



Clustering

Introduction to Data Science

Jordan Boyd-Graber and Michael Paul

SLIDES ADAPTED FROM LAUREN HANNAH

Slides adapted from Tom Mitchell, Eric Xing, and Lauren Hannah

Roadmap

- Classification: machines labeling data for us
- Previously: logistic regression
- This time: SVMs
 - (another) example of linear classifier
 - State-of-the-art classification
 - Good theoretical properties

Thinking Geometrically

- Suppose you have two classes: vacations and sports
- Suppose you have four documents

Sports

Doc₁: {ball, ball, ball, travel}

Doc₂: {ball, ball}

Vacations

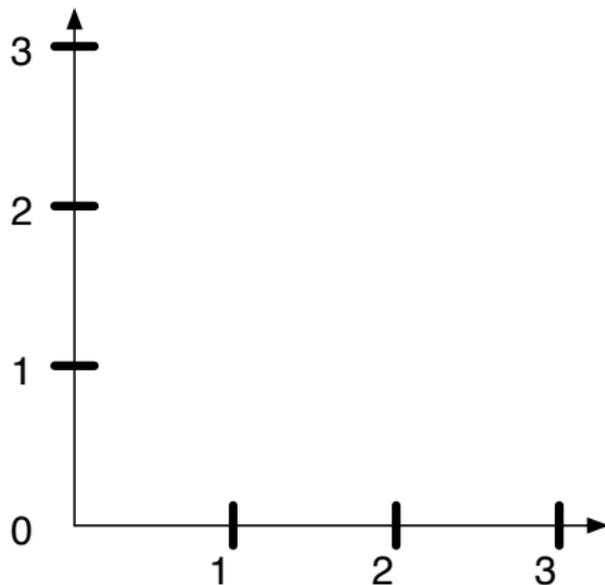
Doc₃: {travel, ball, travel}

Doc₄: {travel}

- What does this look like in vector space?

Put the documents in vector space

Travel



Ball

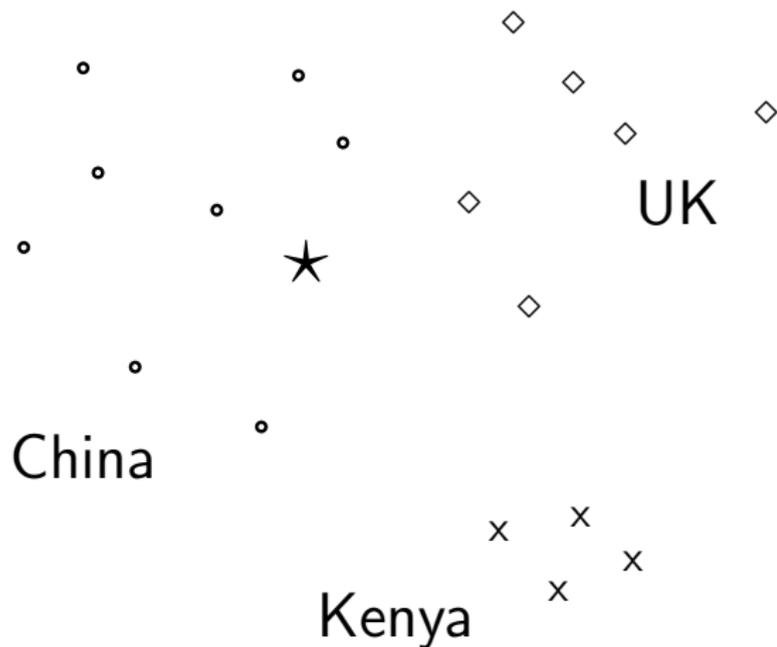
Vector space representation of documents

- Each document is a vector, one component for each term.
- Terms are axes.
- High dimensionality: 10,000s of dimensions and more
- How can we do classification in this space?

