Machine Learning (in 20 minutes or less)

Doug Downey

Machine Learning

- "The study of computer programs that improve automatically with experience"
  T. Mitchell *Machine Learning*, 1998
- Used heavily in:
  - Bioinformatics, robotics, marketing/advertising, recommendations systems, information retrieval, fraud detection, handwriting/speech recognition, etc., etc…

Example Machine Learning Tasks

- How likely is person **x** to default on a loan?
- What is the location of robot **x**?
- Is document **x** about “baseball”?

Learning a function from examples

- **Given:** examples of a function f for various inputs $\mathbf{x}$:
  - $\{(\mathbf{x}^1, f(\mathbf{x}^1)), \ldots, (\mathbf{x}^n, f(\mathbf{x}^n))\}$
- **Goal:** Estimate f
  - Input $\mathbf{x} = (x_1, \ldots, x_d)$; individual features x_i
  - Output $f(\mathbf{x})$

- Probably the most common machine learning task formulation (though there are others)

Learn function from $\mathbf{x} = (x_1, \ldots, x_d)$ to $f(\mathbf{x}) \in \{0, 1\}$
given labeled examples $(\mathbf{x}, f(\mathbf{x}))$

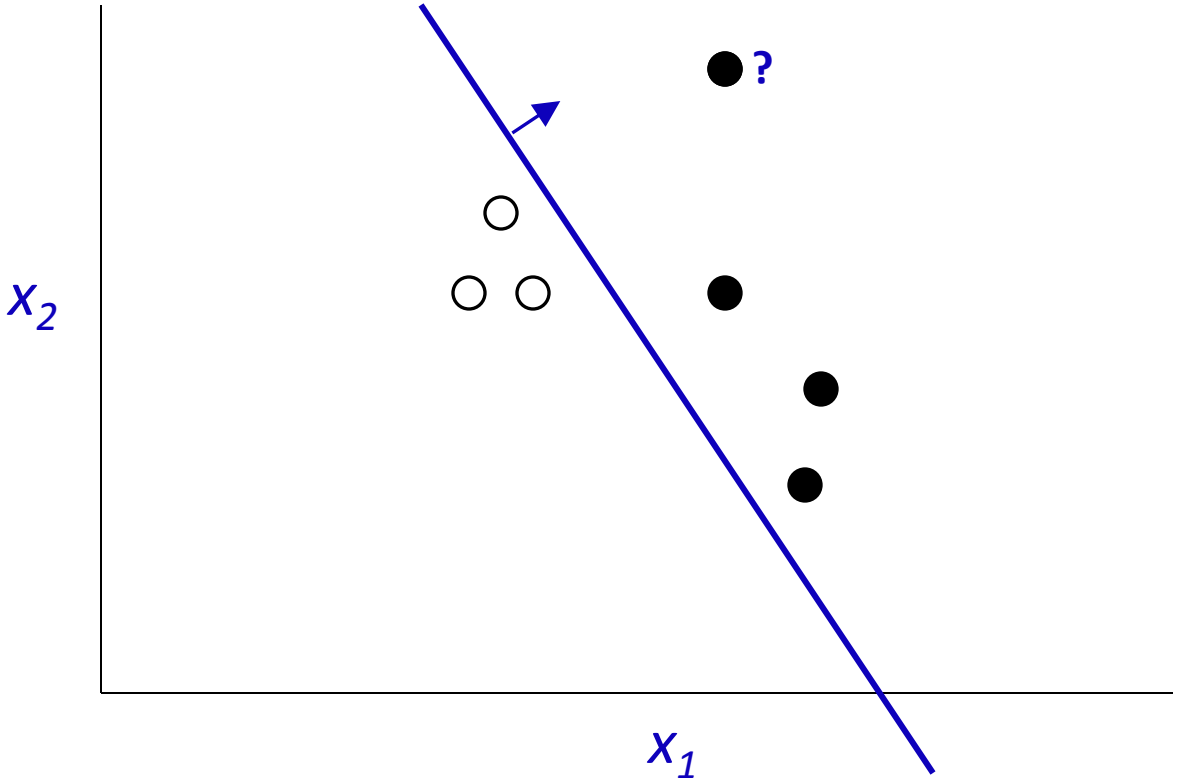

Representation

- In general, inputs and outputs can be
  - Nominal (e.g. Gender)
  - Ordinal (e.g. small, medium, large)
  - Numeric (e.g. Years of Education, probability of credit default, etc.)
- Predicting a nominal output: classification
  - Thus, predicting whether a document is about politics or sports is an instance of **Text Classification**
- Predicting a numeric output: regression (typically continuous)

Which classifier is best?

