Chris Piech
CS 109

Section #8
November 21, 2025

Section 8

1. **Recalibrating an Uncalibrated Model**

You have an uncalibrated binary classification model that outputs values $\hat{p} \in [0, 1]$. These outputs are meant to be the probability that $Y = 1$. However, the outputs from this model are not well-calibrated. For instance, among all examples where $\hat{p} = 0.9$, it was the case that Y was 1 only 70% of the time. To recalibrate the models outputs you decide to use Platt Recalibration, where the corrected probability that $Y = 1$ is:

$$P(Y = 1 \mid \hat{p}) = \sigma(a \cdot \hat{p} - 0.5)$$

$\sigma(z) = 1/(1 + e^{-z})$ is the sigmoid function and a is the parameter of the recalibration model. Here is the partial derivative of the Platt Recalibration model with respect to a:

$$\frac{\partial}{\partial a}\sigma(a \cdot \hat{p} - 0.5) = \sigma(a \cdot \hat{p} - 0.5) \cdot [1 - \sigma(a \cdot \hat{p} - 0.5)] \cdot \hat{p}$$

a. For a new datapoint the uncalibrated model outputs $\hat{p}$ of 0.9. If you use Platt Recalibration with $a = 2$ what is the recalibrated probability that $Y = 1$?

Plug into the formula: $a \cdot \hat{p} - 0.5 = 2(0.9) - 0.5 = 1.3$, so the recalibrated probability is $\sigma(1.3)$.

b. You are given a training dataset with n outputs from the uncalibrated model $(\hat{p}^{(i)}, y^{(i)})$ where $\hat{p}^{(i)}$ is the uncalibrated output and $y^{(i)} \in \{0, 1\}$ is the true binary outcome. Explain how you could estimate the value of a that makes the $y^{(i)}$ values as likely as possible. Solve for any and all partial derivatives required by your answer.

Solution v1 (more explanation):
This problem is related to logistic regression. In both, we can use MLE to estimate parameters, and in both, the likelihood comes from the continuous PMF of the Bernoulli, since here we are still doing binary classification (Y is either 0 or 1):

$$L(a) = \prod_{i=1}^{n} P(Y = 1 \mid \hat{p}^{(i)})^{y^{(i)}} (1 - P(Y = 1 \mid \hat{p}^{(i)}))^{1-y^{(i)}}$$

$$LL(a) = \sum_{i=1}^{n} y^{(i)} \log P(Y = 1 \mid \hat{p}^{(i)}) + (1 - y^{(i)}) \log(1 - P(Y = 1 \mid \hat{p}^{(i)}))$$

To find the value of a that maximizes $LL(a)$, we can take the derivative using the chain rule:

$$\frac{\partial LL(a)}{\partial a} = \frac{\partial LL(a)}{\partial P(Y = 1 \mid \hat{p}^{(i)})} \cdot \frac{\partial P(Y = 1 \mid \hat{p}^{(i)})}{\partial a}$$

The second component is given to us in the problem (it is equivalent to $\sigma(a \cdot \hat{p}^{(i)} - 0.5) \cdot [1 - \sigma(a \cdot \hat{p}^{(i)} - 0.5)] \cdot \hat{p}^{(i)}$). The first term looks the same as in logistic regression:

$$\frac{\partial LL(a)}{\partial P(Y = 1 \mid \hat{p}^{(i)})} = \sum_{i=1}^{n} \left(\frac{y^{(i)}}{P(Y = 1 \mid \hat{p}^{(i)})} - \frac{1 - y^{(i)}}{1 - P(Y = 1 \mid \hat{p}^{(i)})} \right)$$

Using this derivative, you would find the best estimate for a using a gradient ascent.

Solution v2 (less explanation):

To estimate a, set up the log-likelihood for Bernoulli outcomes and differentiate with respect to a. Using the chain rule:

$$\frac{\partial LL(a)}{\partial a} = \sum_{i=1}^{n} \left(\frac{y^{(i)}}{P_i} - \frac{1 - y^{(i)}}{1 - P_i} \right) \cdot \sigma(a\hat{p}^{(i)} - 0.5)\big[1 - \sigma(a\hat{p}^{(i)} - 0.5)\big]\hat{p}^{(i)}$$

where $P_i = P(Y = 1 \mid \hat{p}^{(i)})$.
Then update a using gradient ascent:

$$a \leftarrow a + \eta \frac{\partial LL(a)}{\partial a}$$

This gives the MLE estimate of the recalibration slope a.

2. **Decoding Movement for a Brain-Controlled Prosthetic Leg**
Engineers are designing a brain-controlled prosthetic ankle that infers a user's intended movement from electrical activity in their leg muscles. To train the system, the user performs known movements while electrodes measure muscle activity (EMG), producing labeled data that link muscle signals to intended actions.

