Boosted Image Classification: An Empirical Study

Nicholas R. Howe
Smith College
nhowe@cs.smith.edu

Never the Twain Shall Meet?

Machine Learning

- *Improved performance through boosting & other large-margin techniques.*

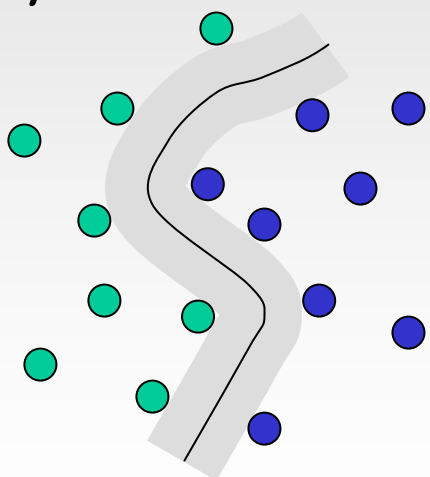

Image Comparison

- *Improved performance through better, more comprehensive image representations.*

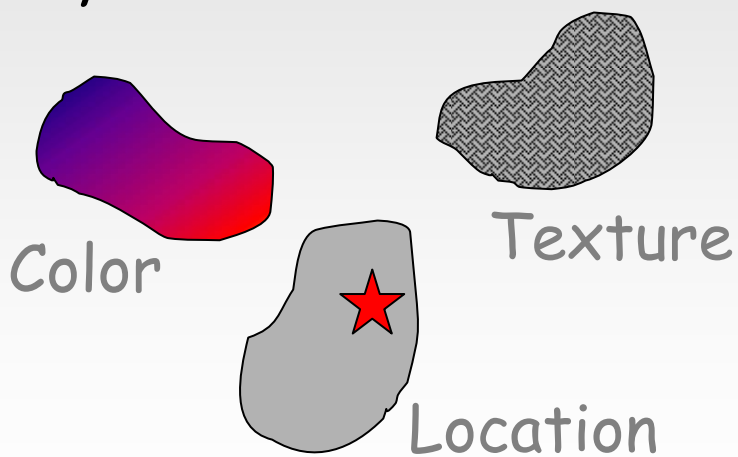

Previous Work

- Tieu and Viola (2000)
  - A good start, but limited
  - Looks at just one candidate image representation
  - Simple, feature-based boosting
    (i.e., decision stumps)
- Need for more comprehensive investigation

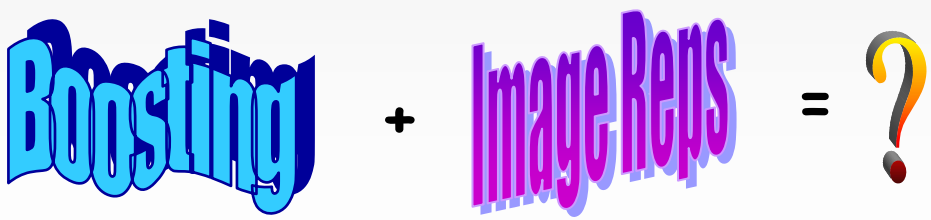

Image Classification is Hard

- Classes are diffuse.
- Features correlate weakly with class.
- High dimension (10K+)

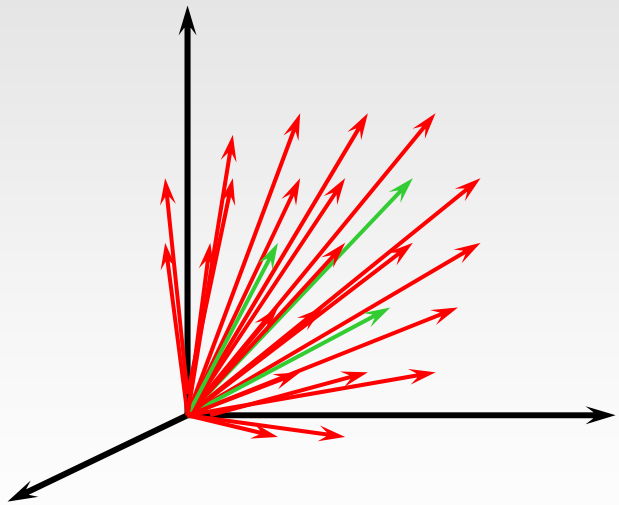

Two Goals of This Work

- Try different ways to apply boosting (i.e., different base classifiers)
- Test boosting with different image representations

Review of Boosting

- Base classifier must score **>50%** on arbitrarily weighted training set.
- Train base classifier using multiple weightings of training data.
- Combined predictions better than single classifier alone.

