

RECITATION 7

LEARNING THEORY AND GENERATIVE MODELS

10-301/10-601: INTRODUCTION TO MACHINE LEARNING

10/21/2021

1 Learning Theory

Some Important Definitions and Theorems

1. Basic notations:

- **True function** (expert/oracle) $c^* : X \rightarrow Y$ (unknown)
- **Hypothesis space** \mathcal{H} and **hypothesis** $h \in \mathcal{H} : X \rightarrow Y$
- Probability Distribution p^* (unknown)
- Training Dataset $S = \{x^{(1)}, \dots, x^{(N)}\}$

2. **True Error (expected risk)**

$$R(h) = P_{x \sim p^*(x)}(c^*(x) \neq h(x))$$

3. **Train Error (empirical risk)**

$$\begin{aligned}\hat{R}(h) &= P_{x \sim S}(c^*(x) \neq h(x)) \\ &= \frac{1}{N} \sum_{i=1}^N \mathbb{1}(c^*(x^{(i)}) \neq h(x^{(i)})) \\ &= \frac{1}{N} \sum_{i=1}^N \mathbb{1}(y^{(i)} \neq h(x^{(i)}))\end{aligned}$$

4. **PAC criterion** is that we produce a high accuracy hypothesis with high probability. More formally,

$$P(\forall h \in \mathcal{H}, |R(h) - \hat{R}(h)| \leq \epsilon) \geq 1 - \delta$$

5. A hypothesis $h \in \mathcal{H}$ is **consistent** with training data S if $\hat{R}(h) = 0$ (zero training error/correctly classify)

6. **Sample Complexity** is the minimum number of training examples N such that PAC criterion is satisfied for a given ϵ (arbitrarily small error) and δ (with high probability)

7. Sample Complexity for 4 Cases: See Figure 1. Note that

	Realizable	Agnostic
Finite $ \mathcal{H} $	Thm. 1 $N \geq \frac{1}{\epsilon} [\log(\mathcal{H}) + \log(\frac{1}{\delta})]$ labeled examples are sufficient so that with probability $(1 - \delta)$ all $h \in \mathcal{H}$ with $\hat{R}(h) = 0$ have $R(h) \leq \epsilon$.	Thm. 2 $N \geq \frac{1}{2\epsilon^2} [\log(\mathcal{H}) + \log(\frac{2}{\delta})]$ labeled examples are sufficient so that with probability $(1 - \delta)$ for all $h \in \mathcal{H}$ we have that $ R(h) - \hat{R}(h) \leq \epsilon$.
Infinite $ \mathcal{H} $	Thm. 3 $N = O(\frac{1}{\epsilon} [\text{VC}(\mathcal{H}) \log(\frac{1}{\epsilon}) + \log(\frac{1}{\delta})])$ labeled examples are sufficient so that with probability $(1 - \delta)$ all $h \in \mathcal{H}$ with $\hat{R}(h) = 0$ have $R(h) \leq \epsilon$.	Thm. 4 $N = O(\frac{1}{\epsilon^2} [\text{VC}(\mathcal{H}) + \log(\frac{1}{\delta})])$ labeled examples are sufficient so that with probability $(1 - \delta)$ for all $h \in \mathcal{H}$ we have that $ R(h) - \hat{R}(h) \leq \epsilon$.

12

Figure 1

- **Realizable** means $c^* \in \mathcal{H}$
 - **Agnostic** means c^* may or may not be in \mathcal{H}
8. **VC dimension** of a hypothesis space \mathcal{H} is the maximum number of points such that there exists at least one arrangement of these points and a hypothesis $h \in \mathcal{H}$ that is consistent with any labelling of this arrangement of points.
 9. If $\text{VC}(\mathcal{H}) = n$, then for all placements of $n + 1$ points, there exists no hypothesis $h \in \mathcal{H}$ that can shatter any of it.

Questions

1. For the following examples, write whether or not the data points can be shattered using a linear classifier
 - 2 points in 1D
 - 3 points in 1D
 - 3 points in 2D
 - 4 points in 2D

How many points can a linear boundary (with bias) classify exactly for d-Dimensions?

2. Consider a semicircle classifier, where all points within the semicircle must equal 1 and all points outside must equal -1 (Rotated and scaled semi-circles are valid)
 - (a) Which of the configurations of 4 points in figure 2 can a semicircle shatter?




Figure 2

(b) What about the configurations of 5 points in figure 3?




Figure 3

3. Let x_1, x_2, \dots, x_n be n random variables that represent binary literals ($x \in \{0, 1\}^n$). Let the hypothesis class H_d denote the class of depth = d decision trees.

Find the minimum number of examples required to learn $h \in H_3$ which guarantees at least 99% accuracy with at least 98% confidence.

