LIDAR

Noel Welsh

18 November 2010

Annoucements

- No lectures need week, as I expect you will all be working towards the demonstrations
- Robot things on the 27th: Circuit Bending and ROBOtic'10

Overview

- Today we're mostly going to look at *lidar*
- What is lidar? Light detection and ranging. Basically, it refers to laser range finders. The techniques are going to discuss apply equally to other range finders, like IR and sonar.

LIDAR

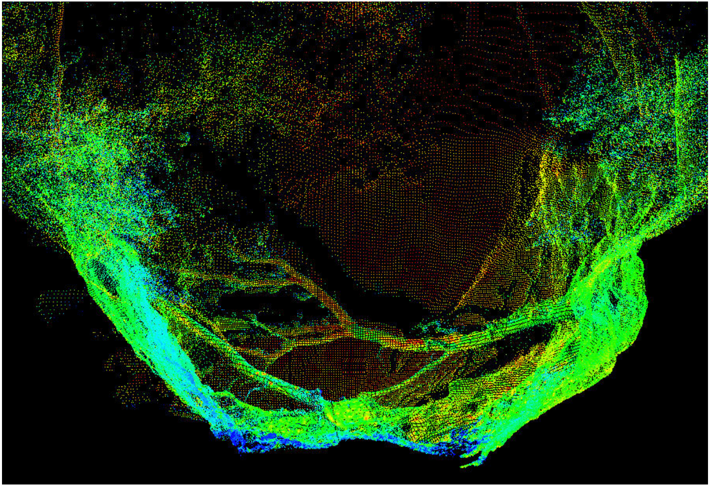

Figure: LIDAR Scan of Rodeo Creek. Taken by J. Toby Minear

LIDAR Details

- Typically we will deal with 2D data, but 3D scanners exist.
- The most common problems involve localisation
  - “Have I been here before”?
  - “What transformation best aligns this scan to a previous one”?

Review of Classification

- We haven't looked at classification in a while, so we're going to review it, looking at how you might build and test an object recogniser based on color detection.
- This is a massive hint on how you should present results in your writeup.
- Source code for the data and analysis will be available.

Data

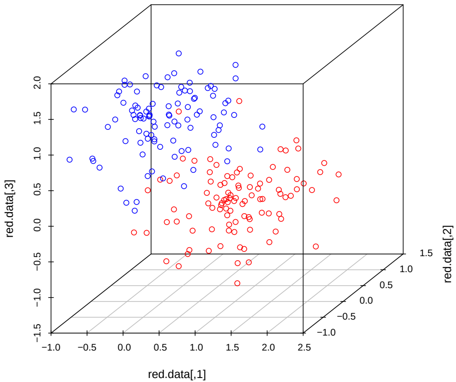

Figure: Example data from a colour detector

Problem

- You have 100 measurements each from a red and a blue object.
- The measurements consists of red, green, and blue values.
- Given this data how do we build and test an object recogniser?

My Solution

- Build a Naive Bayes classifier.
- Split data into test and training set (70/30 split)
- Test on the test data

Naive Bayes Classifier

- Two class classification problem
  - Classes are Red and Blue
  - Observations are RGB triples

$$P(Class|Data) \propto Likelihood \times Prior \quad (1)$$
$$P(Red|R,G,B) \propto P(R,G,B|Red)P(Red) \quad (2)$$
$$P(Blue|R,G,B) \propto P(R,G,B|Blue)P(Blue) \quad (3)$$

Solving the Naive Bayes Classifier

- Each number of samples for each class. Therefore priors the same. Can ignore.
- Assume each class is normally distributed.
  - I.e. assume the likelihood function is a normal distribution
- Problem reduces to estimating the mean and covariance of a 3D normal distribution from the training data for each class.
  - I.e. estimate the mean and covariance from the blue training data and from the red training data.
- Red sample mean = [1.08962796 0.05434472 -0.07618775]
- Blue sample mean = [-0.04437024 0.06112365 1.00726737]
- True means were [1, 0, 0] and [0, 0, 1]
- Covariances estimated similarly.

Testing the Classifier

- Calculate the posterior probability for each class on each of the testing data items.
- Choice the class with highest probability.
- On my data I get two errors from the 60 testing data items
  - Though I rigged the data to be easy to distinguish.
- Could then train again on the full data set, as we've seen the hold-out error is low. Didn't do this, but it is trivial.
- Code will available on the website soon.

LIDAR Techniques

- Recall the two problems
  - “Have I been here before”?
  - “What transformation best aligns this scan to a previous one”?
- The first is a classification problem.
- Techniques for the second can also be used to solve the first
- Two main classes of techniques:
  - Feature (also known as landmark or keypoint) extraction
  - Scan alignment

Feature Extraction

- Represent the data by a few distinctive features. E.g. sharp corners.
- Why?
  - Sparse representation of data
    - Faster search
    - Less storage
  - Features are cheap to compute

Scan Alignment

- Find a transformation (rotation and translation) to minimise distance between two scans.
- From this we can threshold the distance decide if it is a place we've seen before.
- Why?
  - Simple formulation of problem
  - Provides better metric information