Recitation 6
Probabilistic Learning, CNNs, Learning Theory

10-301/10-601: Introduction to Machine Learning

11/02/2022

1 Probabilistic Learning

In probabilistic learning, we are trying to learn a target probability distribution as opposed to a target function. We'll review two ways of estimating the parameters of a probability distribution, as well as one family of probabilistic models: Naive Bayes classifiers.

1.1 MLE/MAP

As a reminder, in MLE, we have

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

For MAP, we have

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

---

1. Imagine you are a data scientist working for an advertising company. The advertising company has recently run an ad and wants you to estimate its performance.

   The ad was shown to N people. Let $Y^{(i)} = 1$ 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 $1 - \theta$.

(a) Note that

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

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

(b) Suppose $N = 100$ and $k = 10$. Calculate $\hat{\theta}_{MLE}$.

(c) 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 $\hat{\theta}_{MAP}$.

(d) How do $\hat{\theta}_{MLE}$ and $\hat{\theta}_{MAP}$ differ in this scenario? Argue which estimate you think is better.

2. Suppose you are an avid Neural and Markov fan who monitors the @neuralthenarwhal Instagram account each day. Suppose you wish to find the probability that Neural or Markov will post at any time of day. Over three days you look on Instagram and find the following number of new posts: $x = [3, 4, 1]$

   A fellow fan tells you that this comes from a Poisson distribution:

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

   Also, you are told that $\theta \sim \text{Gamma}(2, 2)$ — that is, its pdf is:

   $$p(\theta) = \frac{1}{4}\theta e^{-\frac{\theta}{2}},\ \theta > 0$$

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

   (See also https://en.wikipedia.org/wiki/Conjugate_prior)

1.2 Naive Bayes

By applying Bayes' rule, we can model the probability distribution $P(Y|X)$ by estimating $P(X|Y)$ and $P(Y)$.

$$P(Y|X) \propto P(Y)P(X|Y)$$

The Naive Bayes assumption greatly simplifies estimation of $P(X|Y)$ - we assume the features X_d are independent given the label. With math:

$$P(X|Y) = \_\_\_\_\_$$

Different Naive Bayes classifiers are used depending on the type of features.

- Binary Features: Bernoulli Naive Bayes - $X_d \,|\, Y = y \sim \text{Bernoulli}(\theta_{d,y})$
- Discrete Features: Multinomial Naive Bayes - $X_d \,|\, Y = y \sim \text{Multinomial}(\theta_{d,1,y}, \ldots, \theta_{d,K-1,y})$
- Continuous Features: Gaussian Naive Bayes - $X_d \,|\, Y = y \sim \mathcal{N}(\mu_{d,y}, \sigma^2_{d,y})$

We'll walk through the process of learning a Bernoulli Naive Bayes classifier. Consider the dataset below. You are looking to buy a car; the label is 1 if you are interested in the car and 0 if you aren't. There are three features: whether the car is red (your favorite color), whether the car is affordable, and whether the car is fuel-efficient.

Interested?	Red?	Affordable?	Fuel-Efficient?
1	1	1	1
0	0	1	0
0	0	1	1
1	0	0	0
0	0	1	1
0	0	1	1
1	1	1	1
1	1	0	1
0	0	0	0

1. How many parameters do we need to learn?

2. Estimate the parameters via MLE.

3. If I see a car that is red, not affordable, and fuel-efficient, would the classifier predict that I would be interested in it?

4. Is there a problem with this classifier based on your calculations for the previous question? If so, how can we fix it?

5. Now we will derive the decision boundary of a 2D Gaussian Naïve Bayes. Show that this decision boundary is quadratic. That is, show that $p(y = 1 \mid x_1, x_2) = p(y = 0 \mid x_1, x_2)$ can be written as a polynomial function of x_1 and x_2 where the degree of each variable is at most 2. You may fold *unimportant* constants into terms such as C, C', C'', C''' so long as *you are clearly showing each step.*

2 Learning Theory

2.1 PAC Learning

Some Important Definitions

1. Basic notation:
   - Probability distribution (unknown): $X \sim p^*$
   - **True function** (unknown): $c^* : X \to Y$
   - **Hypothesis space** $\mathcal{H}$ and **hypothesis** $h \in \mathcal{H} : X \to Y$
   - Training dataset $\mathcal{D} = \{x^{(1)}, \ldots, x^{(N)}\}$
2. **True Error (expected risk)**

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

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

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

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

$$P(\forall h \in \mathcal{H}, \_\_\_\_\_ \leq \_\_\_\_\_) \geq \_\_\_\_\_$$

Sample Complexity is the minimum number of training examples N such that the PAC criterion is satisfied for a given ϵ and δ

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

- **Realizable** means $c^* \in \mathcal{H}$
- **Agnostic** means c^* may or may not be in $\mathcal{H}$

