6/5
Project Introduction

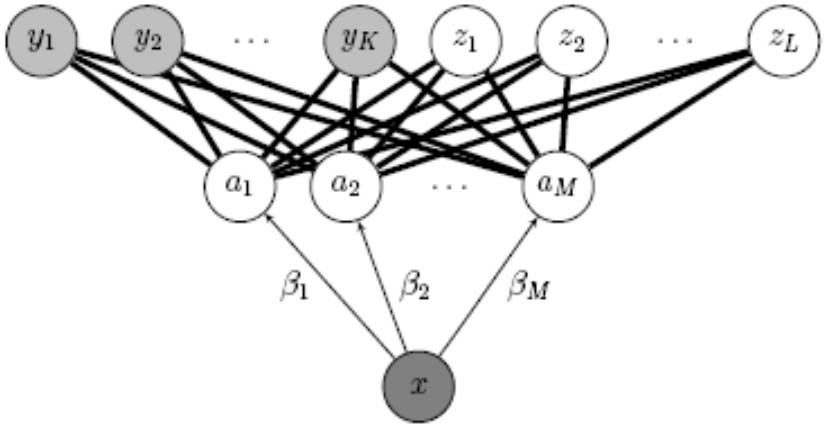
Kevin Tang

Project Background/Motivation

- Motivated by the paper "Learning To Detect Unseen Object Classes by Between-Class Attribute Transfer" by Lampert et al. at CVPR 2009
- Problem: Given a set of object classes to train our system from, can we recognize objects from classes we've never seen before?
- For example, given images from the classes "otter" and "polar bear" to train from. Can we recognize an image of a "zebra"?

Project Background/Motivation

- Solution: Cascade the known and unknown classes with a layer of semantic attributes to relate them



(b) Direct attribute prediction (DAP)

Project Background/Motivation

- Attributes are known and assigned beforehand.
- Example Attributes:

otter	
black:	yes
white:	no
brown:	yes
stripes:	no
water:	yes
eats fish:	yes
polar bear	
black:	no
white:	yes
brown:	no
stripes:	no
water:	yes
eats fish:	yes
zebra	
black:	yes
white:	yes
brown:	no
stripes:	yes
water:	no
eats fish:	no



Project Background/Motivation

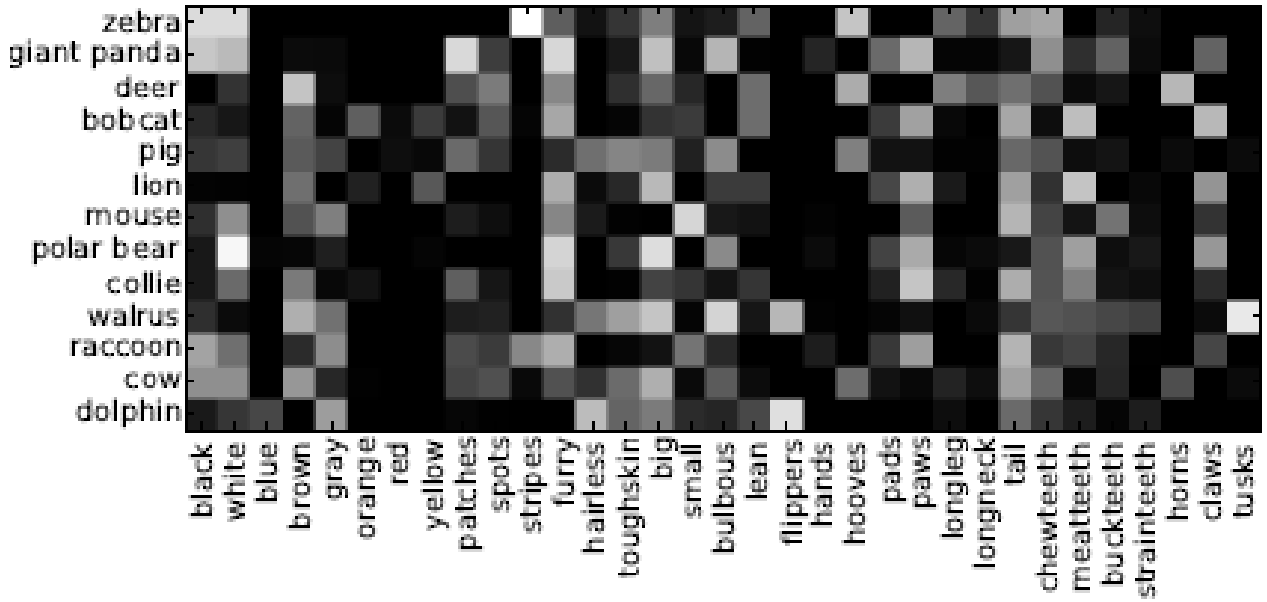
- Total of 50 animal object classes, 85 attributes
- Split the animals into two sets, 40 training classes, and 10 test classes
- Using this framework, trained classifiers to determine if an image has a given attribute
- Mapping these attributes to animals, able to determine which animal these attributes correspond to
- Accuracy of approximately 40.5% for multi-class recognition