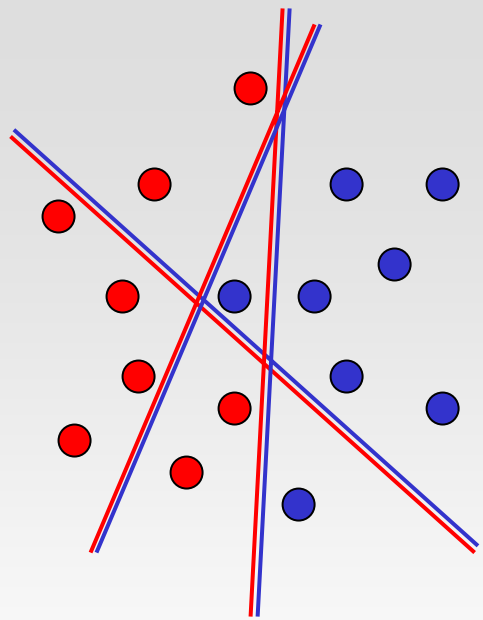

Options for a Base Classifier

- Many standard classifiers are "feature-based".
  (Decision boundaries orthogonal to feature axes.)
- "Vector-based" classifier may suit images better.
  (Decision boundaries are neighborhood around a vector.)

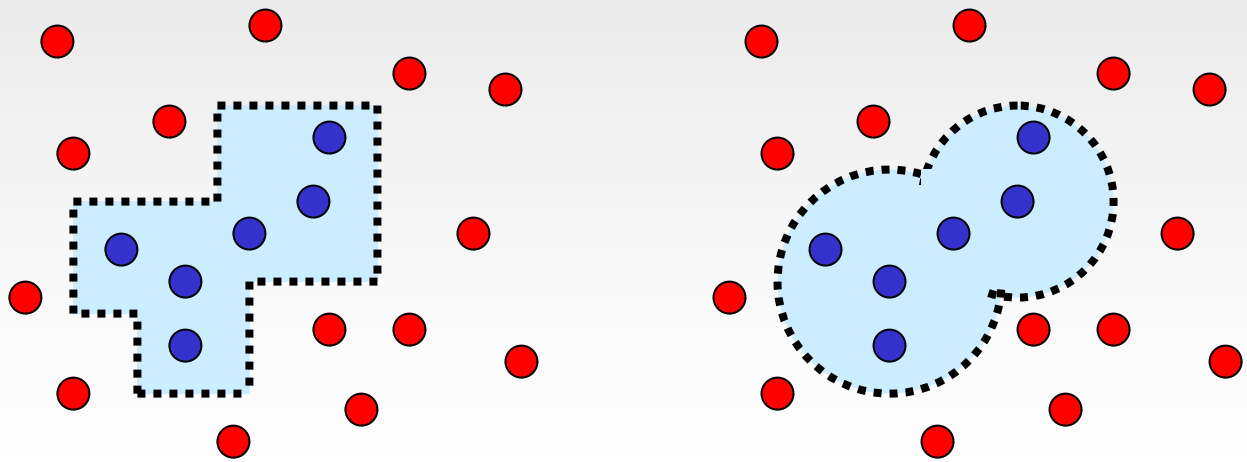

Vector-Based Classifier

$V_p = \sum$ positive instances

$V_n = \sum$ negative instances

$$V_\perp = V_p - \frac{V_n \bullet (V_p + V_n)}{\|V_p + V_n\|}$$

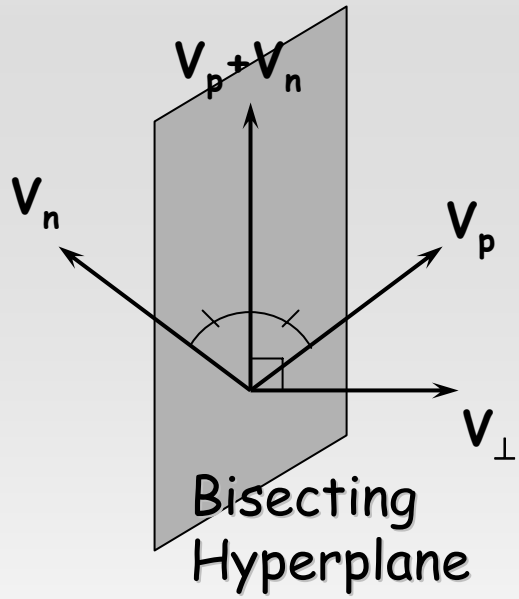

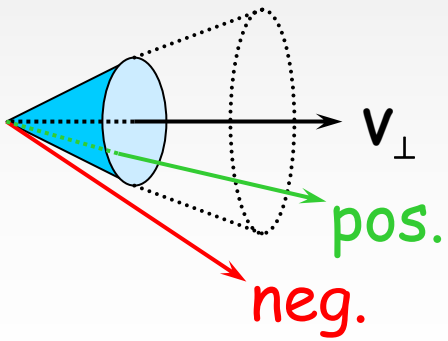

