We've seen how to fit non-linear functions using only least squares linear regression. From next week we will also see more useful things we can do just by manipulating Gaussian distributions. However, the material so far already raises questions that we can't answer with just the tools we've reviewed so far.

- How can we fit several hyperparameters without a massive grid search?
- Can we learn good basis functions from data, rather than choosing them by hand?
- For binary classification using regression the best function should match $P(y=1 | \mathbf{x})$. How can we avoid our fitted functions giving predictions outside $[0, 1]$?

To answer these questions, we'll need to fit non-linear optimization problems. These won't have closed-form solutions, and so we will have to approach them numerically. Many, but not all, optimization problems are solved by *gradient-based methods*, which we begin to look at in this note.

We start by looking at how linear least squares regression can be solved with gradient-based methods, so we can generalize it to other cases. In the next note we'll then fit a better regression model for classification: *logistic regression*.

1 Gradients for linear least squares regression

We can create a vector of residuals (differences between observed values and function values) for a linear regression model as follows:

$$\mathbf{r} = \mathbf{y} - X\mathbf{w}. \quad (1)$$

And previously we noticed that Matlab and NumPy know how to minimize the sum of the square residuals:

$$\mathbf{r}^\top \mathbf{r} = (\mathbf{y} - X\mathbf{w})^\top (\mathbf{y} - X\mathbf{w}) \quad (2)$$

$$= \mathbf{y}^\top \mathbf{y} - 2\mathbf{w}^\top X^\top \mathbf{y} + \mathbf{w}^\top X^\top X \mathbf{w}. \quad (3)$$

We'll now minimize this cost function using calculus.

Our first task is to find the *gradient* of this cost function with respect to the weights. That is, the vector of partial derivatives: $\nabla_{\mathbf{w}} \mathbf{r}^\top \mathbf{r}$. The gradient vector is a function of the weights. At a given position in weight-space, it points in the direction in which a small movement will increase the cost the most. You were asked to show this fact in the background self-test, question 6ii), which has an answer available if you need to review this material.

We can differentiate small matrix/vector expressions by writing them as sums, and using the elementary differentiation rules for scalars. For example:

$$\frac{\partial \mathbf{x}^\top \mathbf{y}}{\partial x_i} = \frac{\partial \sum_j x_j y_j}{\partial x_i} = y_i, \quad \Rightarrow \quad \nabla_{\mathbf{x}} [\mathbf{x}^\top \mathbf{y}] = \mathbf{y}. \quad (4)$$

and

$$\frac{\partial \mathbf{x}^\top A \mathbf{x}}{\partial x_i} = \frac{\partial \sum_{jk} x_j A_{jk} x_k}{\partial x_i} = \sum_k A_{ik} x_k + \sum_j x_j A_{ji}, \quad (5)$$

$$\Rightarrow \quad \nabla_{\mathbf{x}} [\mathbf{x}^\top A \mathbf{x}] = A \mathbf{x} + A^\top \mathbf{x}, \quad \text{or } 2A \mathbf{x} \text{ if } A \text{ is symmetric.} \quad (6)$$

After some experience, you might remember some of these matrix/vector rules. Other such rules can be found in references like *The Matrix Cookbook*. We will discuss a more systematic approach to differentiate large functions in a later note.

For now, we can use the two rules we've derived to differentiate the cost function above:

$$\nabla_{\mathbf{w}} [\mathbf{r}^\top \mathbf{r}] = -2X^\top \mathbf{y} + 2X^\top X \mathbf{w}. \quad (7)$$

2 A closed form solution

The partial derivatives are all zero at an optimum weight vector. If there is unique best setting of the weights, we can solve for where that happens:

$$-2X^T \mathbf{y} + 2X^T X \mathbf{w} = \mathbf{0} \quad (8)$$

$$\Rightarrow X^T X \mathbf{w} = X^T \mathbf{y} \quad (9)$$

$$\Rightarrow \mathbf{w} = (X^T X)^{-1} X^T \mathbf{y}. \quad (10)$$

If (and only if) there isn't a unique solution for the weights that give the minimum square error, then $(X^T X)$ is not invertible and the equation isn't defined.

The above expression is known as the *normal equation* solution of the least squares problem. In mathematical working it's a "closed-form expression", but to fit weights (to some precision) from data, we still need a numerical algorithm and don't necessarily get exactly the right answer. You could implement the normal equations solution as follows:

```
ww = np.linalg.solve(np.dot(X.T, X), np.dot(X.T, yy))
```

It's best practice in implementations to solve linear systems with a dedicated solver¹, rather than inverting the matrix $X^T X$. However even this code doesn't always give as accurate an answer as dedicated routines such as `lstsq`.²

2.1 Uniqueness of the solution

[The website version of this note has a question here.]

2.2 Another way to derive the solution

[The website version of this note has a question here.]

3 Iterative methods

There are also generic algorithms that iteratively improve an initial guess of some model parameters using a cost function and its gradient vectors. While linear regression has specialist solvers, we could apply these generic algorithms anyway. If they don't work in simple cases, they're not likely to be useful more generally!

