Active Machine Learning

Stephen Scott

What is Machine Learning?

- Computer programs that learn from experience
- Given labeled *training data*, infer a *hypothesis:* a function to label new (unlabeled) data
  - Applications: web page/text classification, spam filtering, bioinformatics, data mining, robot control, ...

Simple Example

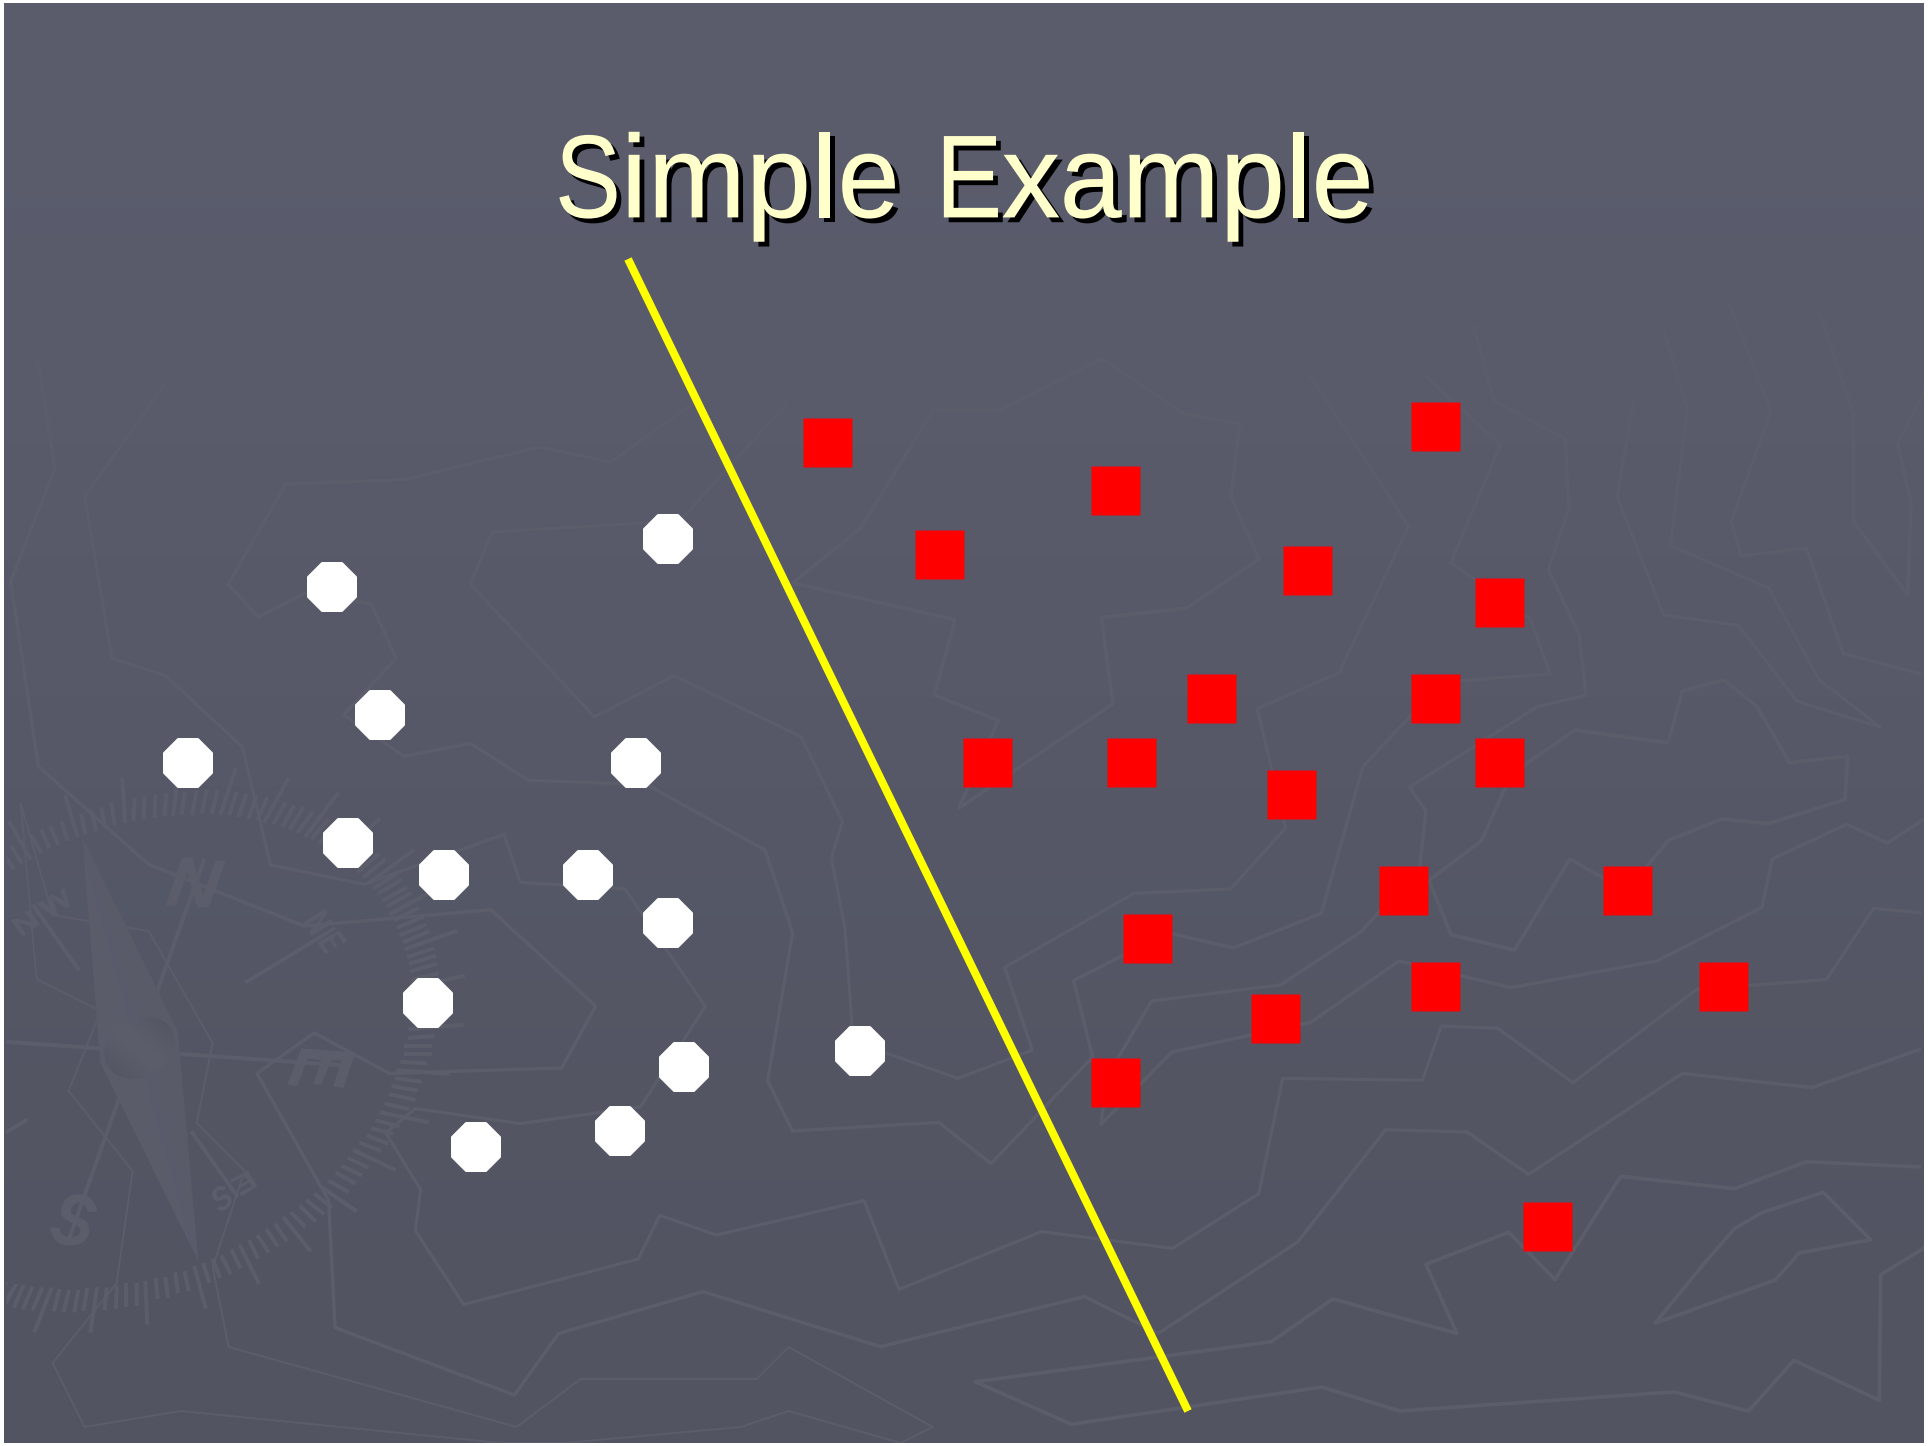

Where Does the Labeled Training Data Come From?

- Acquisition of data often easy, but typically need some benevolent teacher (e.g. human expert) who is patient enough to give the labels

Active Machine Learning

- But what if there are millions of examples (e.g. web pages) and the teacher is not patient enough to label more than a small fraction of your data?

  - Then the learning algorithm must decide *which* small subset of the data is to be labeled

$\Rightarrow$ Active Learning!

Active Machine Learning (cont'd)

- Could just select a random subset of the unlabeled data to label and train on
  - Naïve
  - Can we do better?
- The trick to active learning is to train on the labeled data you have and use the current hypothesis to choose the "best" unlabeled example to label next
  - Then retrain and repeat

Approaches

- Select example that the current hypothesis is least certain about (nearest the boundary)
- Select example that is expected to most increase generalization accuracy
- (And others ...)
- Our current work:
  - Sometimes perform *exploration* by occasionally choosing from far from the boundary rather than near boundary (ICDM 05)
  - Choose examples that are expected to most increase *area under the ROC curve (AUC),* a metric that is generally preferred to generalization accuracy