

Attributes for Classifier Feedback



Amar Parkash and Devi Parikh



You can teach a child by examples...















And on, and on, and on...





Is this a giraffe?



No.



Is this a giraffe?



Yes.



Is this a giraffe?



No.



And on, and on, and on...



Proposed Active Learning Scenario

Current belief

Focused feedback Knowledge of the world

I think this is a giraffe. What do you think?



No, its neck is too short for it to be a giraffe.



Ah! These must not be giraffes either then.

[Animals with even shorter necks]



.....

Feedback on one, transferred to many

Current belief

I think this is a giraffe. What do you think?



Focused feedback Knowledge of the world

No, its neck is too short for it to be a giraffe.



- Learner learns better from its mistakes
- Accelerated discriminative learning with few examples


not be giraffes either then.



.....

Feedback on one, transferred to many

Communication



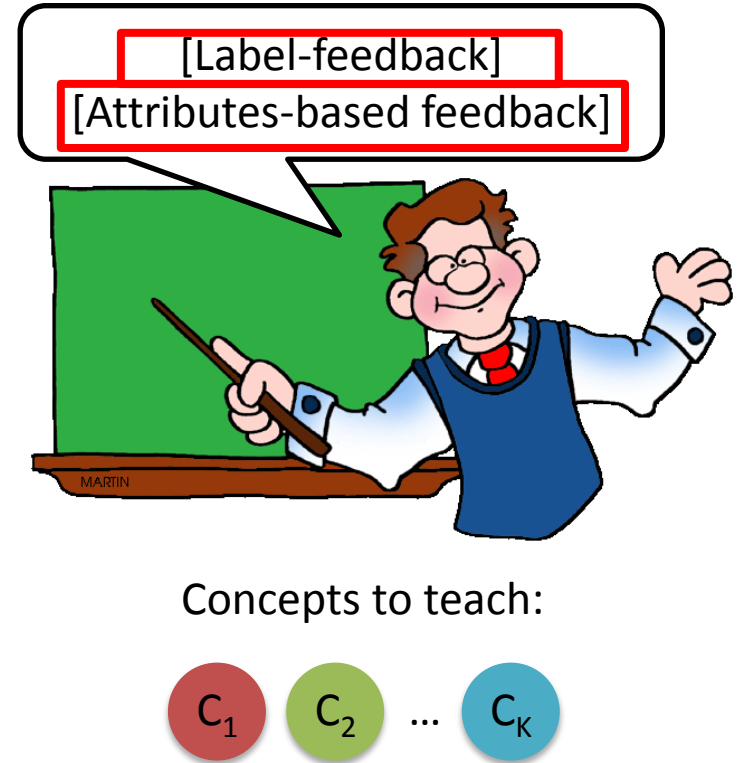
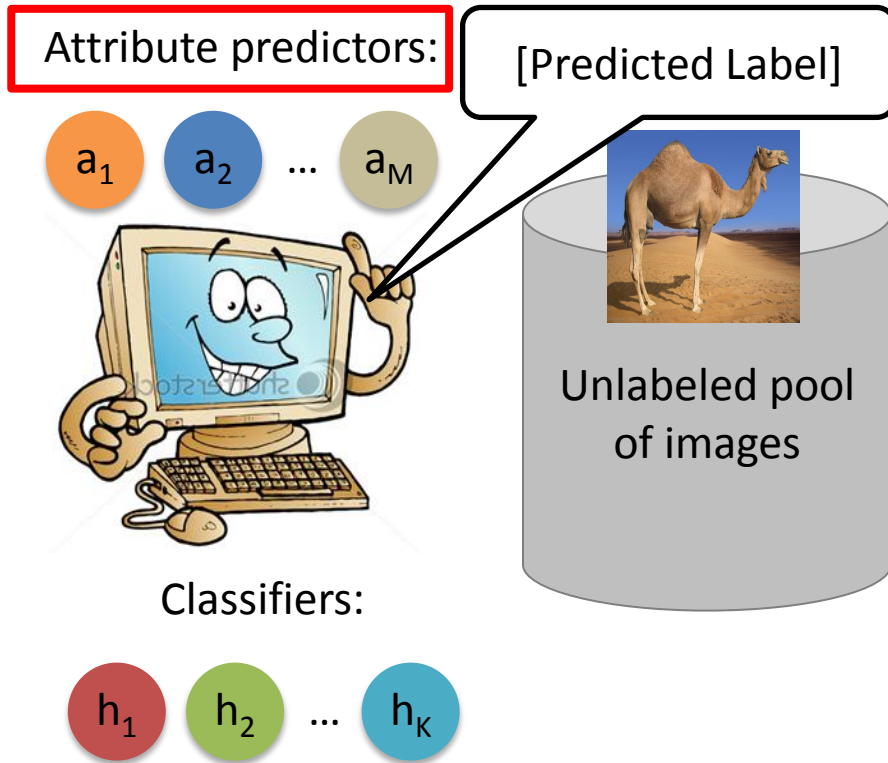
Need a language that is

- Machine understandable
- Human understandable

Attributes!

