

Reading #2: SVMs

Reading #2

- Manning et al 2008: ch 15
 - 15.1: SVM: the linearly separable case
 - 15.2: extensions to the SVM model
 - Soft-margin classification
 - Multi-class SVMs
 - Nonlinear SVMs

- Due: Feb 18 at **11AM**

Questions

- Q1: What is a hyperplane? What does each axis in the feature space represent?
- Q2: What does SVM try to optimize?
- Q3: What is a kernel? What's the benefit of using kernels?
- Q4: What is soft-margin? Why do we need it?