

How many labels do you have?


Some perspectives on gold standard labels

John Duchi

Based on joint work with Chen Cheng and Hilal Asi




The “standard” story in statistics & ML



The big picture

- Excited about the full pipeline of statistical machine learning



Motivation

Dave Donoho, “50 Years of Data Science”

It is no exaggeration to say that the combination of a Predictive Modeling culture together with Common Task Framework is the ‘secret sauce’ of machine learning

Common Task Framework:

1. A publicly available training dataset
2. A set of enrolled competitors whose common task is to infer a class prediction rule from the training data
3. A scoring referee to which competitors submit their prediction rule(s)

ImageNet

The (probably) currently preferred image classification benchmark


Input data X



Goal: assign label Y to
this image
(In this case, Y =
Golden Retriever)

Dataset description: For each of 1000 image categories (e.g. cherry, bow and arrow, golden retriever, dachshund) there are 1000 representative images

ImageNet Progress



Little exaggeration to say deep learning descends from ImageNet

Supervised Learning

The construction of ImageNet isn't really what we teach

- Usual machine learning story

Input data X



Noisy Label Y

“Dog”

“Golden Retriever Puppy”

“Cat”

Machine learning pipeline:

Feed in a bunch of pairs

$$\{(X_i, Y_i)\}_{i=1}^N$$

(magical fitting...)


Output a model

$$\hat{Y} = f(X)$$

ImageNet construction



WordNet hierarchy



Tags ?


bicycle

bike

canon

ImageNet construction

Select all the images that contain a bicycle



ImageNet construction

Select all the images that contain a bicycle



Bicycle object class:

Include X as an example of
 $Y = \text{bicycle}$ if selection
frequency $\geq 70\%$

Repeat for each of 1000 classes

What this is:


Ingenious, clever, surprisingly effective

What this is not:

Noisy labels Y given X


How much does data construction matter?

Even when we're careful, things get weird



Also, methods are quite overconfident in predictions (e.g., predict classes with 90+% certainty)

Remainder of this talk



Steps:

1. Propose a model
2. Analyze the model
3. It makes some predictions: test them!

The model

- Binary classification with m labelers

$$Y \in \{-1, 1\}, \quad X \in \mathbb{R}^d$$


- Symmetric link function

$$\mathbb{P}(Y = y \mid X = x) = \sigma(yx^\top \theta^*)$$

- Data in tuples (n total tuples)

$$(X, Y_1, \dots, Y_m), \quad Y_j \mid X \stackrel{\text{iid}}{\sim} \mathbb{P}(\cdot \mid X)$$

- Covariate vectors $X_i \stackrel{\text{iid}}{\sim} \mathcal{N}(0, I_d)$



The model (continued)

- Margin-based loss ℓ satisfying

$$\ell'(t) = -\sigma(-t)$$


- Loss of parameter θ on (x, y)

$$\ell(\theta^T xy)$$

- E.g. logistic regression

$$\ell(t) = \log(1 + e^t)$$

$$\sigma(t) = \frac{1}{1 + e^{-t}}$$



The two estimators

- Use all the labels

$$L_n(\theta) := \frac{1}{n} \frac{1}{m} \sum_{i=1}^n \sum_{j=1}^m \ell(Y_{ij} X_i^\top \theta)$$

(Log likelihood for multiple labels)

$$\hat{\theta}_n = \operatorname{argmin}_{\theta} L_n(\theta)$$

- Use majority vote

$$L_n^{\text{mv}}(\theta) := \frac{1}{n} \sum_{i=1}^n \ell(\bar{Y}_i X_i^\top \theta)$$

where $\bar{Y}_i = \text{Majority}(Y_{i1}, \dots, Y_{im})$

$$\hat{\theta}_n^{\text{mv}} = \operatorname{argmin}_{\theta} L_n^{\text{mv}}(\theta)$$

Main quantities of interest:

- Calibration error $\|\hat{\theta} - \theta^*\|_2$
- Classification error $\|\hat{u} - u^*\|_2$ where $u = \theta / \|\theta\|_2$ is unit

Convergence of the MLE

$$L_n(\theta) := \frac{1}{n} \frac{1}{m} \sum_{i=1}^n \sum_{j=1}^m \ell(Y_{ij} X_i^\top \theta) \quad L(\theta) = \mathbb{E}[\ell(Y X^\top \theta)]$$
$$\hat{\theta}_n = \operatorname{argmin}_{\theta} L_n(\theta)$$

Theorem:

