


- Announcements:
- Today:
 - Project intro
 - Wrap up SVM and do demo
- Friday: SVM lab
- Next week:
 - Monday: Bayesian classifiers
 - Tuesday: lightning talks
 - Thursday: Mid-term exam
 - Friday: sunset detector lab

Demo

- Software courtesy of <http://ida.first.fraunhofer.de/~anton/software.html> (GNU public license)
- Preview of Lab 4 (posted):
 - Download the Matlab functions that train and apply the SVM.
 - The demo script contains examples of how to call the system
 - Write a similar script to classify data in another toy problem
- Directly applicable to sunset detector

Kernel functions

- Note that a hyperplane (which by definition is linear) in the feature space = a nonlinear boundary in the input space
 - Recall the RBFs
- Note how choice of σ affects the classifier



Comparison with neural nets

- Expensive

- Training can take a *long* time with large data sets. Consider that you'll want to experiment with parameters...
- The classification runtime and space are $O(s)$, where s is the number of support vectors.
- In the worst case, s = size of whole training set (like nearest neighbor)
- But no worse than implementing a neural net with s perceptrons in the hidden layer.
- Empirically shown to have good generalizability even with relatively-small training sets and no domain knowledge.

Preview of demsvm2 (if time)

- Shows differing values of C for positive and negative examples.
- Part of tomorrow's lab