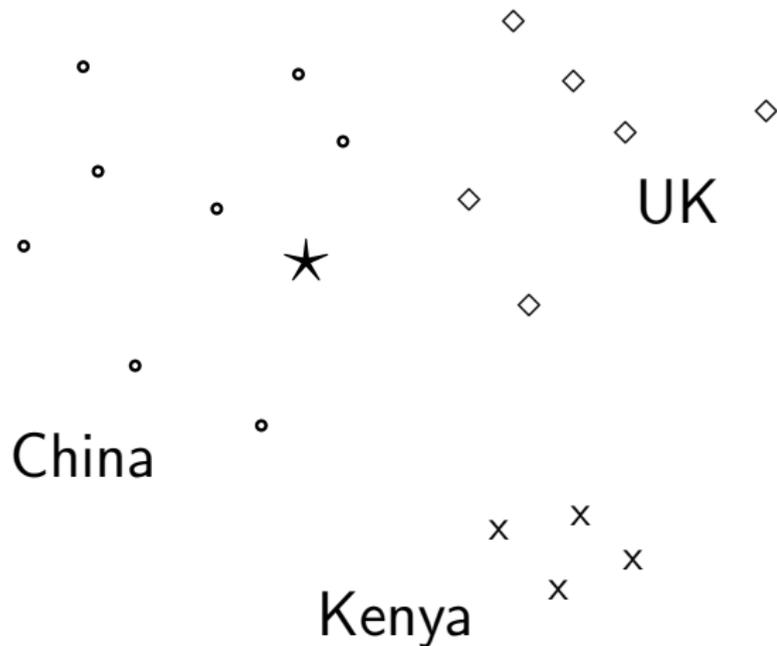
Vector space classification

- As before, the training set is a set of documents, each labeled with its class.
- In vector space classification, this set corresponds to a labeled set of points or vectors in the vector space.
- Premise 1: Documents in the same class form a **contiguous region**.
- Premise 2: Documents from different classes **don't overlap**.
- We define lines, surfaces, hypersurfaces to divide regions.