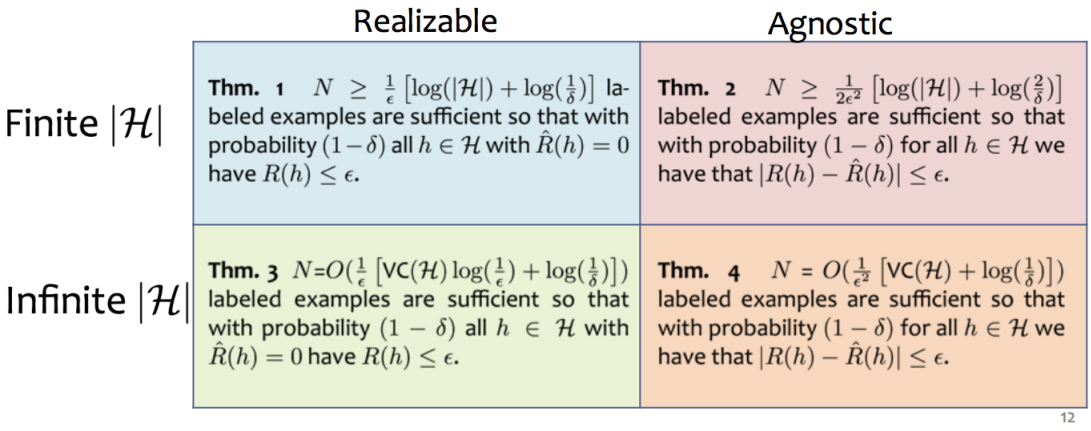

Figure 1: Sample Complexity for 4 Cases

The **VC dimension** of a hypothesis space $\mathcal{H}$, denoted $\text{VC}(\mathcal{H})$ or $d_{VC}(\mathcal{H})$, 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.

To show that $\text{VC}(\mathcal{H}) = n$:

-
-

Questions

1. For the following examples, write whether or not there exists a dataset with the given properties that can be shattered by 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. In the below table, state in which case the sample complexity of the hypothesis falls under.

<table>
<tr><th>Problem</th><th>Hypothesis Space</th><th>Realizable/ Agnostic</th><th>Finite/ Infinite</th></tr>
<tr><td>A binary classification problem, where the data points are linearly separable</td><td>Set of all linear classifiers</td><td></td><td></td></tr>
<tr><td>Predict whether it will rain or not based on the following dataset:
<table><tr><th>Temp</th><th>Humid</th><th>Wind</th><th>Rain?</th></tr><tr><td>High</td><td>Yes</td><td>Yes</td><td>Yes</td></tr><tr><td>Low</td><td>Yes</td><td>No</td><td>No</td></tr><tr><td>Low</td><td>No</td><td>Yes</td><td>Yes</td></tr><tr><td>High</td><td>No</td><td>No</td><td>Yes</td></tr></table></td><td>A decision tree with max depth 2, where each node can only split on one feature, and the features cannot be repeated along a branch</td><td></td><td></td></tr>
<tr><td>Classifying a set of real-valued points where the underlying data distribution is unknown</td><td>Set of all linear classifiers</td><td></td><td></td></tr>
<tr><td>A binary classification problem on a given set of data points, where the data is not linearly separable</td><td>K-nearest neighbour classifier with Euclidean distance as distance metric</td><td></td><td></td></tr>
</table>

3. Consider a rectangle classifier (i.e. the classifier is uniquely defined 3 points $x_1, x_2, x_3 \in \mathbb{R}^2$ that specify 3 out of the four corners), where all points within the rectangle must equal 1 and all points outside must equal -1

   (a) Which of the configurations of 4 points in figure 2 can a rectangle shatter?

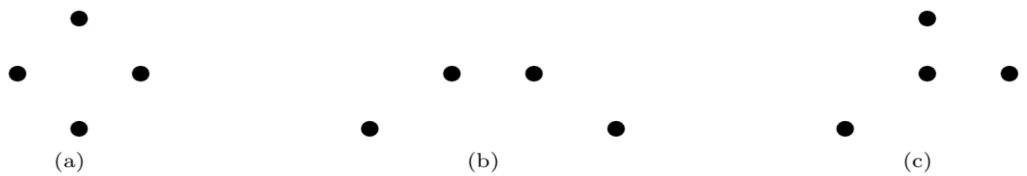

Figure 2

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

Figure 3

4. Let $x_1, x_2, ..., x_n$ be n random variables that represent binary literals ($x \in \{0,1\}^n$). Let the hypothesis class $\mathcal{H}_n$ denote the conjunctions of no more than n literals in which each variable occurs at most once. Assume that $c^* \in \mathcal{H}_n$.

   Example: For $n = 4$, $(x_1 \wedge x_2 \wedge x_4), (x_1 \wedge \neg x_3) \in \mathcal{H}_4$

   Find the minimum number of examples required to learn $h \in \mathcal{H}_{10}$ which guarantees at least 99% accuracy with at least 98% confidence.