- Mid-level shareable
- Visual
- Semantic

Proposed Active Learning





Any feature space
Any discriminative learning algorithm

Relative Attributes

[Parikh and Grauman, ICCV 2011]

Attribute predictors:



$$r_m(\mathbf{x}_i) = \mathbf{w}_m^T \mathbf{x}_i$$

Parameters \swarrow \searrow Image features


$$O_m: \left\{ \left(\begin{array}{c} \text{Image 1} \\ \text{Image 2} \end{array} \right), \dots \right\}$$



Openness \rightarrow

Attributes-based Feedback

Attribute predictors:





Forest



Unlabeled pool of images

No,
It is too open to be a forest




Not Forest

Openness

Attributes-based Feedback

Attribute predictors:



Forest





Unlabeled pool of images

No,
It is too open to be a forest




Classifiers:



Proposed Active Learning

Attribute predictors:




[Predicted Label]



Unlabeled pool of images



Classifiers:



[Label-feedback]

[Attributes-based feedback]



Concepts to teach:




Label-based Feedback

- Not our contribution
- Experiment with different scenarios
- Benefits of attributes-based feedback
 - Small when label-based is very informative
 - Large when label-based feedback is weak

Label-based Feedback


- Accept: Yes, this is a forest.
 - Strong: It is not anything else
Example: Classification
 - Weak: It can be other things
Example: Annotation



Label-based Feedback



- Reject:
 - Strong: No, it is a coast.
Example: Classification with few classes
 - Weak: No, this is not a forest.
Example: Large-scale classification
Example: Biased binary classification
- 4 different scenarios in experiments



Datasets

- Datasets and relative attribute predictors from [Parikh and Grauman, ICCV 2011]

Datasets


- Faces:
 - 8 celebrity categories
 - 11 attributes (chubby, white, etc.)



[Kumar et al., ICCV 2009]

Datasets

- Scenes:
 - 8 categories
 - 6 attributes (open, natural, etc.)



[Oliva and Torralba, IJCV 2001]

Settings

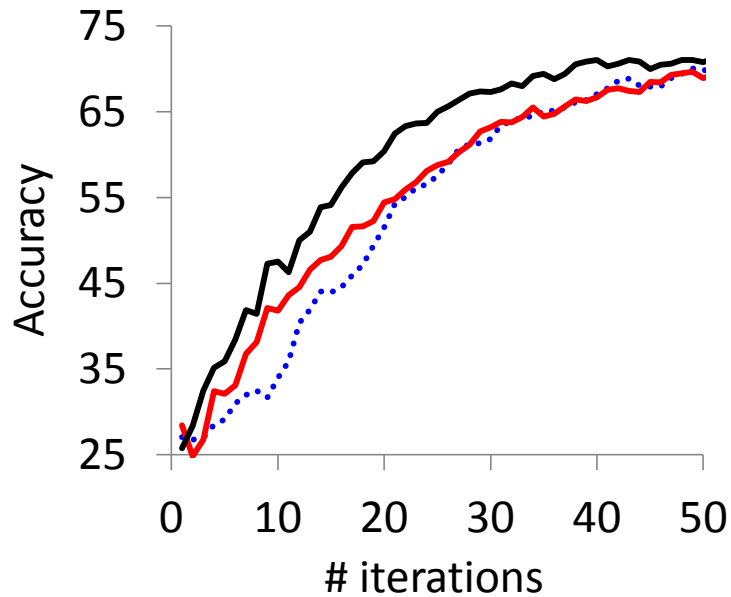
- Feedback from MTurk
- Features:
 - Raw image features (gist, color)
 - Attribute scores
- Category classifiers: SVM with RBF kernel
- Results on
 - 2 datasets x 2 features x 4 label-feedback scenarios
 - Show 2 here, rest in paper.



This image doesn't have enough perspective to be a street scene.

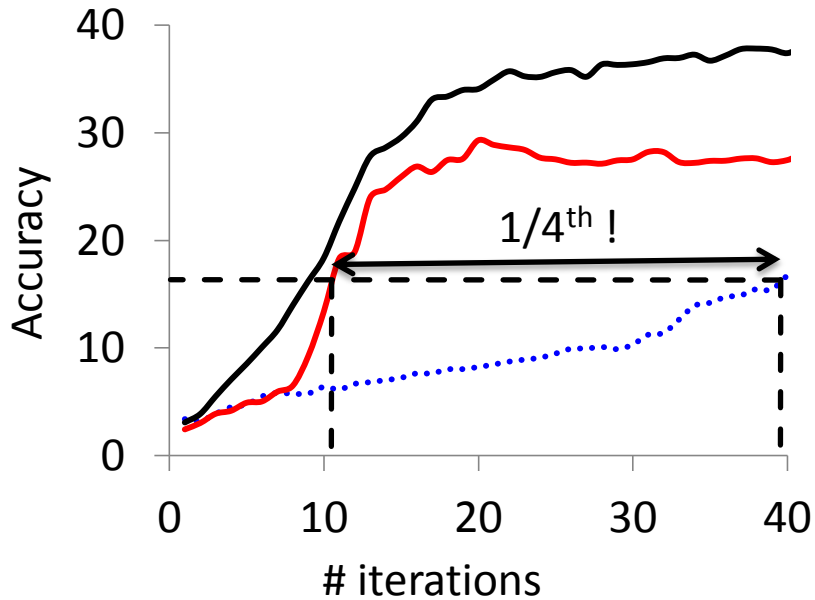
Results

Faces, Attribute features, Strong label-feedback



Results

Scenes, Image features, Weak label-feedback



More results in the paper.

Conclusion

- Attributes for providing classifier feedback
- Novel learning paradigm with enhanced human-machine communication
- Discriminative learning + domain knowledge
- Learning with few examples

- Connections to semi-supervised learning
 - Shrivastava, Singh and Gupta: Up Next!



Thank you!