←Instances within some angular radius of $V_\perp$ are classified as positive.

Image Representations

- Correlogram (Huang et. al.)

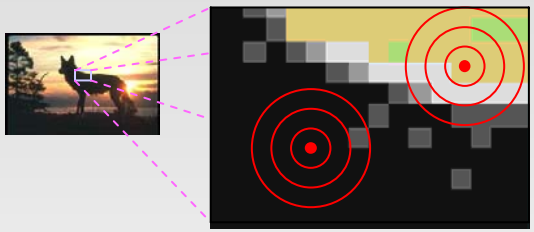

- Stairs (Howe & Huttenlocher)

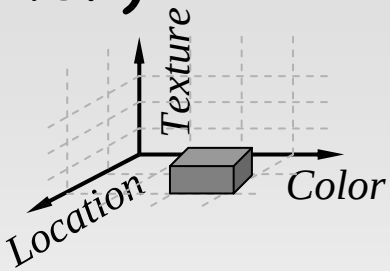

- Tieu-Viola

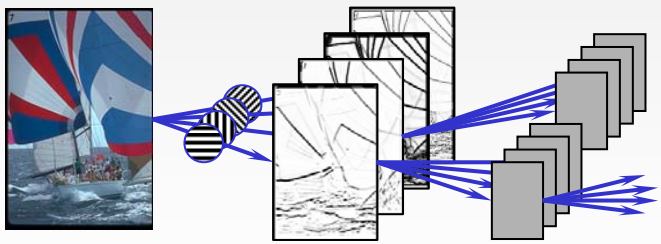

- Histogram (Swain & Ballard)

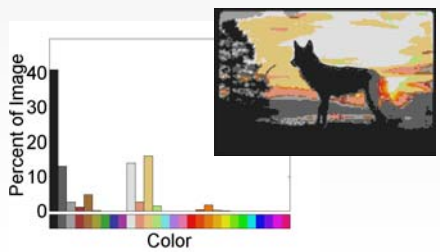

Evaluation Mechanism

- 20K images (Corel)
- 5 categories
- 5x2 cross validation
- Unboosted control:
  k-Nearest Neighbor (kNN)




$\Rightarrow$ ROC curves

Comparison based on area under curve.

Comparison: Image Reps

- Correlograms do best, T-V worst.

