Active Learning Recitation

Colin White, Nupur Chatterji,
Kenny Marino

Outline

- Review Active Learning Model
- Active Learning Algorithms
- Active Learning Practice Problems

Active Learning

- Supervised Learning: training set contains labeled examples

Massive amounts of data, could be costly or time-consuming to label

- Active Learning: Do we need to have labels for *every* datapoint?

Batch Active Learning

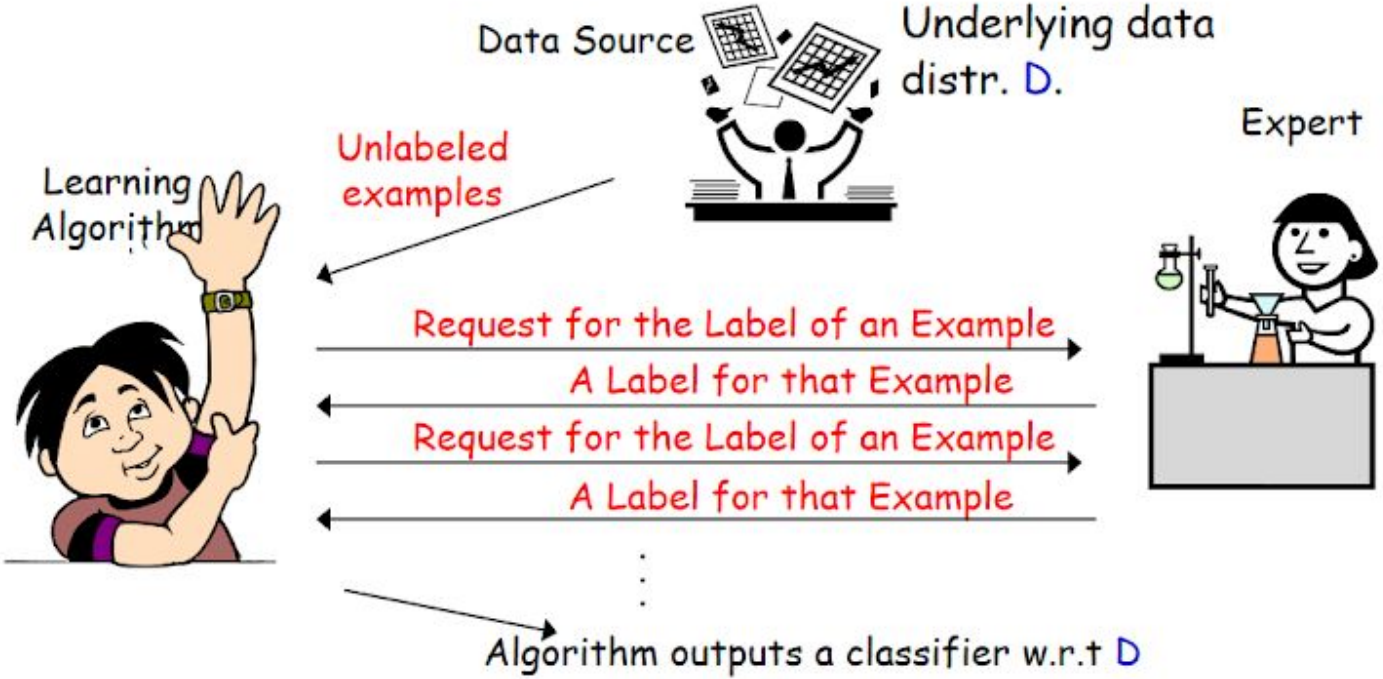

- Learner can choose specific examples to be labeled.
- Goal: use fewer labeled examples [pick informative examples to be labeled].

Motivating Example

- Threshold functions on the real line: $h_w = 1(x > w)$

Active algorithm: Get *N* unlabeled examples

- *Binary search* to find the correct threshold. *O(*log *n)* labels

Passive supervised learning: N labels to find OPT, Ω(1/ε) to get ε error

Active Learning: *O(*log *n)* labels to find OPT, *O(*log 1/ε*)* to get ε error

Common Technique in Practice

Active SVM seems to be quite useful in practice.

[Tong & Koller, ICML 2000; Jain, Vijayanarasimhan & Grauman, NIPS 2010]

Algorithm (batch version)

Input $S_u=\{x_1, \ldots, x_{m_u}\}$ drawn i.i.d from the underlying source D

Start: query for the labels of a few random x_is.