Project Outline

- Extend upon Lampert's idea of semantic attributes
- Lampert made available his "Animals with Attributes" dataset and attributes
- Consists of 30475 images with at minimum 92 images per class
- Class-attribute matrix designating the strength of each attribute for each class

Project Outline

- Small subset of class-attribute matrix

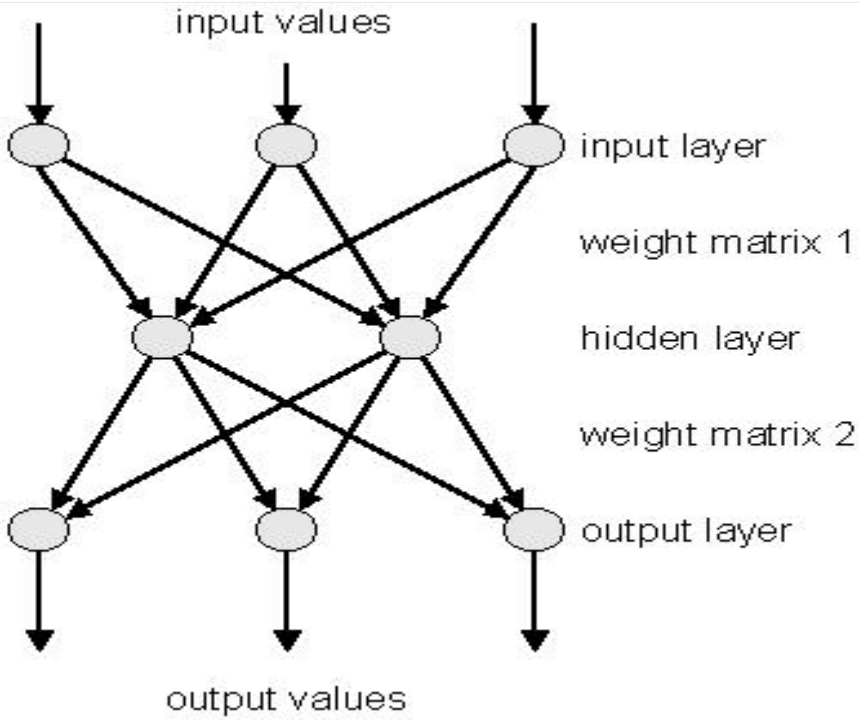


Project Outline

- From a representation point of view, idea is to model the idea of attribute classifiers with a Multi-Layer Perceptron (MLP)
- The hidden nodes represent the attribute classifiers
- Each hidden node represents the classifier for an arbitrary attribute that divides the object classes

Project Outline

Multi-Layer Perceptron:



Project Outline

- Use the idea of attributes to assist in multiple “one-shot recognition” cases
- One-shot recognition refers to being able to recognize a new object class with just one training image
- This isn’t easy, since we don’t have a set of training examples to train our system on
- Some previous work by Fei-Fei Li (Stanford) and Andrew Ng (Stanford)

Project Outline

- Solution: transfer learning, which is transferring learnt parameters from an unrelated classification problem to our problem at hand
- Our idea is to train a system to know it only has few examples to learn from, and use this framework to optimize and choose attributes that can best classify objects
- Hope that with a system trained in this manner, it can perform competitive one-shot recognition

Progress so far

- This week, read a large amount of papers and literature related to learning semantic attributes, one-shot recognition, and transfer learning
- Implemented a very preliminary version of the method used in Lampert's paper using MLPs and Netlab
- Not working very well, classification of ~91% for the training set, and ~15% for test images taken from the training set
- Learned a lot about MLPs

Next steps

- Figure out the math behind optimization idea Marshall has for training the system
- Implement idea and test it on the “Animals with Attributes” dataset provided by Lampert
- Compare the classification accuracy with Lampert’s accuracy, as well as accuracies obtained by other one-shot recognition systems (Fei-Fei Li, Andrew Ng)