One-shot Recognition

7/17

- Kevin Tang (Cornell)
- Marshall Tappen (UCF)
- Rahul Sukthankar (Intel/CMU)

This Week

- Obtained actual images from Lampert
- Understand data better
  - Varying parameters to create plots
  - SVM test on hold out sets
- Redo previous experiments
  - Using all features (cq, lss, phog, rgsift, sift, surf)
  - Using all images
- Refined one-shot clustering for new results

Image Examples



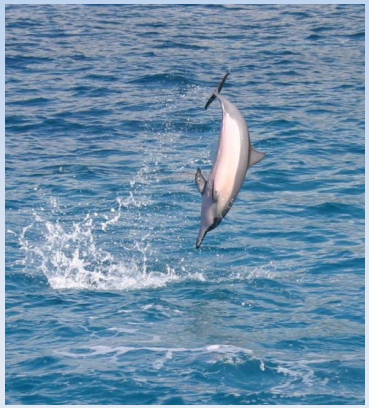


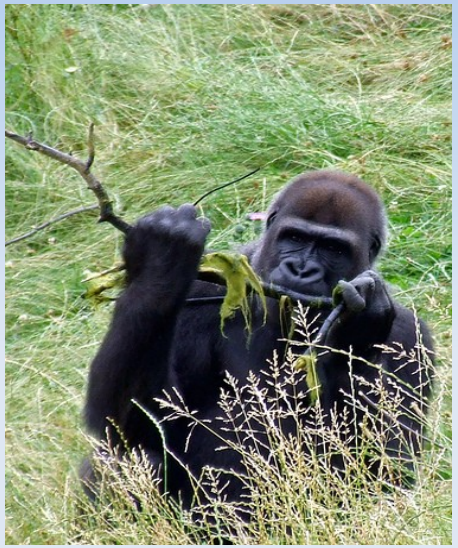


Image Examples







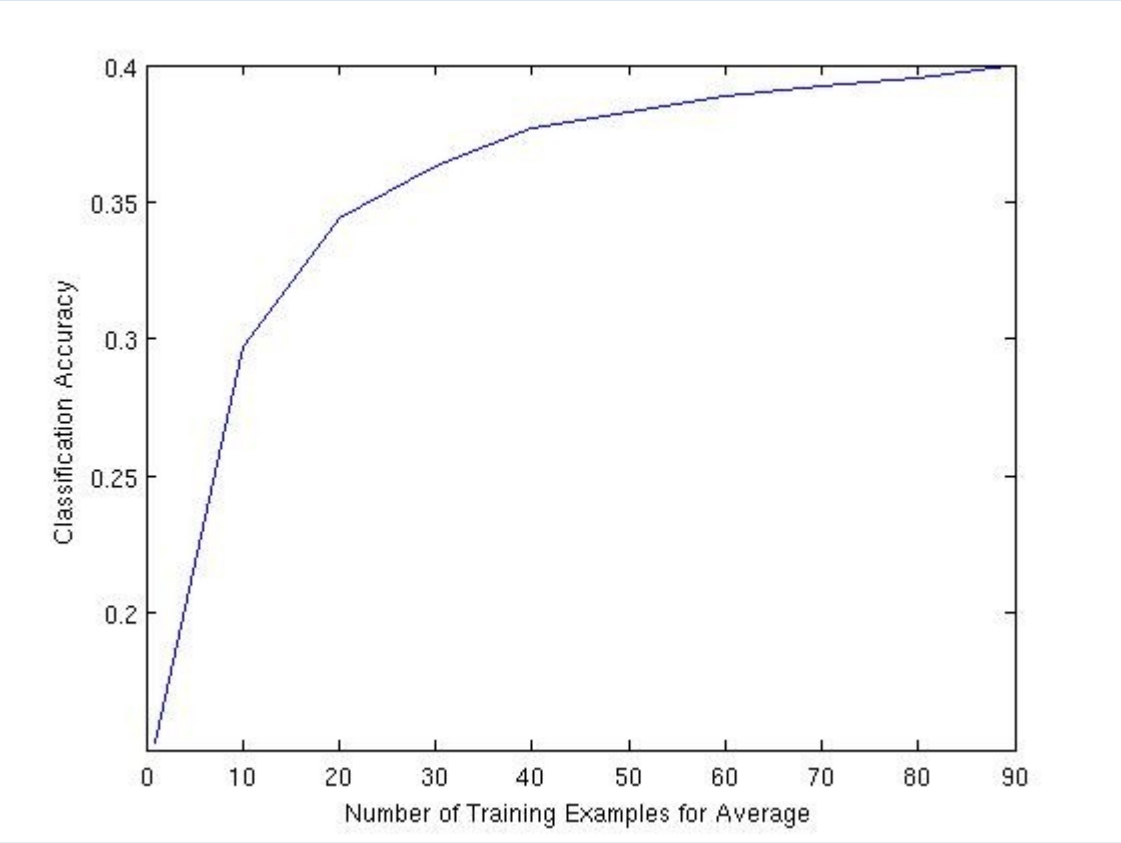






Understanding the Data

- How does the classification accuracy vary with the number of examples we average to create the training example?