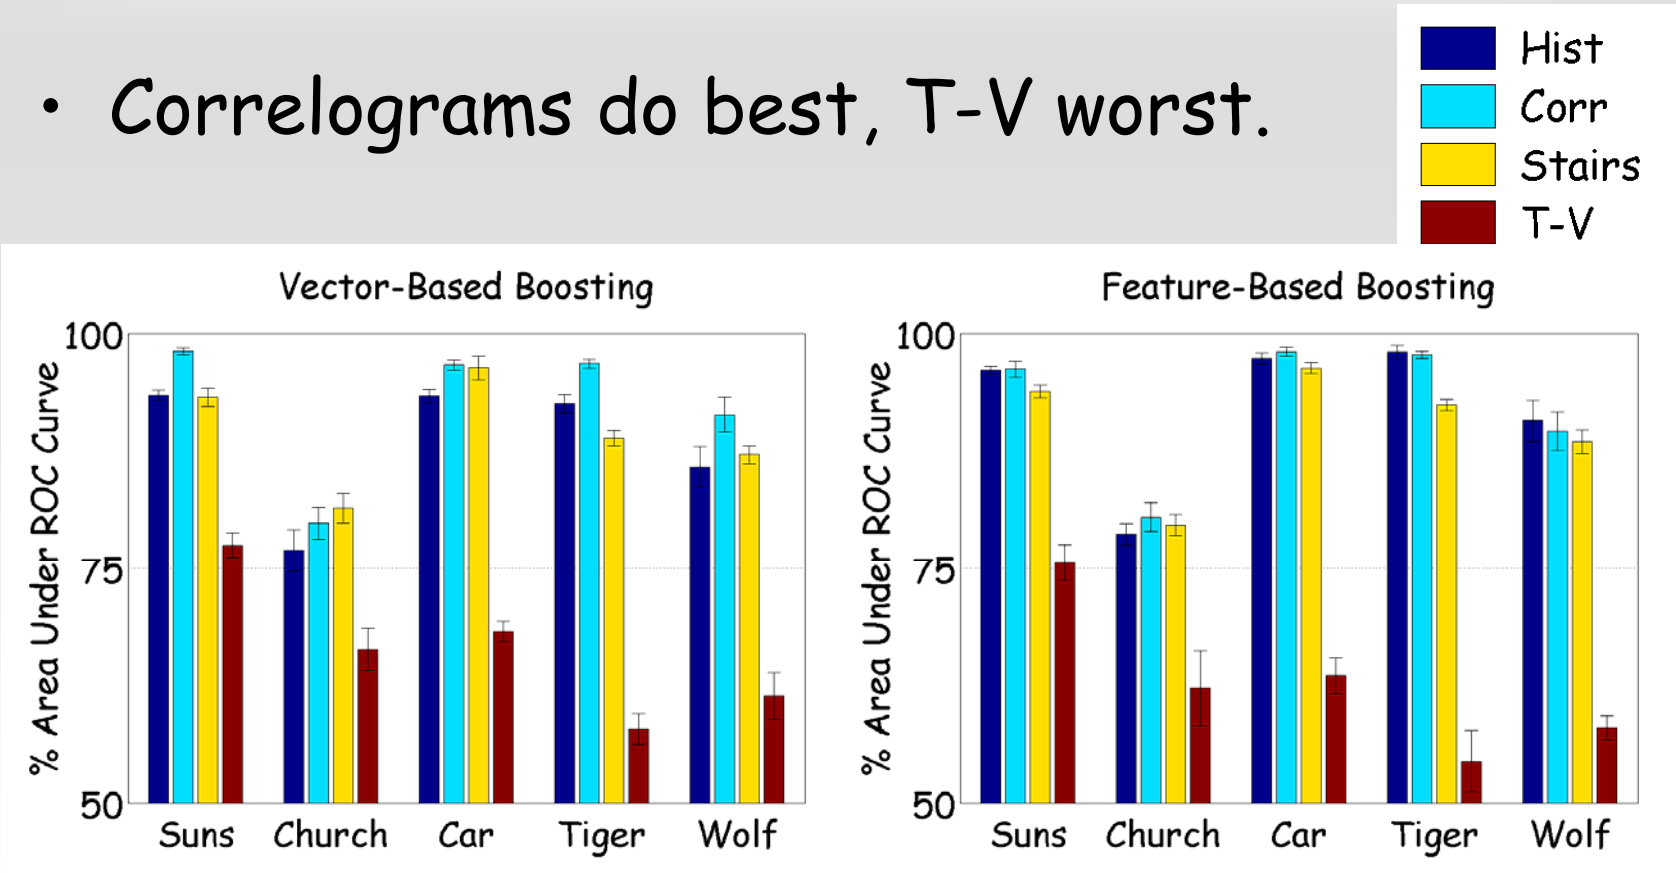

Comparison: Base Classifier

- Best method varies with size of feature space.