A naive way to use the gradient $\nabla_{\mathbf{w}}[\mathbf{r}^T \mathbf{r}]$ is the steepest-descent method:

$$\mathbf{w} \leftarrow \mathbf{w} - \eta \nabla_{\mathbf{w}}[\mathbf{r}^T \mathbf{r}], \quad (11)$$

which uses a small step size η . The parameters are moved in the direction that makes the biggest immediate change. The rule is applied repeatedly, with the gradient re-evaluated before each update. There are several methods (for example provided by SciPy) that will converge faster. Sometimes *much* faster.

It may help to rewrite the gradient to understand what the steepest descent rule is doing.

$$\nabla_{\mathbf{w}}[\mathbf{r}^T \mathbf{r}] = 2X^T (X\mathbf{w}) - 2X^T \mathbf{y} \quad (12)$$

$$= 2X^T (\mathbf{f} - \mathbf{y}) \quad (13)$$

$$= 2 \sum_n \mathbf{x}^{(n)} (f^{(n)} - y^{(n)}). \quad (14)$$

1. The Cholesky-based solver in SciPy is specialized for solving with a positive-definite matrix, and would be faster here.

2. See *Solving Least Squares Problems*, Lawson and Hanson (1974), Chapter 19, which argues that it's better to use a QR decomposition of the data matrix than losing precision by forming the summary $X^T X$. Murphy's textbook also champions the QR approach. Some library routines also use a singular value decomposition (SVD). I must say I'm not totally convinced this care is necessary. For noisy and regularized machine learning problems I've tried, the normal equations approach seems fine, and is $\sim 10\times$ faster on my machine. However, I still frequently use the least squares solvers: they're still quite fast, convenient, and I don't have to keep checking whether the normal equations approach will be accurate.

For each example we look at the ‘prediction error’, $(f - y)$. The weights are pulled most in the direction of inputs that had large prediction errors, to reduce misfit in those directions.

In *Stochastic Gradient Descent* (SGD) we take just one example at a time (perhaps at random, or visit the examples in order). Each example gives a crude one-sample *Monte Carlo* approximation of the gradient sum:

$$\frac{1}{N} \nabla_{\mathbf{w}} [\mathbf{r}^\top \mathbf{r}] \approx 2\mathbf{x}^{(n)}(f^{(n)} - y^{(n)}). \quad (15)$$

We take a small step in minus this direction. Then pick another example and repeat. Each time we see an example, we move the weights in a direction proportional to just one of the input vectors.

If we have 100,000 datapoints, we perform 100,000 updates in one sweep through the dataset (a ‘training epoch’). In the traditional (batch) steepest gradient descent method, we only perform one update after looking at the whole dataset once. In the limit of an infinite stream of data, SGD can fit a model as the data comes in. A traditional batch method never gets started, as it can never compute *the* gradient.

Once we have a working gradient-based optimization procedure³, we can apply it to problems beyond linear regression. We need to identify a suitable cost function, which no longer needs to be least squares, and then obtain its gradients.

[The website version of this note has a question here.]

4 Further Reading

Bishop covers the least squares solution in Section 3.1.1. The equations include a basis function expansion (replacing X with Φ), and motivates the least squares cost from the log-likelihood of a Gaussian model.

Murphy derives the least squares solution in Chapter 7 slightly more slowly. This description also uses a Gaussian model to motivate the least squares cost function.

Barber Chapter 17 starts with linear regression.

Numerical Recipes (Press et al.) will tell you that steepest gradient descent is a bad algorithm, and describes more sophisticated alternatives. See the section on Conjugate Gradient Methods in Multidimensions (section 10.6 of the second edition or 10.8 of the third edition). The books have some nice descriptions, but I would stay clear of the code (better alternatives with free licenses are available). However, they only describe ‘batch’ methods, that look at an entire dataset before making each update. *Stochastic* (steepest) gradient descent is not such a bad algorithm, especially for machine learning. <https://arxiv.org/abs/1606.04838> is one recent survey.

5 Check your understanding

Optional questions for those with time to be thorough on their second reading:

- In the normal equation solution code, why was the suggested code:

```
np.linalg.solve(np.dot(X.T, X), np.dot(X.T, yy))
```

rather than

```
np.dot(np.linalg.solve(np.dot(X.T, X), X.T), yy) # ?
```

Do they give the same answer? Is there something else to consider?

3. And we might find one slightly more sophisticated than described above.

- One misguided way to add an “extra” feature would be to copy an existing one, i.e., create an extra column of the training inputs X equal to an existing column. Explain why there would be multiple settings of the weights that have the same least squares training cost. Explain why the normal equation approach becomes ill-defined. What would happen if you ran a gradient-based optimizer to fit the weights?
- You could answer a generalized version of the previous question, where the “extra” feature is an arbitrary linear combination of all the other features. The training inputs X would contain a column that was a linear combination of the other columns.
- Instead of applying the trick we used before to fit L2 regularized models, can you generalize the derivation of the normal equation to explicitly cover this case? (Many textbooks have the answer to this problem.)