

Lab 8: Model fitting

Not graded, just practice

Katie Schuler

2024-10-24

1 Model fitting

1. True or false, gradient descent and ordinary least squares are both iterative optimization algorithms.
 - (A) True
 - (B) False
2. What cost function have we been using to perform our gradient descent?
 - (A) standard deviation
 - (B) bootstrapping
 - (C) sum of squared error
 - (D) absolute error
3. True or false, when performing gradient descent on the model given by the equation $y = w_0 + w_1x_1 + w_2x_2$, we might arrive at a local minimum and miss the global one.
 - (A) True
 - (B) False
4. Which of the following would work to estimate the free parameters of a **nonlinear** model?
 - (A) gradient descent

- (B) ordinary least squares solution
 - (C) both work
5. True or false, in gradient descent, we search through all possible parameters in the parameter space.
- (A) True
 - (B) False

2 Model fitting in R

Questions 6-9 refer to the code and output below, performing gradient descent with `optim`:

```
optim(data = data, par = c(0,0), fn=SSE, method = "STGD")
```

```
$par
```

```
[1] 3.37930046 0.06683237
```

```
$value
```

```
[1] 959.4293
```

```
$counts
```

```
[1] 6
```

```
$convergence
```

```
[1] 0
```

6. How many steps did the gradient descent algorithm take? _____
7. What was the sum of squared error of the optimal parameters? _____
8. What coefficients does the algorithm converge on?
 - (A) 3.37930046, 0.06683237
 - (B) 0, 0
 - (C) 959.4293
 - (D) 6, 0

- (E) all of the above
9. What parameters were used to initialize the algorithm?
- (A) 3.37930046, 0.06683237
 - (B) 0, 0
 - (C) 959.4293
 - (D) 6, 0

Questions 10-12 refer to the output below from `lm()`:

Call:

```
lm(formula = y ~ x, data = data)
```

Coefficients:

(Intercept)	x
3.37822	0.06688

10. Use R notation to write the model specification.

```
y ~ x # this works (implicit intercept)
```

```
y ~ 1 + x # this also works (explicit intercept)
```

11. Given the model is specified by the equation $y = w_0 + w_1x_1$, what is the parameter estimate for $w_0 = \underline{\hspace{2cm}}$ and $w_1 = \underline{\hspace{2cm}}$.
12. True or false, for this model, `optim()` with gradient descent would converge on the same parameter estimates?
- (A) True
 - (B) False