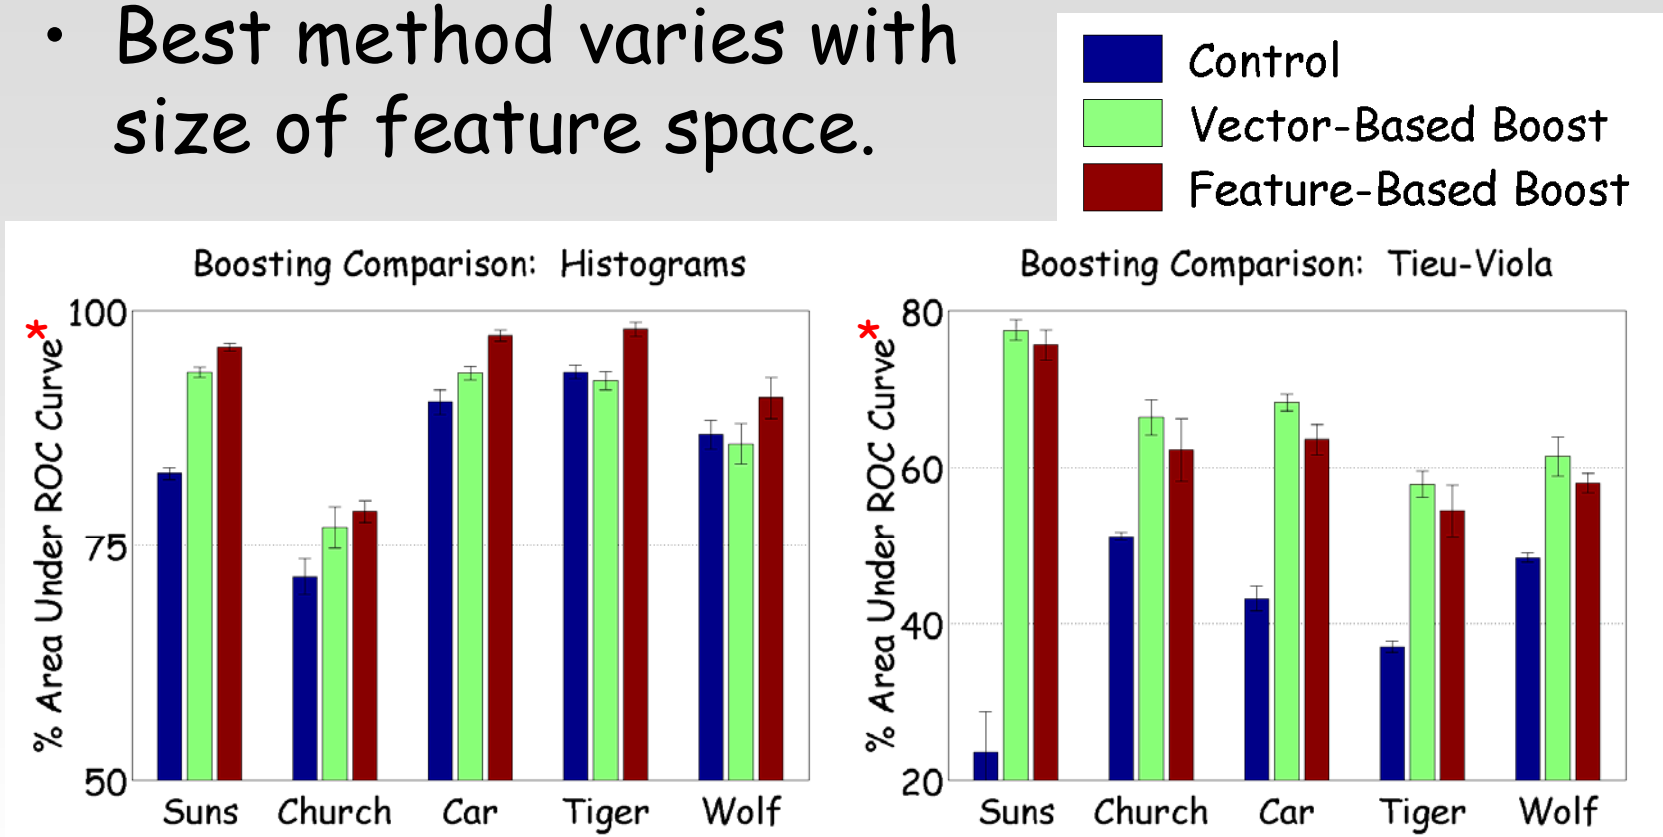

* Note differing y axes

More on Base Classifier

- Mid-sized feature spaces show fewer trends.