At each time step, two muscle sensors are recorded:

- S_{TA}: tibialis anterior ("lift up" muscle),
- S_{GA}: gastrocnemius ("press down" muscle).

Each sensor reading is labeled as either `Active` (A) or `Quiet` (Q). The user can intend one of three movements:

$$U = \text{lift foot up}, \qquad D = \text{press foot down}, \qquad N = \text{neutral/relax}.$$

During calibration, the engineers measured how often each sensor fired while the user intended each movement. They found:

- When the user intends Up, sensor TA is Active 90% of the time and sensor GA is Active 20% of the time.
- When the user intends Down, sensor TA is Active 10% of the time and sensor GA is Active 85% of the time.
- When the user intends Neutral, sensor TA is Active 10% of the time and sensor GA is Active 10% of the time.

Engineers also found that when walking, a user spends about 30% of the time intending "Up," about 30% intending "Down," and the remaining 40% in "Neutral." Engineers model the two muscle sensors as independent once the users intended movement is specified.

Question. We observe S_{TA} = Active and S_{GA} = Active, which intendeed movement $M \in \{U, D, N\}$ is most likely?

First, we can formalize the information that has been provided to us in the problem statement.
Given Information:

- Prior probabilities: $P(U) = 0.30$, $P(D) = 0.30$, $P(N) = 0.40$
- Likelihoods for sensor TA:

$$P(S_{TA} = A \mid U) = 0.90, \quad P(S_{TA} = A \mid D) = 0.10, \quad P(S_{TA} = A \mid N) = 0.10$$

- Likelihoods for sensor GA:

$$P(S_{GA} = A \mid U) = 0.20, \quad P(S_{GA} = A \mid D) = 0.85, \quad P(S_{GA} = A \mid N) = 0.10$$

Objective: Find the most likely movement given both sensors are active, we want to find the movement that results in the highest probability given what we have observed about the sensors:

$$M^* = \arg \max_{M \in \{U,D,N\}} P(M \mid S_{TA} = A, S_{GA} = A)$$

Approach: Apply Bayes' theorem:

$$P(M \mid S_{TA} = A, S_{GA} = A) = \frac{P(S_{TA} = A, S_{GA} = A \mid M) \cdot P(M)}{P(S_{TA} = A, S_{GA} = A)}$$

Since the denominator is constant across all movements, we can focus on only maximizing the numerator:

$$M^* = \arg \max_{M \in \{U,D,N\}} P(S_{TA} = A, S_{GA} = A \mid M) \cdot P(M)$$

Calculations:
We can use the definition of conditional independence applied here to simplify our calculations: $P(S_{TA}, S_{GA} \mid M) = P(S_{TA} \mid M) \cdot P(S_{GA} \mid M)$ Using conditional independence:

$$P(S_{TA} = A, S_{GA} = A \mid U) = 0.90 \times 0.20 = 0.18$$
$$P(S_{TA} = A, S_{GA} = A \mid D) = 0.10 \times 0.85 = 0.085$$
$$P(S_{TA} = A, S_{GA} = A \mid N) = 0.10 \times 0.10 = 0.01$$

Computing the unnormalized posteriors:

$$P(U) \cdot P(S_{TA} = A, S_{GA} = A \mid U) = 0.30 \times 0.18 = 0.054$$
$$P(D) \cdot P(S_{TA} = A, S_{GA} = A \mid D) = 0.30 \times 0.085 = 0.0255$$
$$P(N) \cdot P(S_{TA} = A, S_{GA} = A \mid N) = 0.40 \times 0.01 = 0.004$$

Answer: Since $0.054 > 0.0255 > 0.004$, the most likely intended movement is **Up (U)**.

Normalization: We can calculate the exact probabilities using normalization.

To obtain the exact posterior probabilities, we normalize by dividing each unnormalized posterior by their sum:

First, compute the normalizing constant (evidence):

$$\begin{aligned} P(S_{TA} = A, S_{GA} = A) &= \sum_{M \in \{U,D,N\}} P(S_{TA} = A, S_{GA} = A \mid M) \cdot P(M) \\ &= 0.054 + 0.0255 + 0.004 \\ &= 0.0835 \end{aligned}$$

Then, compute the normalized posteriors:

$$P(U \mid S_{TA} = A, S_{GA} = A) = \frac{0.054}{0.0835} \approx 0.647$$
$$P(D \mid S_{TA} = A, S_{GA} = A) = \frac{0.0255}{0.0835} \approx 0.305$$
$$P(N \mid S_{TA} = A, S_{GA} = A) = \frac{0.004}{0.0835} \approx 0.048$$

Verification: $0.647 + 0.305 + 0.048 = 1.000$ ✓

$$P(U \mid \text{both active}) \approx 64.7\%, \quad P(D \mid \text{both active}) \approx 30.5\%, \quad P(N \mid \text{both active}) \approx 4.8\%$$