Classes in the vector space

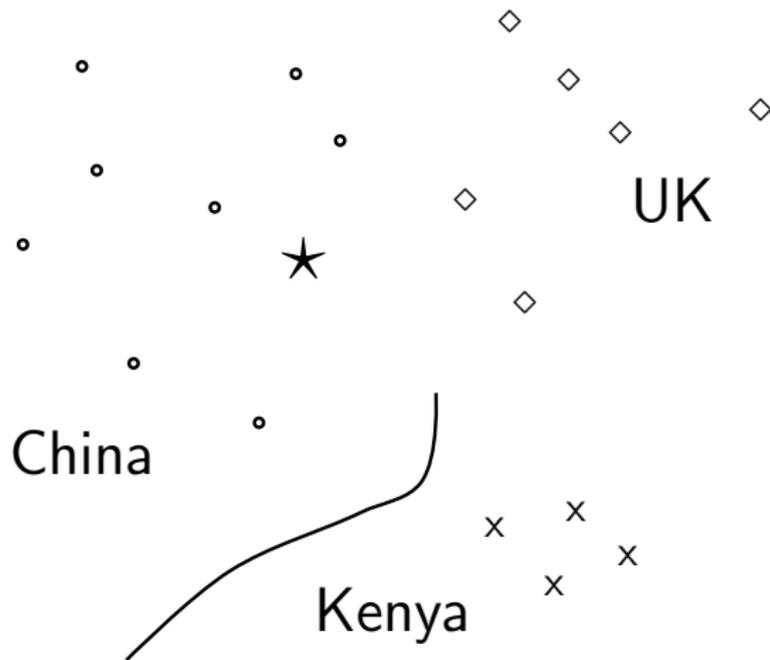


Classes in the vector space



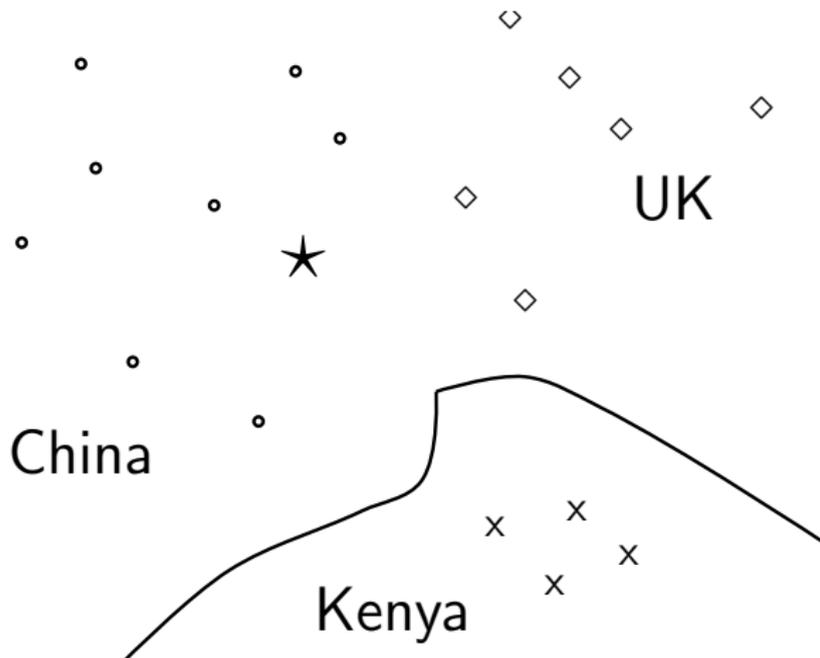
Should the document \star be assigned to China, UK or Kenya?