2 MLE/MAP

As a reminder, in MLE, we have

$$\begin{aligned}\hat{\theta}_{MLE} &= \arg \max_{\theta} p(\mathcal{D}|\theta) \\ &= \arg \min_{\theta} -\log(p(\mathcal{D}|\theta))\end{aligned}$$

For MAP, we have

$$\begin{aligned}\hat{\theta}_{MAP} &= \arg \max_{\theta} p(\theta|\mathcal{D}) \\ &= \arg \max_{\theta} \frac{p(\mathcal{D}|\theta)p(\theta)}{\text{Normalizing Constant}} \\ &= \arg \max_{\theta} p(\mathcal{D}|\theta)p(\theta) \\ &= \arg \min_{\theta} -\log(p(\mathcal{D}|\theta)p(\theta))\end{aligned}$$

Imagine you are a data scientist working for an advertising company. The advertising company has recently run an ad and they want you to estimate its performance. The ad was shown to N people. $Y^{(i)} = 1$ if person i clicked on the ad and 0 otherwise. Thus $\sum_i^N y^{(i)} = k$ people decided to click on the ad. Assume that the probability that the i -th person clicks on the ad is θ and the probability that the i -th person does not click on the ad is $1 - \theta$.

1. Note

$$p(\mathcal{D}|\theta) = p((Y^{(1)}, Y^{(2)}, \dots, Y^{(k)})|\theta) = \theta^k(1 - \theta)^{N-k}$$

Calculate $\hat{\theta}_{MLE}$.

2. Suppose $N = 100$ and $k = 10$. Calculate $\hat{\theta}_{MLE}$.
3. Your coworker tells you that $\theta \sim \text{Beta}(\alpha, \beta)$. That is:

$$p(\theta) = \frac{\theta^{\alpha-1}(1 - \theta)^{\beta-1}}{B(\alpha, \beta)}$$

Recall from lecture that $\hat{\theta}_{MAP}$ for a Bernoulli random variable with a Beta prior is given by:

$$\hat{\theta}_{MAP} = \frac{k + \alpha + 1}{N + \alpha + \beta - 2}$$

Suppose $N = 100$ and $k = 10$. Furthermore, you believe that in general people click on ads about 6 percent of the time, so you, somewhat naively, decide to set $\alpha = 6 + 1 = 7$, and $\beta = 100 - 6 + 1 = 95$. Calculate $\hat{\theta}_{MAP}$.

4. How do $\hat{\theta}_{MLE}$ and $\hat{\theta}_{MAP}$ differ? Argue which estimate you think is better.

5. Suppose a rental car service operates in your city. Drivers can drop off and pick up cars anywhere inside the city limits. You can find and rent cars using an app. Suppose you wish to find the probability that you can find a rental car within a short distance of your home address at any time of day. Over three days you look at the app and find the following number of cars within a short distance of your home address:

$$x = [3, 4, 1].$$

Your coworker tells you that this comes from a Poisson distribution. That is:

$$p(x|\theta) = \frac{e^{-\theta}\theta^x}{x!}$$

Also, you are told that

$$\theta \sim \text{Gamma}(2, 2)$$

That is:

$$p(\theta) = \frac{1}{\Gamma(\alpha)\beta^\alpha} \theta^{(\alpha-1)} e^{-\frac{\theta}{\beta}}, \quad \theta > 0$$

Calculate $\hat{\theta}_{MAP}$.

(Writeup from https://en.wikipedia.org/wiki/Conjugate_prior#Practical_example)

3 Convolutional Neural Network

3.1 Concepts

What are filters?

What are convolutions?

What are the shapes of the input and output tensors¹ of a CNN layer?

¹tensors are the generalization of matrices to multi-dimensions

What are some benefits of CNNs over fully connected (also called dense) layers?

3.2 Parameters

Suppose that we want to classify images that belong to one of ten possible classes (i.e. [cat, dog, bird, turtle, ..., horse]). The images come in RGB format (one channel for each color), and are downsampled to dimension 128x128.

Figure 4 illustrates one such image from the MS-COCO dataset².



Figure 4: Image of a horse from the MS-COCO dataset, downsampled to 128x128

We construct a Convolutional Neural Network that has the following structure: the input is first max-pooled with a 2x2 filter with stride 2 and 3 output channels. The results are then sent to a convolutional layer that uses a 17x17 filter of stride 1 and 12 output channels. Those values are then passed through a max-pool with a 3x3 filter with stride 3 and also 12 output channels. The result is then flattened and passed through a fully connected layer (ReLU activation) with 128 hidden units followed by a fully connected layer (softmax activation) with 10 hidden units. We say that the final 10 hidden units thus represent the categorical probability for each of the ten classes. With enough labeled data, we can simply use some optimizer like SGD to train this model through backpropagation.

Note: By default, please assume we have bias terms in all neural network layers unless explicitly stated otherwise.

1. **Draw a diagram that illustrates the channels and dimensions of the tensors before and after every neural net operation.**
2. **How many parameters are in this network for the convolutional components?**

²<https://cocodataset.org/>

3. How many parameters are in this network for the fully connected (also called dense) components?
4. From these parameter calculations, what can you say about convolutional layers and fully connected layers in terms of parameter efficiency³? Why do you think this is the case?

³the ratio between the number of parameters from some layer type and the total number of parameters.

3.3 Links

Visualization of convolutional filter sweep steps https://github.com/vdumoulin/conv_arithmetic

Visualization of convolutional filter smooth sweep with outputs <https://www.youtube.com/watch?v=f0t-0CG79-U>

Visualization of neural network layer outputs <http://cs231n.stanford.edu/>

The architecture used there is (conv → relu → conv → relu → pool) x3 → fc → softmax