Intelligent Audio Systems: A review of the foundations and applications of semantic audio analysis and music information retrieval

Jay LeBoeuf
Imagine Research
jay{at}imagine-research.com

July 2008

These lecture notes contain hyperlinks to the CCRMA Wiki.

On these pages, you can find additional supplement the lecture material found in the class - providing extra tutorials, support, references for further reading, or demonstration code snippets for those interested in a given topic .

Click on the  symbol on the lower-left corner of a slide to access additional resources.

WIKI REFERENCES...

Review from Day 4

- What are the 3 major components of a MIR system?
- Define Temporal Attack time and how it might be used to characterize a sound…
- What's "Spectral Brightness"?

k-means is an algorithm that does "hard clustering." What does that mean?

- How did the lab go?

Supervised vs. Unsupervised

- Unsupervised - “clustering”
- Supervised – binary classifiers (2 classes)
- Multiclass is derived from binary

Classification

- Classification – class labels (discrete or nominal)
- Regression – models continuous-valued function

80%

SVM

[draw on board]

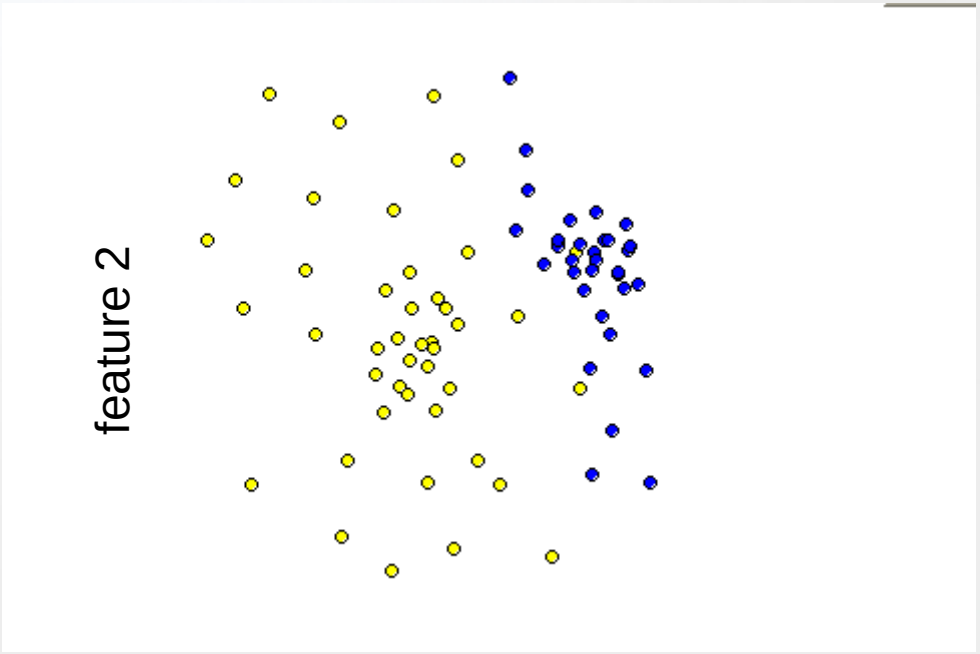
feature 2
feature 1

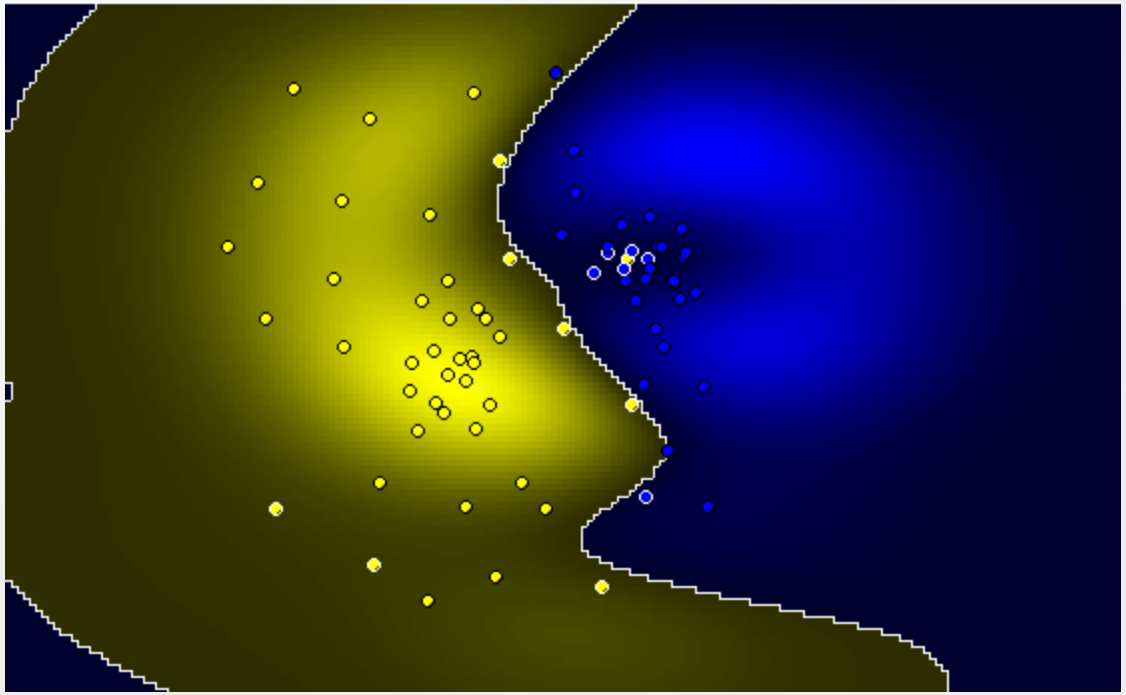

SVM

- Hyperplane separates the data from the two classes with a “maximum margin”.
- Support Vectors - are those data points that the margin pushes up against
- SVM training is guaranteed to find the global minimum of the cost function.
- Less experience needed - fewer parameters to tune

- >> svmdemo

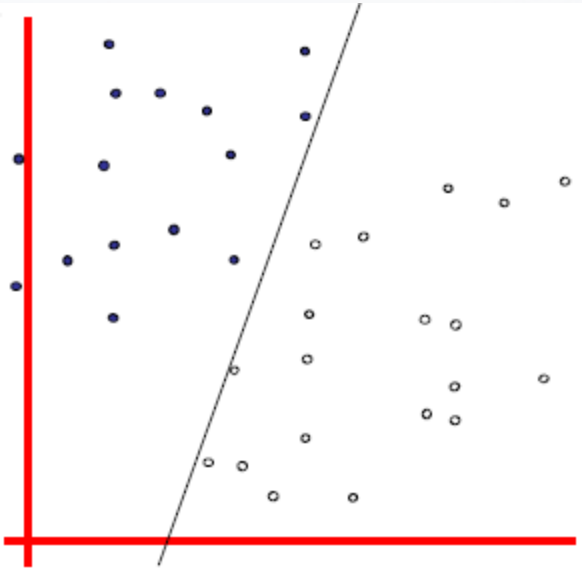

How would you classify this data?

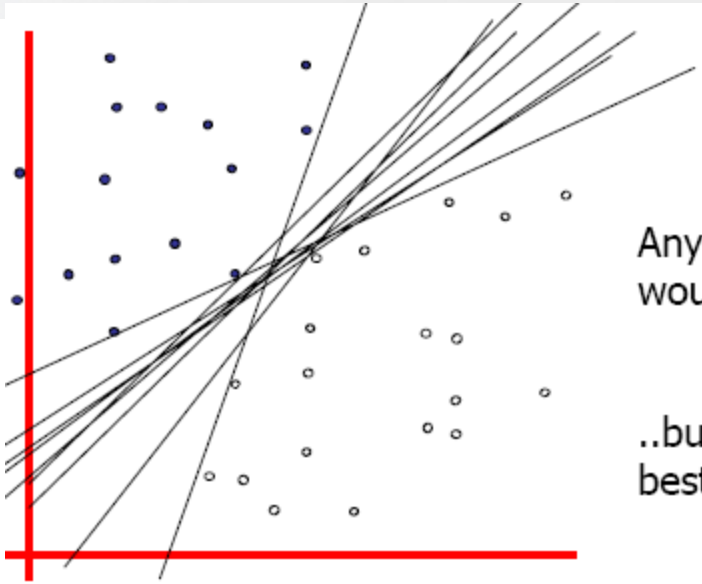

Any of these would be fine..

..but which is best?

From : http://www.autonlab.org/tutorials/svm15.pdf

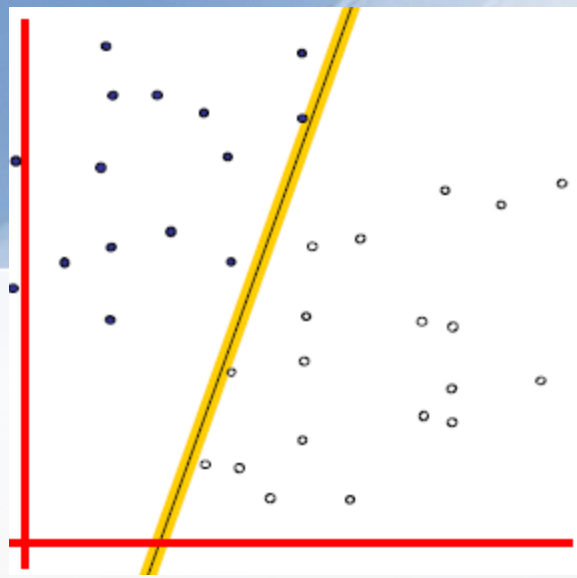

Define the margin of a linear classifier as the width that the boundary could be increased by before hitting a datapoint.