Classes in the vector space



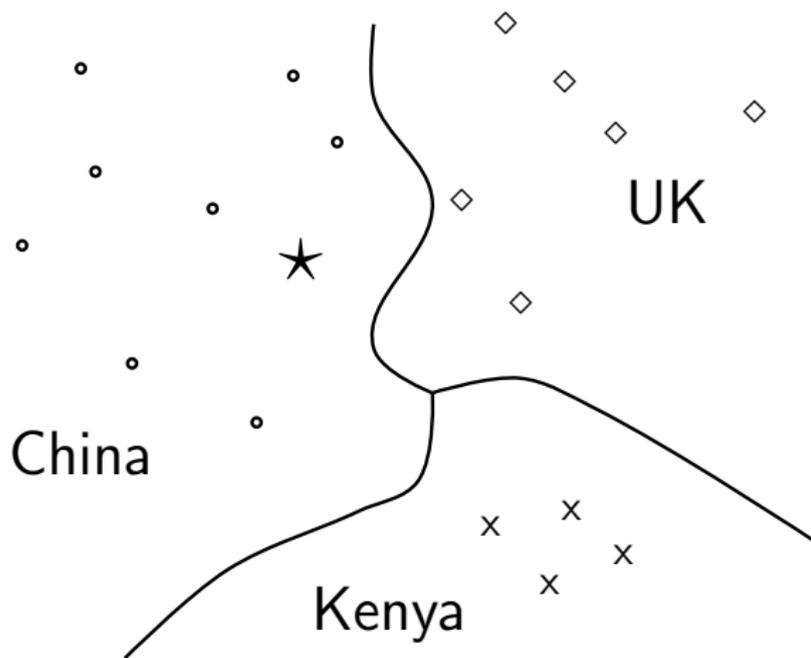
Find separators between the classes

Classes in the vector space



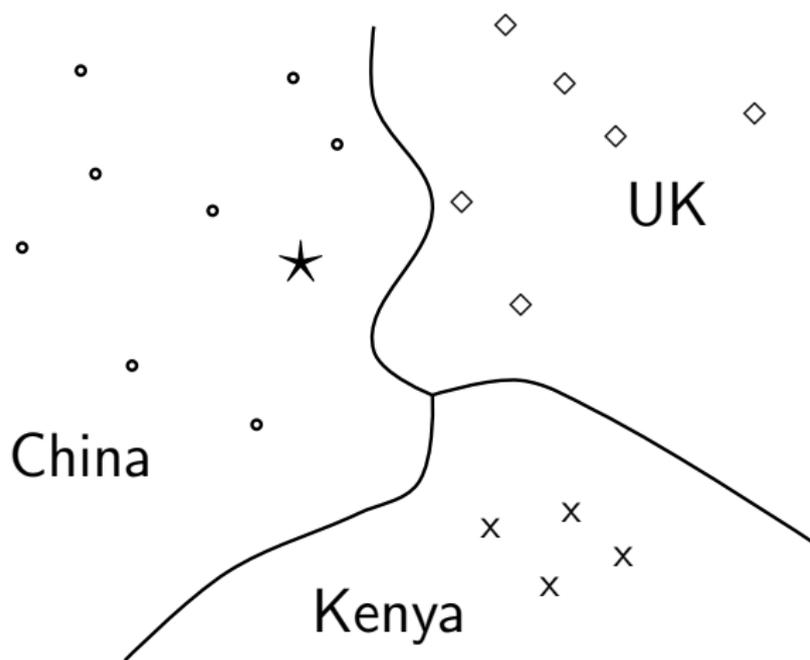
Find separators between the classes

Classes in the vector space



Based on these separators: \star should be assigned to China

Classes in the vector space



How do we find separators that do a good job at classifying new documents like \star ? – Main topic of today

Linear classifiers

- Definition:
 - A linear classifier computes a linear combination or weighted sum $\sum_i \beta_i x_i$ of the feature values.
 - Classification decision: $\sum_i \beta_i x_i > \beta_0$? (β_0 is our bias)
 - ... where β_0 (the threshold) is a parameter.
- We call this the **separator** or **decision boundary**.
- We find the separator based on training set.
- Methods for finding separator: logistic regression, linear SVM
- Assumption: The classes are **linearly separable**.

Linear classifiers

- Definition:
 - A linear classifier computes a linear combination or weighted sum $\sum_i \beta_i x_i$ of the feature values.
 - Classification decision: $\sum_i \beta_i x_i > \beta_0$? (β_0 is our bias)
 - ... where β_0 (the threshold) is a parameter.
- We call this the **separator** or **decision boundary**.
- We find the separator based on training set.
- Methods for finding separator: logistic regression, linear SVM
- Assumption: The classes are **linearly separable**.
- Before, we just talked about equations. What's the geometric intuition?

A linear classifier in 1D



- A linear classifier in 1D is a point x described by the equation $\beta_1 x_1 = \beta_0$

A linear classifier in 1D



- A linear classifier in 1D is a point x described by the equation $\beta_1 x_1 = \beta_0$
- $x = \beta_0 / \beta_1$

A linear classifier in 1D



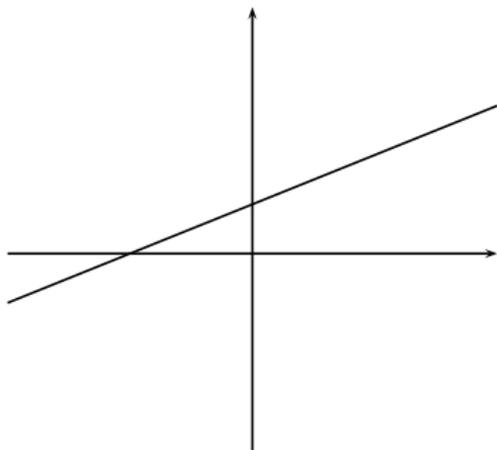
- A linear classifier in 1D is a point x described by the equation $\beta_1 x_1 = \beta_0$
- $x = \beta_0 / \beta_1$
- Points (x_1) with $\beta_1 x_1 \geq \beta_0$ are in the class c .