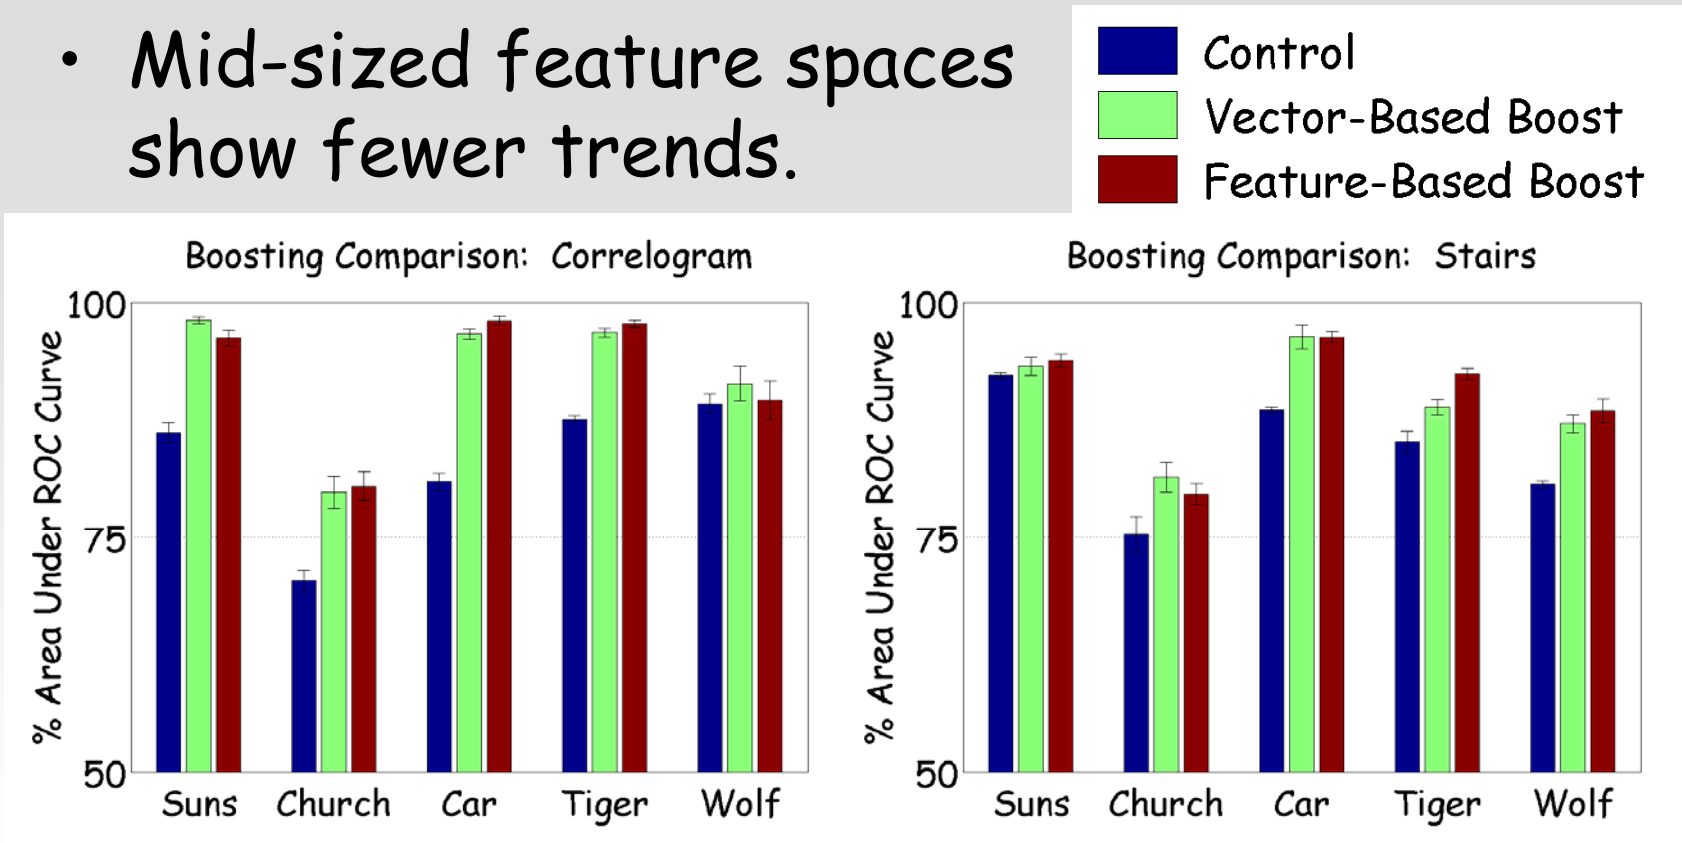

Conclusion

- Boosting works with a range of image representations. (No surprise!)
- Boosted correlogram is most successful representation.
- Best base classifier varies with size/complexity of feature space.