One-shot Recognition

7/24

Kevin Tang

This Week

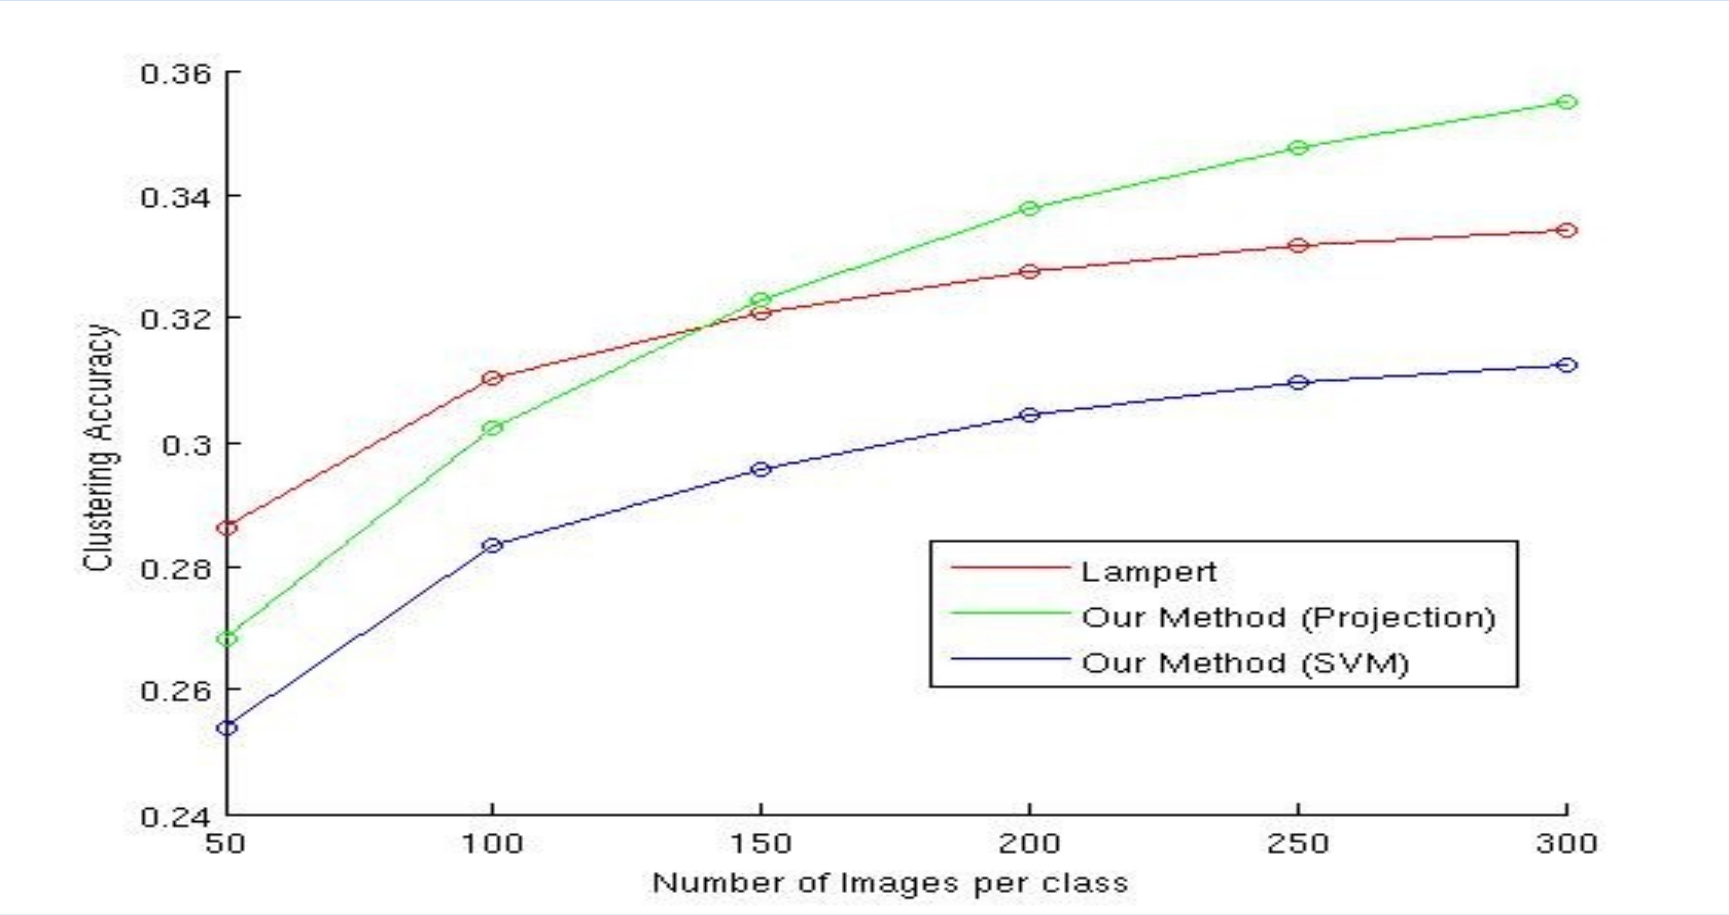
- Ran some benchmark tests on raw features (no transformation)
- Refined SVMs for better results
- Evaluate classification of 10 testing classes when all 50 classes are possible classifications
- Find sets of training/testing examples that perform better
- Comparison of our attribute space with Lampert's attribute space