A linear classifier in 1D



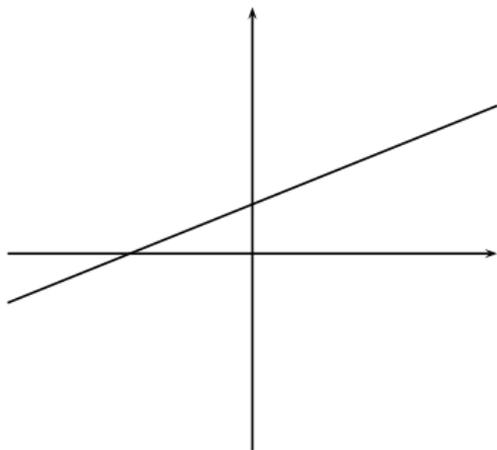
- A linear classifier in 1D is a point x described by the equation $\beta_1 x_1 = \beta_0$
- $x = \beta_0 / \beta_1$
- Points (x_1) with $\beta_1 x_1 \geq \beta_0$ are in the class c .
- Points (x_1) with $\beta_1 x_1 < \beta_0$ are in the complement class \bar{c} .

A linear classifier in 2D



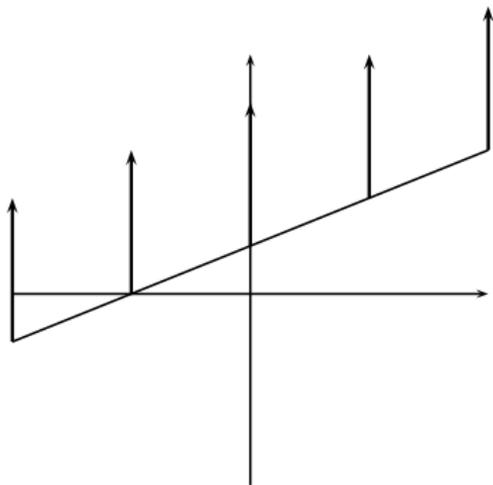
- A linear classifier in 2D is a line described by the equation $\beta_1 x_1 + \beta_2 x_2 = \beta_0$

A linear classifier in 2D



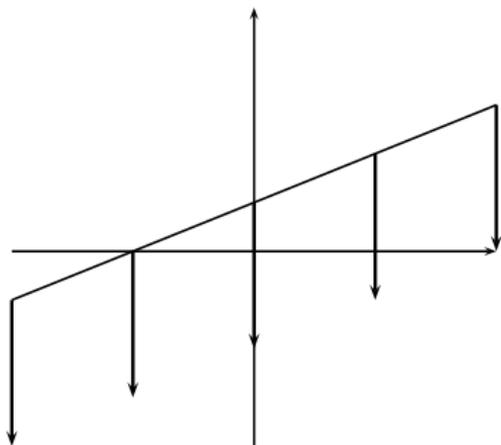
- A linear classifier in 2D is a line described by the equation $\beta_1 x_1 + \beta_2 x_2 = \beta_0$
- Example for a 2D linear classifier

A linear classifier in 2D



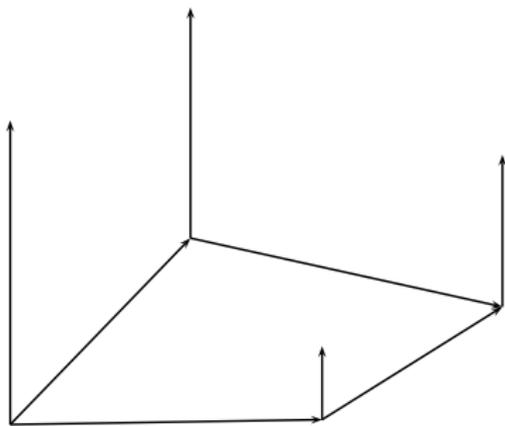
- A linear classifier in 2D is a line described by the equation $\beta_1 x_1 + \beta_2 x_2 = \beta_0$
- Example for a 2D linear classifier
- Points $(x_1 \ x_2)$ with $\beta_1 x_1 + \beta_2 x_2 \geq \beta_0$ are in the class c .