Under the well-specified model, we have asymptotic normality

$$\sqrt{n} \left(\hat{\theta}_n - \theta^\star \right) \xrightarrow{d} \mathcal{N} \left(0, \frac{1}{m} \nabla^2 L(\theta^\star)^{-1} \operatorname{Cov}(\dot{\ell}_{\theta^\star}) \nabla^2 L(\theta^\star)^{-1} \right)$$

and

$$\sqrt{n} \left(\hat{u}_n - u^\star \right) \xrightarrow{d} \mathcal{N} \left(0, \frac{1}{m \|\theta^\star\|_2^2} (I - u^\star u^{\star\top}) \right)$$

Convergence of majority vote

$$L_n^{\text{mv}}(\theta) := \frac{1}{n} \sum_{i=1}^n \ell(\bar{Y}_i X_i^\top \theta) \quad \text{where } \bar{Y}_i = \text{Majority}(Y_{i1}, \dots, Y_{im})$$

Decompose $X = u^* Z + (I - u^* u^{*\top})X = u^* Z + W$

Theorem: Under the model, we have “overconfident” convergence

$$\hat{\theta}_n^{\text{mv}} \xrightarrow{p} t_m u^* \quad \text{where} \quad t_m \asymp \sqrt{m}$$

and asymptotic normality

$$\begin{aligned} \sqrt{n} (\hat{u}_n^{\text{mv}} - u^*) &\xrightarrow{d} \text{N}\left(0, \frac{1}{t_m^2} H(t_m)^\dagger C(t_m) H(t_m)^\dagger\right) \\ &\stackrel{\text{dist}}{=} \text{N}\left(0, \frac{c(1 + o_m(1))}{\sqrt{m}} (I - u^* u^{*\top})\right) \end{aligned}$$

for matrices $H(t) = \frac{1}{4t} \mathbb{E}[W W^\top](1 + o(1))$ and $C(t) = \frac{c}{t} \mathbb{E}[W W^\top](1 + o(1))$

Robustness of majority vote

Theorem: With misspecified link, we have “overconfident” convergence

$$\hat{\theta}_n^{\text{mv}} \xrightarrow{p} t_m u^* \quad \text{where} \quad t_m \asymp \sqrt{m}$$

and asymptotic normality (for fixed Σ)

$$\sqrt{n} (\hat{u}_n^{\text{mv}} - u^*) \xrightarrow{d} \text{N} \left(0, \frac{1 + o_m(1)}{\sqrt{m}} \Sigma \right)$$

Take home messages:

- Majority vote is (unfixably) uncalibrated and overconfident
- More robust (doesn't matter if the link is correct)
- Less efficient when the link *is* correct

Extensions: semiparametric estimates

- Corrected estimator: fit the model, refit the link, refit the model

$$\hat{\theta}^{\text{mv}} = \operatorname{argmin} L_n^{\text{mv}}(\theta)$$

$$\hat{\sigma} = \operatorname{argmin} \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m (2\sigma(\hat{u}^\top X_i Y_{ij}) - 1 - Y_{ij})^2$$


Theorem: Under appropriate conditions,

$$\hat{\theta} = \operatorname{argmin} \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m \ell_{\hat{\sigma}}(\theta^\top X_i Y_{ij})$$

is efficient: $\sqrt{n} \left(\hat{\theta} - \theta^* \right) \xrightarrow{d} \mathcal{N} \left(0, \frac{1}{m} I(\theta^*)^{-1} \right)$

More link-estimation

- Setting: individual links
- Estimate each link, then model



$$\ell_{\vec{\sigma}}(\theta; X_i, Y_i) = \frac{1}{m} \sum_{j=1}^m \ell_{\sigma_j}(\theta^\top X_i Y_{ij})$$

Theorem: so long as estimated links are consistent at all

$$\hat{\theta}_n = \operatorname{argmin} \frac{1}{n} \sum_{i=1}^n \ell_{\vec{\sigma}_n}(\theta; X_i, Y_i)$$

satisfies

$$\sqrt{n} \left(\hat{\theta}_n - \theta^\star \right) \xrightarrow{d} \mathbf{N} \left(0, \frac{1}{m} \Sigma_\star \right)$$


Experimental results

If our model is reasonable, it should make real predictions



- BlueBirds: *Indigo Bunting* versus *Blue Grosbeak* [Welinder, Branson, Perona, Belongie 10]





- CIFAR-10H: soft labels of CIFAR-10 test set [Peterson, Battleday, Griffiths, Russakovsky 19]



Experimental results: bluebirds



Experimental results: CIFAR-10H




Semi-synthetic experiment



- Re-generate labels for CIFAR-10 using ResNet18 networks



- Vary accuracy to simulate labeler accuracy



Semi-synthetic results



Conclusions and next steps

- Interesting to think about dataset construction: a place for statistics to lay down some intellectual foundations
- Would obtaining data with (human) perceptual uncertainty help build better prediction methods?
- Currently limited datasets like those above: develop datasets to drive progress we want to see
- Fun to make (theoretical) predictions that can be wrong
- Implications for LLMs and generative models: substantial work to find data?