For $t = 1, \ldots.$,

- Find w_t the max-margin separator of all labeled points so far.
- Request the label of the example closest to the current separator: minimizing $|x_i \cdot w_t|$.
  (highest uncertainty)

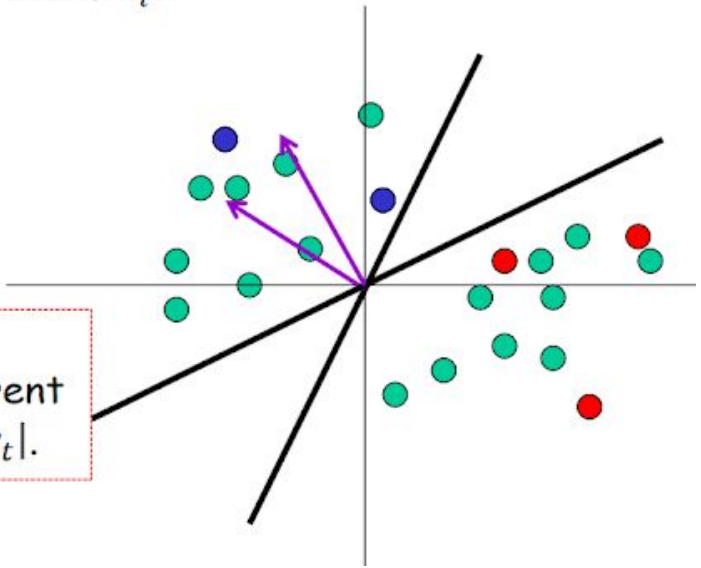

Active Learning in Practice

One caveat: sampling bias

- Bias created because of querying strategy
- As time goes on, the sample is less and less representative of the true source
- Sampling bias is observed in practice

Disagreement Based Active Learning [CAL92]

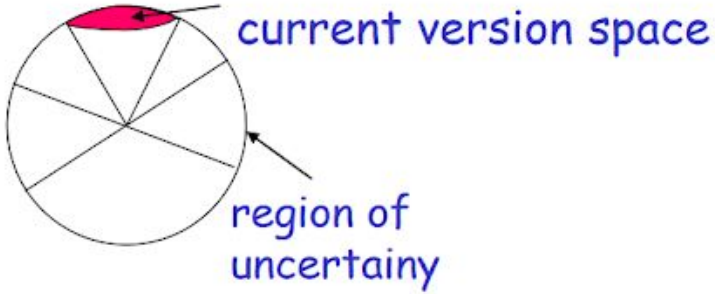

Algorithm:

Query for the labels of a few random x_is.

Let H_1 be the current version space.

For $t = 1, \ldots,$

Pick a few points at random from the current region of disagreement $\mathrm{DIS}(H_t)$ and query their labels.

Let H_{t+1} be the new version space.

Region of uncertainty [CAL92]

- Current version space: part of C consistent with labels so far.
- "Region of uncertainty" = part of data space about which there is still some uncertainty (i.e. disagreement within version space)

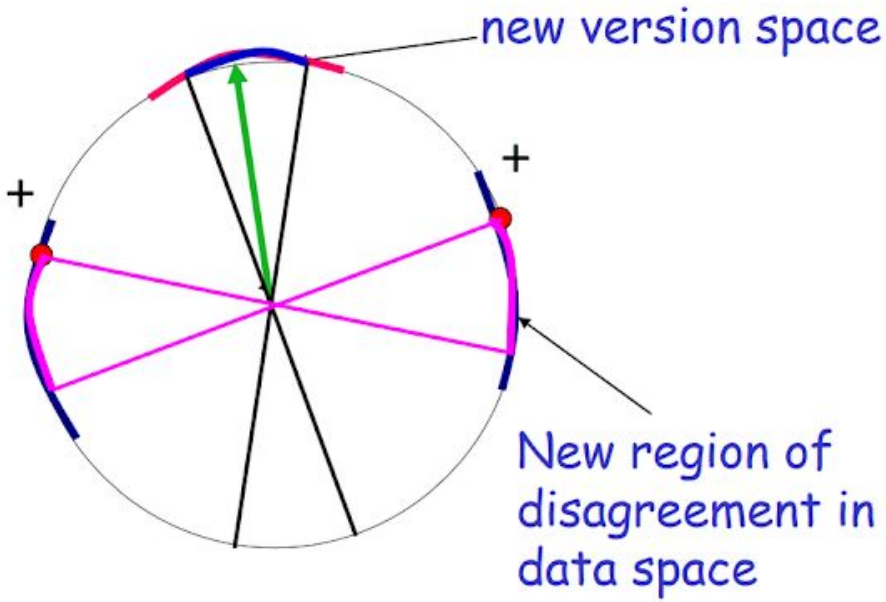

Practice Questions

In this problem you will design an active learning algorithm for finding a consistent linear separator passing through the origin when the data is on the unit circle in 2 dimensions. That is, given a dataset $S = \{x_1, \ldots, x_n\}$ with $\|x_i\| = 1$ for all $i = 1, \ldots, n$, your goal is to find a consistent classifier of the form $h(x) = \text{sign}(w^\top x)$. Assume we are in the realizable setting.

