

# The Two Faces of Active Learning

Sanjoy Dasgupta

University of California, San Diego

The active learning model is motivated by scenarios in which it is easy to amass vast quantities of unlabeled data (images and videos off the web, speech signals from microphone recordings, and so on) but costly to obtain their labels. Like supervised learning, the goal is ultimately to learn a classifier. But like unsupervised learning, the data come unlabeled. More precisely, the labels are hidden, and each of them can be revealed only at a cost. The idea is to query the labels of just a few points that are especially informative about the decision boundary, and thereby to obtain an accurate classifier at significantly lower cost than regular supervised learning.

There are two distinct narratives for explaining when active learning is helpful. The first has to do with efficient search through the hypothesis space: perhaps one can always explicitly select query points whose labels will significantly shrink the set of plausible classifiers (those roughly consistent with the labels seen so far)? The second argument for active learning has to do with exploiting cluster structure in data. Suppose, for instance, that the unlabeled points form five nice clusters; with luck, these clusters will be pure and only five labels will be necessary!

Both these scenarios are hopelessly optimistic. But I will show that they each motivate realistic models that can effectively be exploited by active learning algorithms. These algorithms have provable label complexity bounds that are in some cases exponentially lower than for supervised learning. I will also present experiments with these algorithms, to illustrate their behavior and get a sense of the gulf that still exists between the theory and practice of active learning.

This is joint work with Alina Beygelzimer, Daniel Hsu, John Langford, and Claire Monteleoni.