

WEEKLY PARTICIPATION 4: CONVEXITY OF LOGISTIC REGRESSION

I made the claim that all the optimization problems we have seen for fitting ML models so far are in fact convex optimization problems. Let's work through the general template for proving this by showing that binary logistic regression as seen in Homework 1,

$$\min_{\omega' = [\omega; \omega_0] \in \mathbb{R}^{d+1}} \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-y_i(\omega_0 + \omega^T \mathbf{x}_i))) + \lambda \|\omega\|_2^2$$

is convex.

- Rewrite the objective by writing the data-fitting terms as the composition of the logsumexp function and affine functions of ω' , and the regularization term as the composition of a linear function of ω' and the squared Euclidean norm.
- Identify some functions that we showed (or claimed) in class are convex, and use your knowledge of how to combine functions to get other convex functions, to argue that the objective is convex.
- Complete the argument that this is a convex optimization problem.