SVM Results

- Improved SVM performance
  - Changed from linear to RBF kernels
  - Performed cross-validation to find best parameters
- One-shot classification
  - Chance Accuracy: 10%
  - Raw features (PCA to 500 dim): ~14.07%
  - Lampert's Accuracy: ~29.09%
  - Our Accuracy (Projection): ~27.41%
  - Our Accuracy (1v1 SVM's): ~27.08%

SVM Results

- Clustering classification
  - Chance Accuracy: 10%



New Experiment

- Previously, classified 10 testing classes
  - Chance accuracy of 10%
- New test of classifying same images from 10 testing classes, but possibility of being any of the 50 training and testing classes
  - Chance accuracy of 2%
- With Lampert's attribute formulation, should make no difference in accuracies
  - All animals have different attributes

Results

- One-shot classification
  - Chance Accuracy: 2%
  - Lampert's Accuracy: ~2%
  - Our Accuracy (Projection): ~7.48%
  - Our Accuracy (1v1 SVM's): ~9.11%
- We significantly outperform Lampert's performance on this test

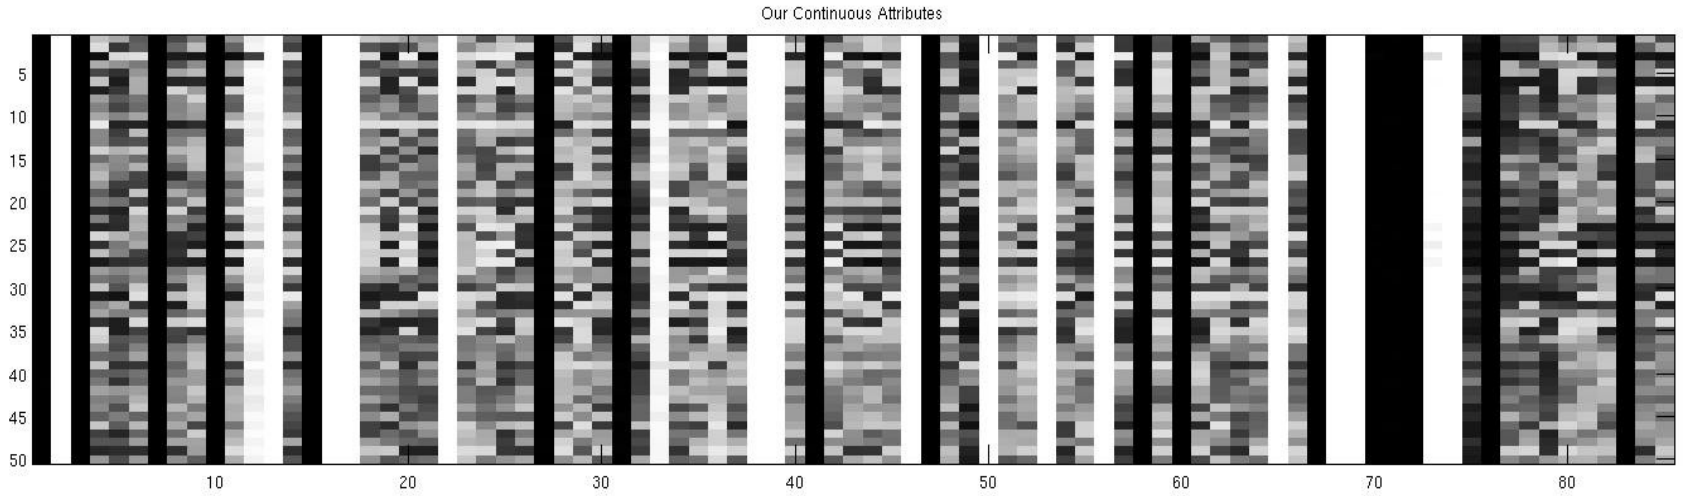
More Experiments

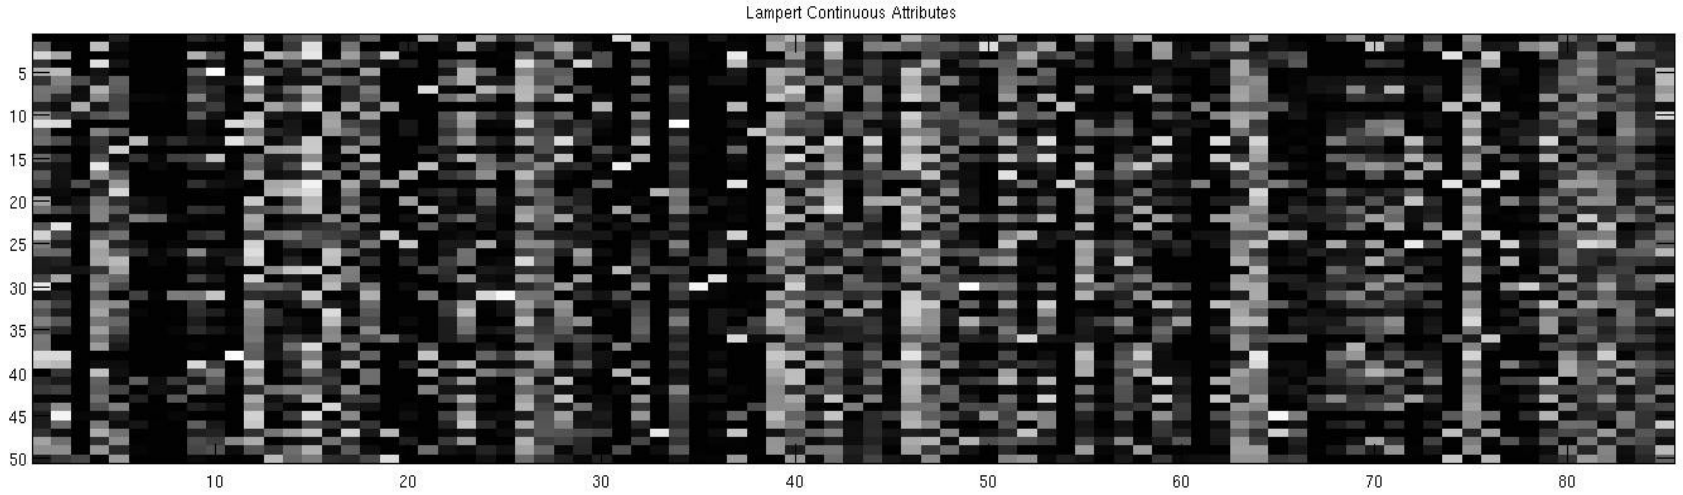
- Do the training/testing examples we randomly choose make a difference?
- Saved the random sets that produced good results
  - After enough iterations, all testing images are part of a good set
  - Training image selection is important
  - Need more tests and analysis