Learn function from $\mathbf{x} = (x_1, \ldots, x_d)$ to $f(\mathbf{x}) \in \{0, 1\}$
given labeled examples $(\mathbf{x}, f(\mathbf{x}))$

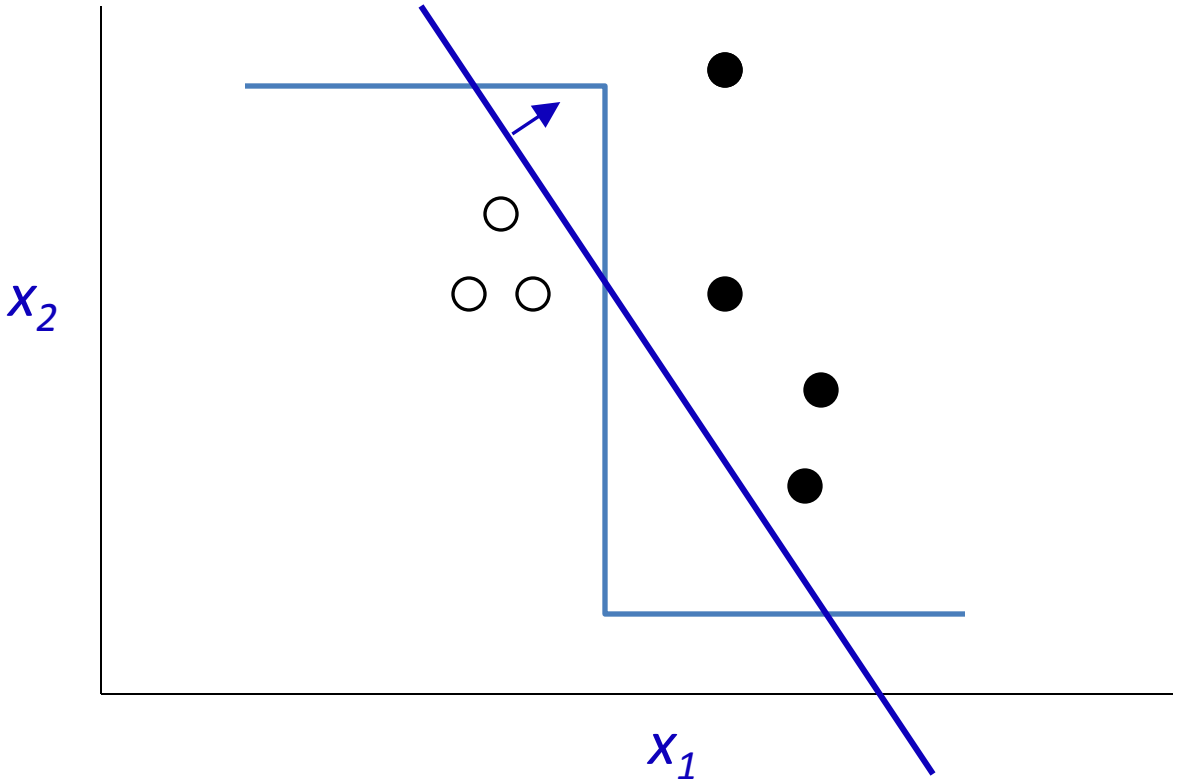

Which classifier is best?

Answer: you don’t know

Solutions:

1) try many and compare
2) Use domain knowledge

What does it mean for an ML algorithm to perform well?

- Metrics
  - Lots of possibilities
  - Classification: **accuracy**, precision, recall, cost, etc.
    - Accuracy = fraction of examples **x** where algorithm's predicted $f(\mathbf{x})$ matches true classification
  - Regression: mean squared error, etc.

What does it mean for an ML algorithm to perform well?