(a) [8 pts.] First, suppose that our data lies only on the *top half* of the circle (e.g., see Figure 3a). In 1–2 sentences, describe an algorithm for finding a consistent linear separator passing through the origin using $O(\log n)$ label queries. Hint: this problem is very similar to learning a consistent threshold function for data on the real line.

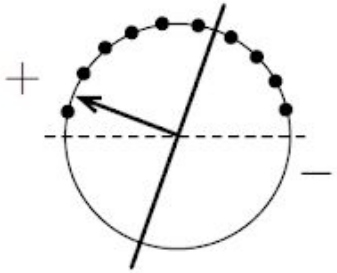

(a) Example data on the top half of the circle.

(b) [**Extra Credit 4 pts.**] Describe in 1–3 sentences how to extend your algorithm from part (a) so that it works for data anywhere on the circle. See Figure 3b for an example dataset. Hint: use the fact that for any point v $\text{sign}(w^{*\top}(-v)) = -\text{sign}(w^{*\top}v)$ (where w^* is the target weight vector) to try to reduce this case to the case studied in part (a).

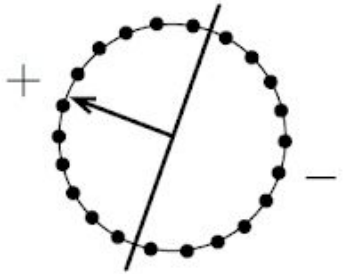

(b) Example data on the circle.

Practice Questions

- Assume the data lies in one dimension, and your goal is to find a consistent interval classifier of the form $h_{[a,b]}=1(a<x<b)$? Assume we are in the realizable setting. What is the smallest label complexity you can have?
- Assume the data lies in two dimensions, and your goal is to find a consistent two-dimensional threshold classifier of the form $h_{a,b}=1(x_1>a,x_2>b)$? Assume we are in the realizable setting. What is the minimum label complexity you can have?

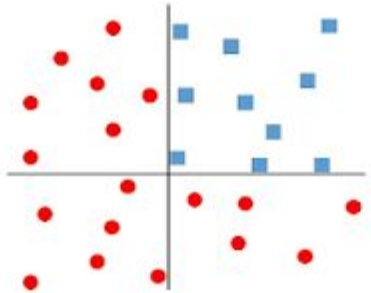

Practice Questions

- Now suppose we want to learn a class of rectangles in 2-D. A rectangle predicts positive on points inside and negative on points outside. Assume we are in the realizable setting.

  - What is the version space?
  - What is the region of disagreement?
  - Can we imply other labels?

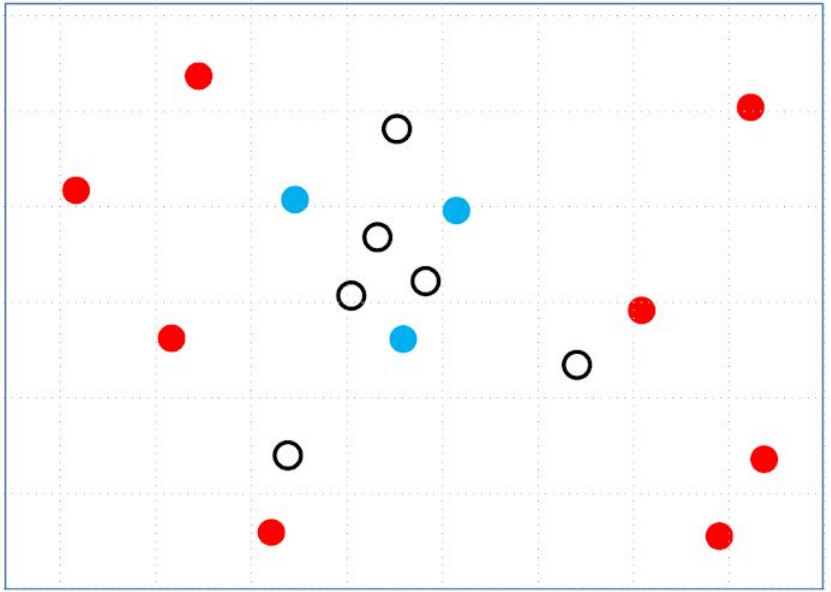

Questions?