

# Part III: Active Learning Beyond Label Feedback

Chicheng Zhang

University of California San Diego

June 21, 2017

## More Complex Feedback

- ▶ So far: feedback used = labels or abstentions from a single annotator
- ▶ Human beings are capable of much more complex feedback
- ▶ How to utilize such feedback?

# Outline

1. Active Learning using More Complex Queries

2. Queries from Labelers with Varying Expertise

## Challenge of Active Learning: Rare Classes [Das05, AP11]

- ▶ Impossible to account for **rare classes** if never observed
- ▶  $\Omega(\frac{1}{\epsilon})$  label complexity



balanced classes



rare class

## Challenge of Active Learning: Small Disjuncts [AP11]

- ▶ Classes are spread in “small islands”
- ▶ Need to find all rare subclasses to get an accurate classifier

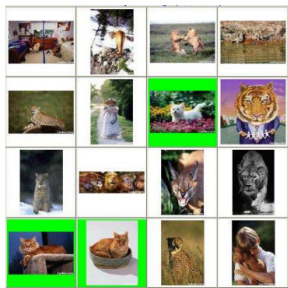


## Complex Queries: Addressing the Rare-class problem

- ▶ In active learning, rare classes often leads to problems
- ▶ Remedy: more complex queries
  - ▶ Class-conditional queries (CCQ) [CTGC05, BH12]
  - ▶ Search queries [AP10, BHLZ16]

# Class-Conditional Queries (CCQ) [BH12]

- ▶ Oracle answers questions e.g.: “Show a cat among these images” [CTGC05]



CCQ



“Cat”

## Key Observation: Seed examples + Label queries

$H =$  indicator functions for subintervals of  $[0, 1] \subset \mathbb{R}$ ,

$\mathcal{D}_X =$  uniform on  $[0, 1]$

- ▶ Positive class is rare
- ▶ In noiseless setting, **need a seed positive example.**






## Key Observation: Seed examples + Label queries

$H =$  indicator functions for subintervals of  $[0, 1] \subset \mathbb{R}$ ,

$\mathcal{D}_X =$  uniform on  $[0, 1]$


- ▶ Positive class is rare
- ▶ In noiseless setting, **need a seed positive example.**



## Key Observation: Seed examples + Label queries

$H$  = indicator functions for subintervals of  $[0, 1] \subset \mathbb{R}$ ,  
 $\mathcal{D}_X$  = uniform on  $[0, 1]$

- ▶ Positive class is rare
- ▶ In noiseless setting, **need a seed positive example**.




Armed with seed positive example, and negative examples, can use **binary search** to find interval boundaries via **label queries**.

## Key Observation: Seed examples + Label queries

$H$  = indicator functions for subintervals of  $[0, 1] \subset \mathbb{R}$ ,

$\mathcal{D}_X$  = uniform on  $[0, 1]$

- ▶ Positive class is rare
- ▶ In noiseless setting, **need a seed positive example**.



Armed with seed positive example, and negative examples, can use **binary search to find interval boundaries** via **label queries**.

Search queries [AP10, BHLZ16]: Similar idea but weaker than class conditional queries

# Outline

1. Active Learning using More Complex Queries
2. Queries from Labelers with Varying Expertise

# Motivation

- ▶ What if we have auxiliary information? - as an extra label oracle




Oracle: expensive but correct




Weak labeler: cheap, sometimes wrong

## What Makes a Labeler Weak?

- ▶ Noise (See e.g. [KOS13])
- ▶ Bias (This Work)



Labels by Oracle  $\mathcal{O}$



Labels by Weak Labeler  $\mathcal{W}$

## Formal Model [ZC15]

Given:

- ▶ Access to unlabeled examples drawn from  $D_{\mathcal{X}}$
- ▶ **Abilities to query oracle  $\mathcal{O}$**
- ▶ **Abilities to query weak labeler  $\mathcal{W}$**

Goal:

- ▶ Get a classifier  $\hat{h}$  with excess error  $\epsilon$  **wrt**  $\mathcal{O}$  with probability  $1 - \delta$

Label Complexity:

- ▶ How many label queries ( $m(\epsilon, \delta)$ ) **to**  $\mathcal{O}$  are needed to achieve this goal?


## How do We Address Weakness?

- ▶ Learn where weak and strong labelers differ
- ▶ Run standard active learning
  - ▶ Query  $\mathcal{O}$  in difference region
  - ▶ Query  $\mathcal{W}$  outside difference region
- ▶ Problem: may be statistically inconsistent




# Statistical Inconsistency

- ▶ False Negatives (incorrectly predict  $\mathcal{O}$  and  $\mathcal{W}$  agree) lead to wrong annotations




Actual Labels




Annotation using  $h^*$  as difference classifier

# Statistical Inconsistency

- ▶ False Negatives (incorrectly predict  $\mathcal{O}$  and  $\mathcal{W}$  agree) lead to wrong annotations




Actual Labels




Annotation using  $h^*$  as  
difference classifier

## Solution

Learn where  $\mathcal{O}$  and  $\mathcal{W}$  differ subject to low false negative (FN) rates




Actual Labels




Annotation using  $h^*_{FN}$  as  
difference classifier

## Solution

Learn where  $\mathcal{O}$  and  $\mathcal{W}$  differ subject to low false negative (FN) rates



Actual Labels






Annotation using  $h^*_{FN}$  as  
difference classifier




# Label Complexity

- ▶ Training the difference classifier over entire space does not save labels
- ▶ Solution: train difference classifiers in disagreement regions only (at each phase)
- ▶ Label complexity for the rest of active learning can also be established



# References I

-  Josh Attenberg and Foster Provost.  
Why label when you can search?: alternatives to active learning for applying human resources to build classification models under extreme class imbalance.  
*In Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 423–432. ACM, 2010.
-  Josh Attenberg and Foster Provost.  
Inactive learning?: difficulties employing active learning in practice.  
*ACM SIGKDD Explorations Newsletter*, 12(2):36–41, 2011.
-  Maria Florina Balcan and Steve Hanneke.  
Robust interactive learning.  
*In Conference on Learning Theory*, pages 20–1, 2012.

## References II

-  Alina Beygelzimer, Daniel J Hsu, John Langford, and Chicheng Zhang.  
Search improves label for active learning.  
*In Advances in Neural Information Processing Systems*, pages 3342–3350, 2016.
-  Edward Y Chang, Simon Tong, Kingshy Goh, and Cheng-Wei Chang.  
Support vector machine concept-dependent active learning for image retrieval.  
2005.
-  S. Dasgupta.  
Coarse sample complexity bounds for active learning.  
*In NIPS*, 2005.

## References III

-  David R Karger, Sewoong Oh, and Devavrat Shah.  
Efficient crowdsourcing for multi-class labeling.  
*ACM SIGMETRICS Performance Evaluation Review*,  
41(1):81–92, 2013.
-  Chicheng Zhang and Kamalika Chaudhuri.  
Active learning from weak and strong labelers.  
In *Advances in Neural Information Processing Systems*, pages  
703–711, 2015.