Learn function from $\mathbf{x} = (x_1, \ldots, x_d)$ to $f(\mathbf{x}) \in \{0, 1\}$
given labeled examples $(\mathbf{x}, f(\mathbf{x}))$

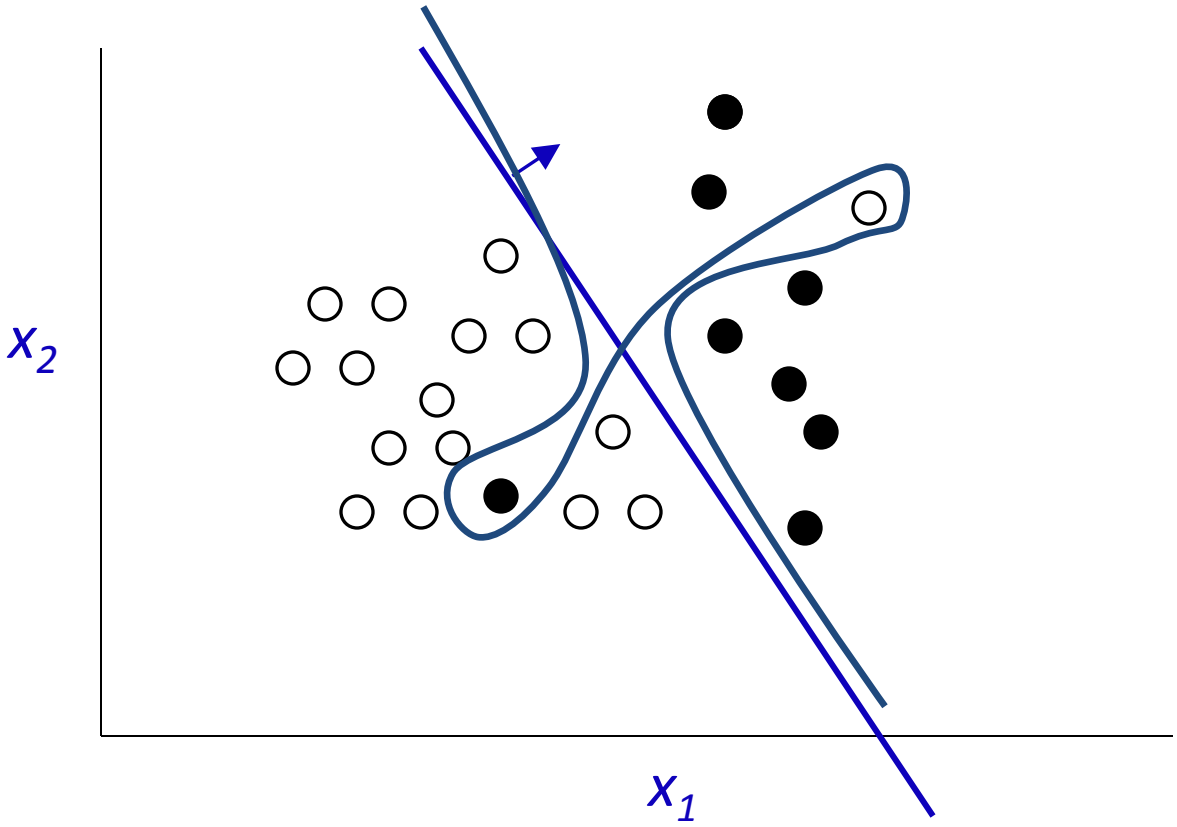

What does it mean for an ML algorithm to perform well?

- We want to know how our algorithm will perform on *new* inputs
  - So, test on a set of examples from disjoint from training (e.g. 80% train, 20% test)

How to do Machine Learning

1) Pick a feature representation for your task
2) Compile data
3) Choose a machine learning algorithm
4) Train the algorithm
5) Evaluate the results
6) *Probably: go to (1)*

How to do Machine Learning

1) Pick a feature representation for your task
2) Compile data
3) Choose a machine learning algorithm
4) Train the algorithm
5) Evaluate the results
6) *Probably: go to (1)*

Feature Engineering

- The art of machine learning
  - Features should be predictive and (relatively) independent
- How likely is person **x** to default on a loan?
  - FICO score
  - Income
  - Education Level
  - Assets
  - ~~Social Security Number~~
  - …

“Bag of words”