Attributes

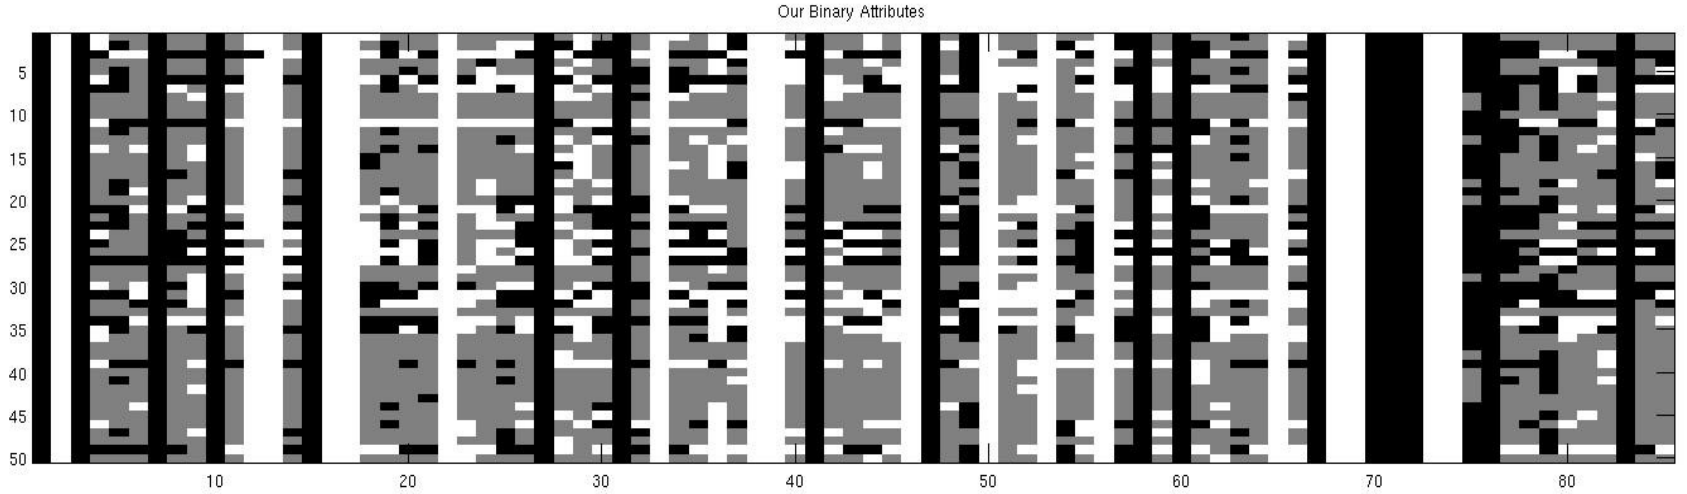
- Originally, Lampert connected animal classes by an 85 dimensional layer of attributes for the 50 animals
  - These attributes are human-defined
- We project the animals into an 85 dimensional space
  - Average the projected features for each animal
  - Compare our ”attributes” with Lampert's attributes

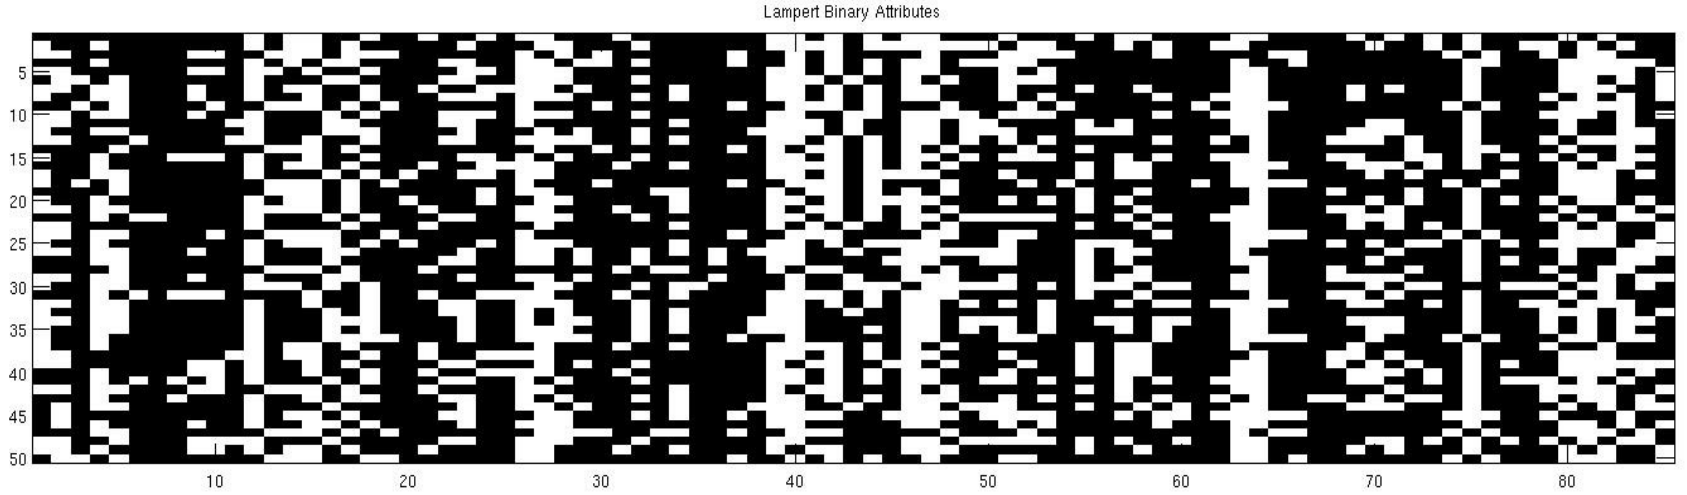
Continuous Attribute Space





Binary Attribute Space





Next Steps

- Better understanding of which training examples seem to help
- Deeper analysis of our attribute space
- Figure out story behind the paper and write it
- Lampert has joined us as a co-author
  - Test more of the data and results he has
  - See what ideas he can bring to the table

Questions?

Any Questions?