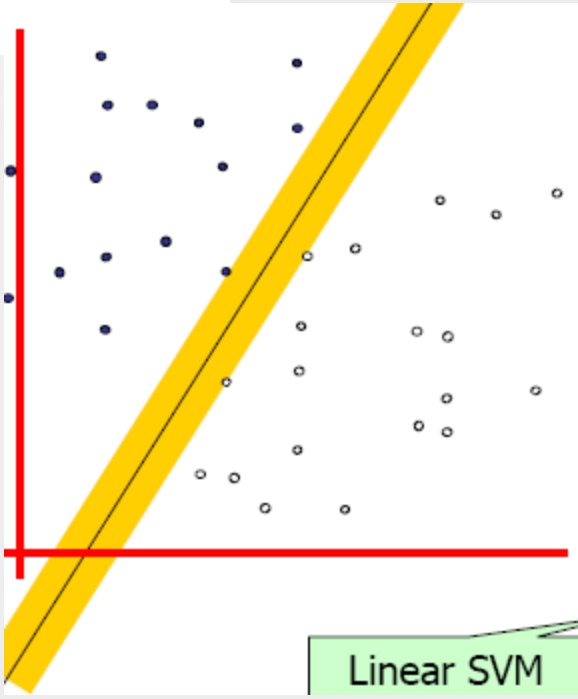

The maximum margin linear classifier is the linear classifier with the, um, maximum margin.

This is the simplest kind of SVM (Called an LSVM)

Linear SVM

From : http://www.autonlab.org/tutorials/svm15.pdf

SVM Parameters

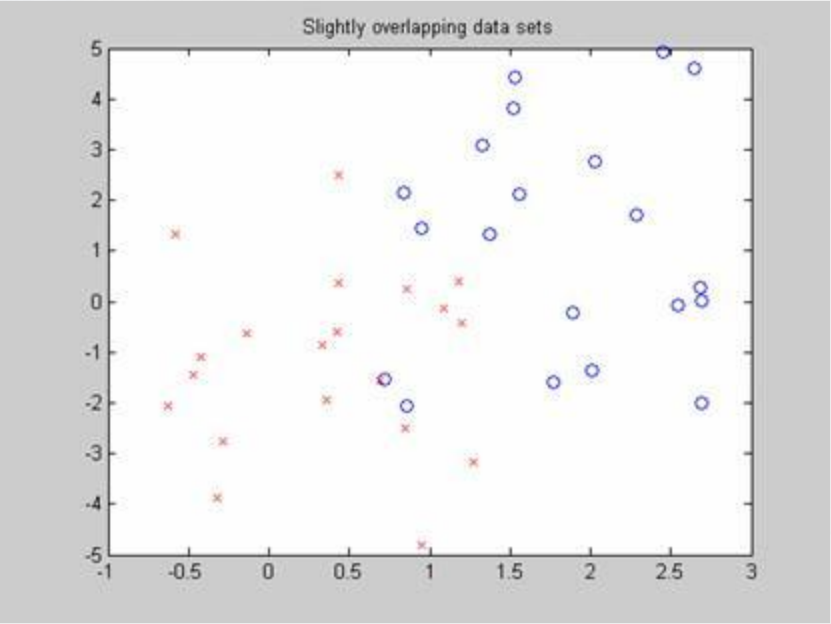

What effect do the parameters of an radial-basis-function SVM have on the separating the two data sets?

Using the RBF kernel, we have to choose values of :

gamma = degree of curviness of the hyperplane / complexity of the contour

C = allowance for points to overlap into each other's class

Video 1

Video 2

RBF Parameters: C and gamma

- Grid search using cross-validation to find the best one. Coarse then fine grid search.
- e.g., 2-5, 2-3, … 2+15, gamma = 2-15, 2-13, 2+3
- Why grid search
  - Psychological (If you *have time* for brute force… why chance it on approximations or heuristics)
  - Since there are only 2 params, grid search isn't all the different from advanced estimation techniques
  - Easily parallelized (C and gamma are independent)
- Large datasets
  - Random sample as approximation

SVM Parameters

- Whew!
- Grid search for finding the optimum parameters.
- You can manually tweak to reduce F+ or F- rate, but is generally not necessary or wanted.

- You can get approx. probability information, too. (Distance from the margin)

Practical Guide to SVM: The Lab

- Feature selection?
- Scale feature data
  - Save scaling stats so we can scale the test data to be in the same range
- Feature format
- Class labels {1,-1} or {0,1}
- Kernels (linear, polynomial, RBF, sigmoid)
- Find best C and gamma (cross-validation)
- Train with entire training set
- Test with validation or test set

- easy.py or grid.py

One-class SVM

- Binary classifiers rely on positive and negative examples of training data.
- One-class classifiers, however, only rely on positive examples. Great for models where the negative examples are not easily definable. (e.g., a classifier that detects "funky" sounds)
- Parameter: **ν** ("nu")

One-class SVM

- **ν** equals the % of training examples that you are willing to get wrong. (e.g., 10% error rate on training set is **ν** of 0.1)

SVM References

Libsvm and Libsvm Tools

- http://www.csie.ntu.edu.tw/~cjlin/libsvm/
- http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/

SVM Practical (How to get good results without cheating)

- http://www.kyb.tuebingen.mpg.de/bs/people/weston/svmpractical/