3 Convolutional Neural Networks

3.1 Concepts

1. What are filters?
    - Filters (also called kernels) are feature extractors in the form of a small matrix used in convolutional neural layers. They usually have a width, height, depth, stride, padding, channels (output) associated with them.
2. What are convolutions?
    - We sweep the filter around the input tensor and take element-wise product sums based on factors such as filter size, stride, padding. These output product-sums form a new tensor, which is the output of a convolutional layer.
3. How do we calculate the output shape of a convolution?
    - Given input width W_{in}, kernel width K_W, padding P, and stride S, the output width W_{out} can be calculated as:

$$W_{out} = \lfloor \frac{W_{in} - K_w + 2 \times P}{S} \rfloor + 1$$

    .
    - Output height can be calculated similarly.
4. What are some benefits of CNNs over fully connected (also called dense) layers?
    - Good for image-related machine learning (learns the kernels that do feature engineering)
    - Pseudo translational invariance
    - Parameter efficient
5. How does the number of channels vary through convolutional networks?
    - Each convolution filter will have as many channels as the input, and there will be as many filters as there are output channels.
    - Pooling and activations often maintain the number of channels.

3.2 Dance Dance Convolution

Consider the following 4 x 4 image and 2x2 filter below.

1	3	-2	4
0	8	6	5
2	1	-9	0
4	-1	3	7

1	2
-2	-1

1. Assume that there is no padding and stride = 1. What are the dimensions of the output, and what is the value in the bottom right corner of the output image?

2. Now assume that we having padding = 1. Given that, what are the new dimensions of the output, and the new value in the bottom right corner?

3.3 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 dataset[1].

[1] https://cocodataset.org/

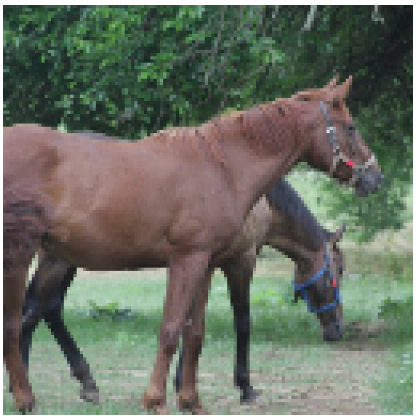

Figure 4: Image of a horse from the MS-COCO dataset, downsampled to 128x128

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 backprogation.

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

1. Fill the table below with channels and dimensions of the tensors before and after every neural net operation.

Layer / Operation	**Shape**
Input	3@128 × 128
maxpool-1	(a)
conv	(b)
maxpool-2	(c)
flatten	(d)
fully-connected-1	(e)
ReLU	(f)
fully-connected-1	(g)
softmax	(h)

2. Draw a diagram that illustrates the above table.

3. How many parameters are in this network for the convolutional components?

4. How many parameters are in this network for the fully connected (also called dense) components?

5. From these parameter calculations, what can you say about convolutional layers and fully connected layers in terms of parameter efficiency[2]? Why do you think this is the case?

3.4 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-OCG79-U

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

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

[2]the ratio between the number of parameters from some layer type and the total number of parameters.

4 Recurrent Neural Networks

4.1 Sample RNN

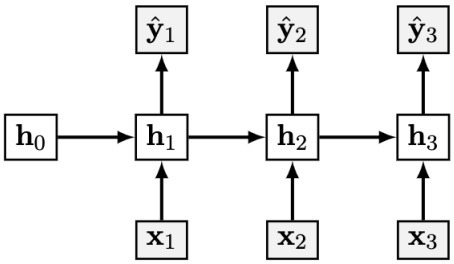

Where the layers and their corresponding weights are given below:

$$\mathbf{x}_t \in \mathbb{R}^3 \qquad \mathbf{W}_{hx} \in \mathbb{R}^{4\times 3}$$
$$\mathbf{h}_t \in \mathbb{R}^4 \qquad \mathbf{W}_{yh} \in \mathbb{R}^{2\times 4}$$
$$\mathbf{y}_t, \hat{\mathbf{y}}_t \in \mathbb{R}^2 \qquad \mathbf{W}_{hh} \in \mathbb{R}^{4\times 4}$$

$$\begin{aligned}\hat{\mathbf{y}}_t &= \sigma(\mathbf{o}_t)\\ \mathbf{o}_t &= \mathbf{W}_{yh}\mathbf{h}_t\\ \mathbf{h}_t &= \psi(\mathbf{z}_t)\\ \mathbf{z}_t &= \mathbf{W}_{hh}\mathbf{h}_{t-1} + \mathbf{W}_{hx}\mathbf{x}_t\end{aligned}$$

Where σ and ψ are activations.

1. Redraw the above diagram in a compact form such that we don't need to unroll it across several timesteps.

4.2 Concepts

1. What are recurrent neural networks?
    - According to Wikipedia [3], a recurrent neural network (RNN) can be characterized by connections between nodes creating a cycle. Outputs from some nodes

[3] Article linked here.

can affect subsequent computations. This allows it to exhibit temporal dynamic behavior.

- the recurrent nature makes them useful when the input is sequential (or temporal).

2. How do they use both inputs and previous outputs?
   - Hidden nodes have two sets of weights, one to process input from the previous layer, and one to process their own outputs from the previous timestep.
3. How do we optimize RNNs?
   - Applying chain rule to the 'unrolled' RNN (as above) is no different than a regular feed forward neural network aside from the fact that the same parameters are repeated throughout the network at each timestep.
   - Called as backpropagation through time (BPTT).