Understanding the Data

- How does the number of images used affect the SVM performance?
  - Used different numbers of training images
  - Tested SVM's on a hold-out test set
  - Found that performance increases by ~10%
- Previously used 80 images from each class
  - 80 * 50 = 4,000 images
  - Total of ~30,000 images in the dataset
  - Should draw from all images

New Experiments

- With magnitude of data in mind, used all data and all features
- Concatenate features together
  - Results in a ~10,000 dimension vector
- Learn logistic projection
  - PCA down to 500 dimension for run-time purposes
  - Logistic projection from 500 down to 85 dimensions
- 1v1 SVM's
  - Throw all the data into the SVM's

Results

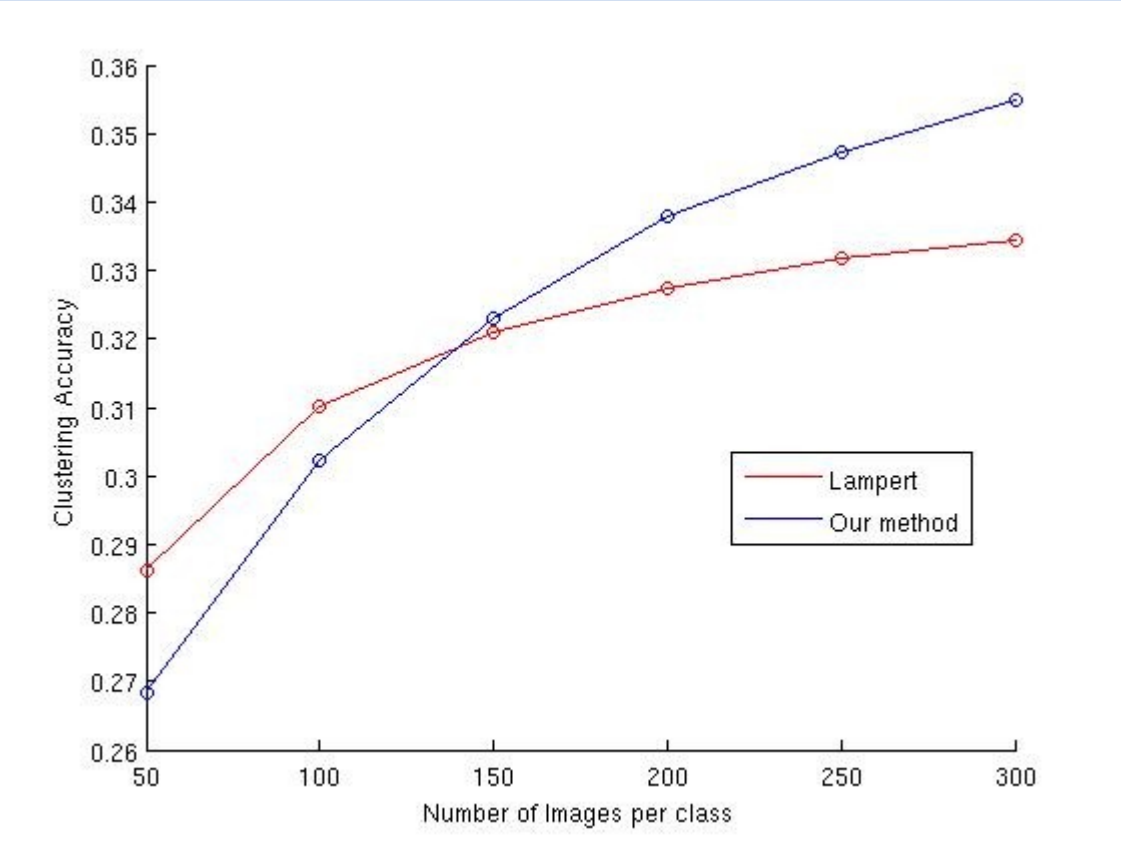
- One-shot classification (Logistic projection)
  - Chance Accuracy: 10%
  - Lampert's Accuracy: ~30.75% mean, ~43.48% max
  - Our Accuracy: ~27.41% mean, ~51.00% max
- One-shot classification (1v1 SVM's)
  - Our Accuracy: ~22.70% mean, ~45.00% max
  - Linear kernels, Platt scaling

Results

- One-shot clustering (all images)
  - Chance Accuracy: 10%
  - Lampert's Accuracy: ~37.64% mean, ~38.50% max
  - Our Accuracy: ~46.28% mean, ~48.04% max
- Problem with clustering
  - There are unequal numbers of images for the various classes
  - As low as 303, as high as 942 images
  - Should cluster with fixed number of images

Clustering

- Tested clustering with random samples (10,000 iterations) of fixed number of images



Discussion

- With no previous knowledge of animal attributes, we perform comparably to Lampert
  - In addition, our maximum accuracies outperform Lampert's
- However, in defense of Lampert's method
  - Works in classification with no training examples
  - Wasn't meant for one-shot recognition, although his ideas should theoretically extend to performing well on one-shot as well
- 1v1 SVM's still need more work

Next Steps

- Tweak parameters for better results
  - Dimension we are using PCA down to
  - Dimension we are projecting down to
- Perform in-depth tests and benchmarks
  - Which training images give us better performance
  - Gather details for better understanding and analysis
- 1v1 SVM's
  - Use outputs as inputs to logistic projection
  - Optimize Chi-squared kernels (3 weeks to run tests)

Questions?

Any Questions?