A linear classifier in 2D



- A linear classifier in 2D is a line described by the equation $\beta_1 x_1 + \beta_2 x_2 = \beta_0$
- Example for a 2D linear classifier
- Points $(x_1 \ x_2)$ with $\beta_1 x_1 + \beta_2 x_2 \geq \beta_0$ are in the class c .
- Points $(x_1 \ x_2)$ with $\beta_1 x_1 + \beta_2 x_2 < \beta_0$ are in the complement class \bar{c} .

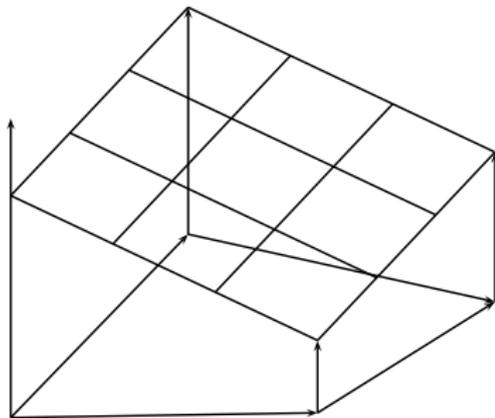
A linear classifier in 3D



- A linear classifier in 3D is a plane described by the equation

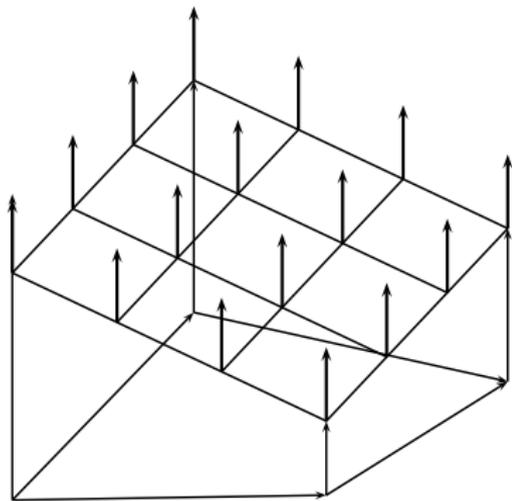
$$\beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 = \beta_0$$

A linear classifier in 3D



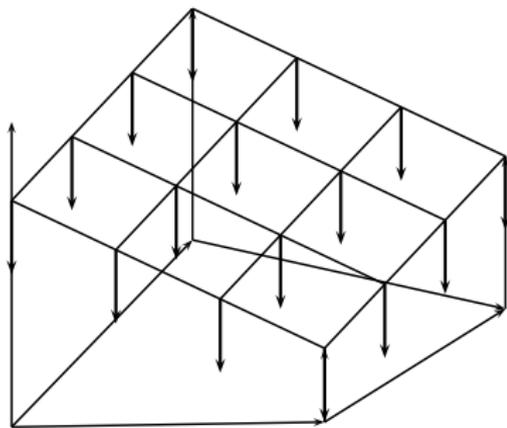
- A linear classifier in 3D is a plane described by the equation
$$\beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 = \beta_0$$
- Example for a 3D linear classifier

A linear classifier in 3D



- A linear classifier in 3D is a plane described by the equation
$$\beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 = \beta_0$$
- Example for a 3D linear classifier
- Points $(x_1 \ x_2 \ x_3)$ with
$$\beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 \geq \beta_0$$
are in the class c .

A linear classifier in 3D



- A linear classifier in 3D is a plane described by the equation
$$\beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 = \beta_0$$
- Example for a 3D linear classifier
- Points $(x_1 \ x_2 \ x_3)$ with
$$\beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 \geq \beta_0$$
are in the class c .
- Points $(x_1 \ x_2 \ x_3)$ with
$$\beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 < \beta_0$$
are in the complement class \bar{c} .