Deep Active Learning

2018.9.28

Outline

Cost-Effective Active Learning for Deep Image Classification

Generative Adversarial Active Learning

Adversarial Sampling for Active Learning

Active Learning

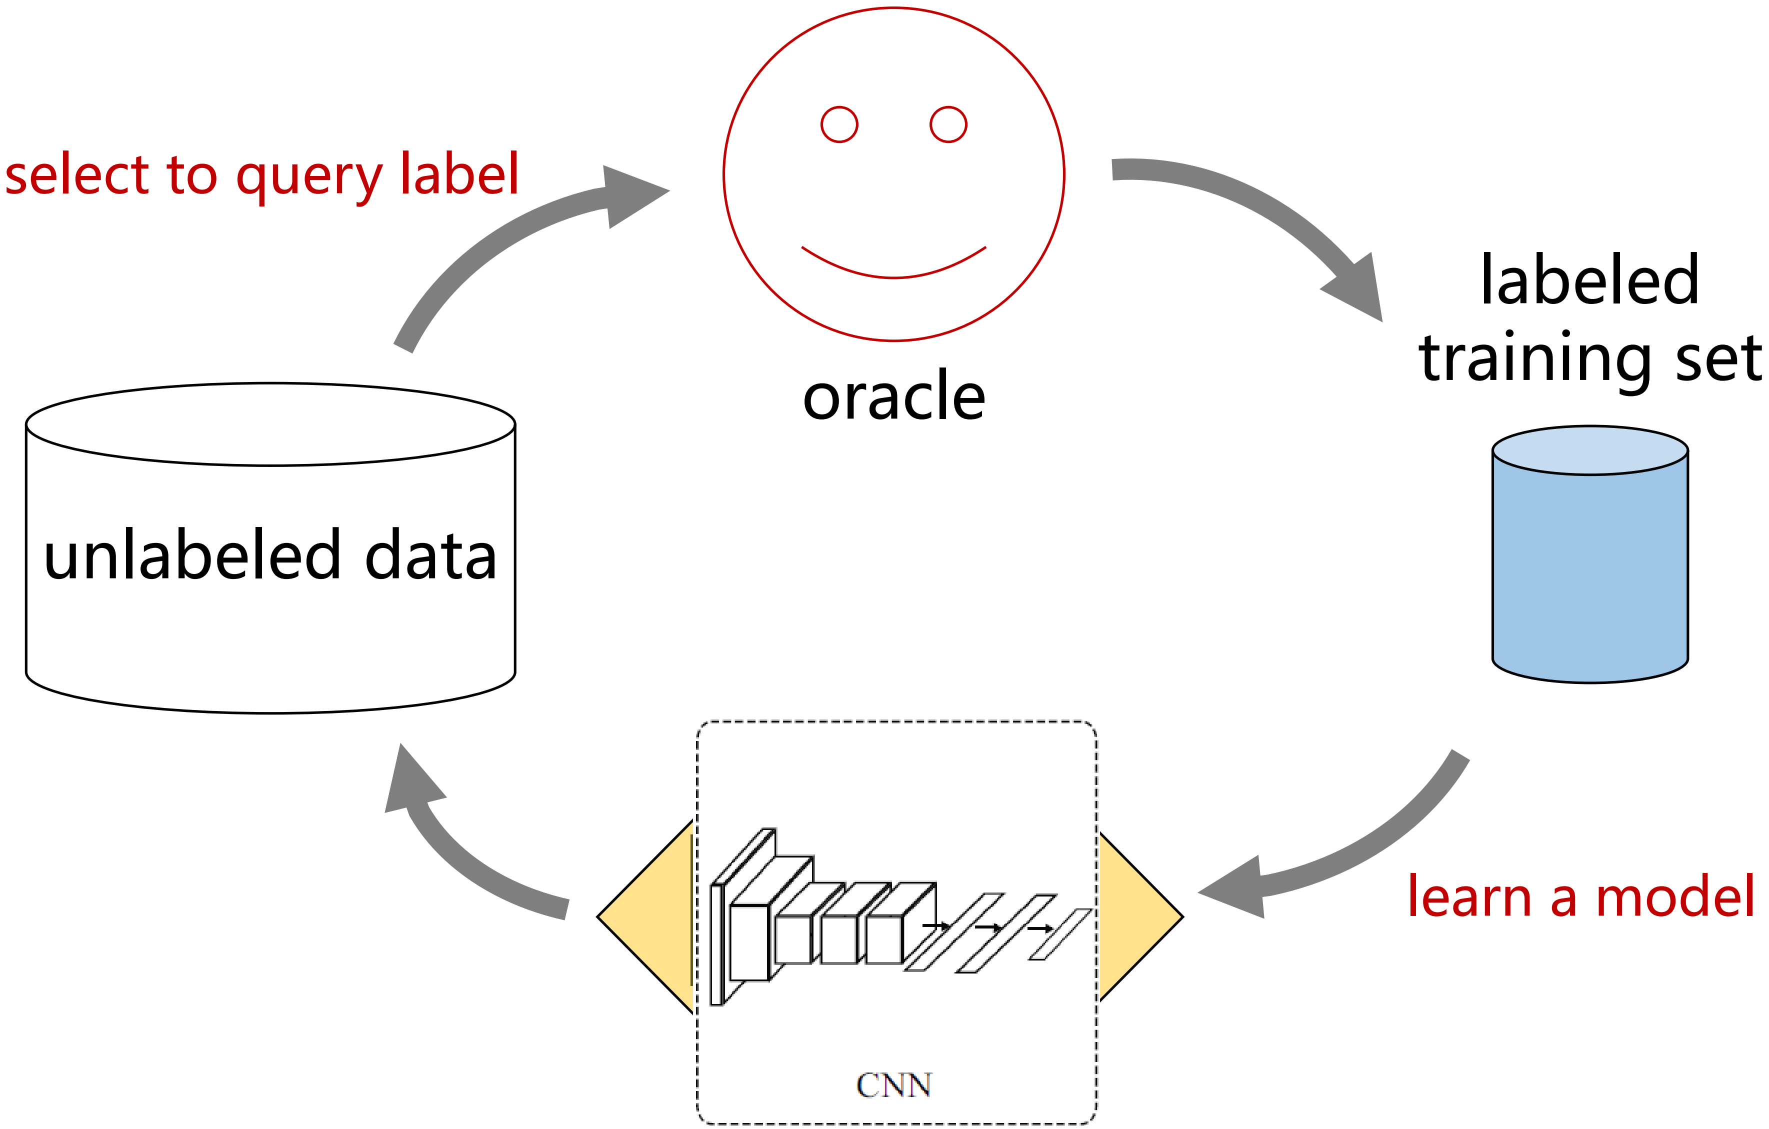

Motivation

- Active Learning vs Deep Learning

  - Minority queried samples are insufficient for training deep CNNs
    - AL usually select only a few of the most informative samples in each step
    - Difficult to obtain proper feature representations by fine-tuning CNNs with these minority informative samples.

  - The process pipelines of AL and CNNs are inconsistent with each other
    - The feature learning and classifier training are jointly optimized in CNNs.
    - AL strategies to select samples assume that the feature representation is fixed

Active Learning for Deep Image Classification

- Cost-Effective Active Learning (CEAL) framework

The majority high confidence samples conduce to learn more discriminative feature representation .

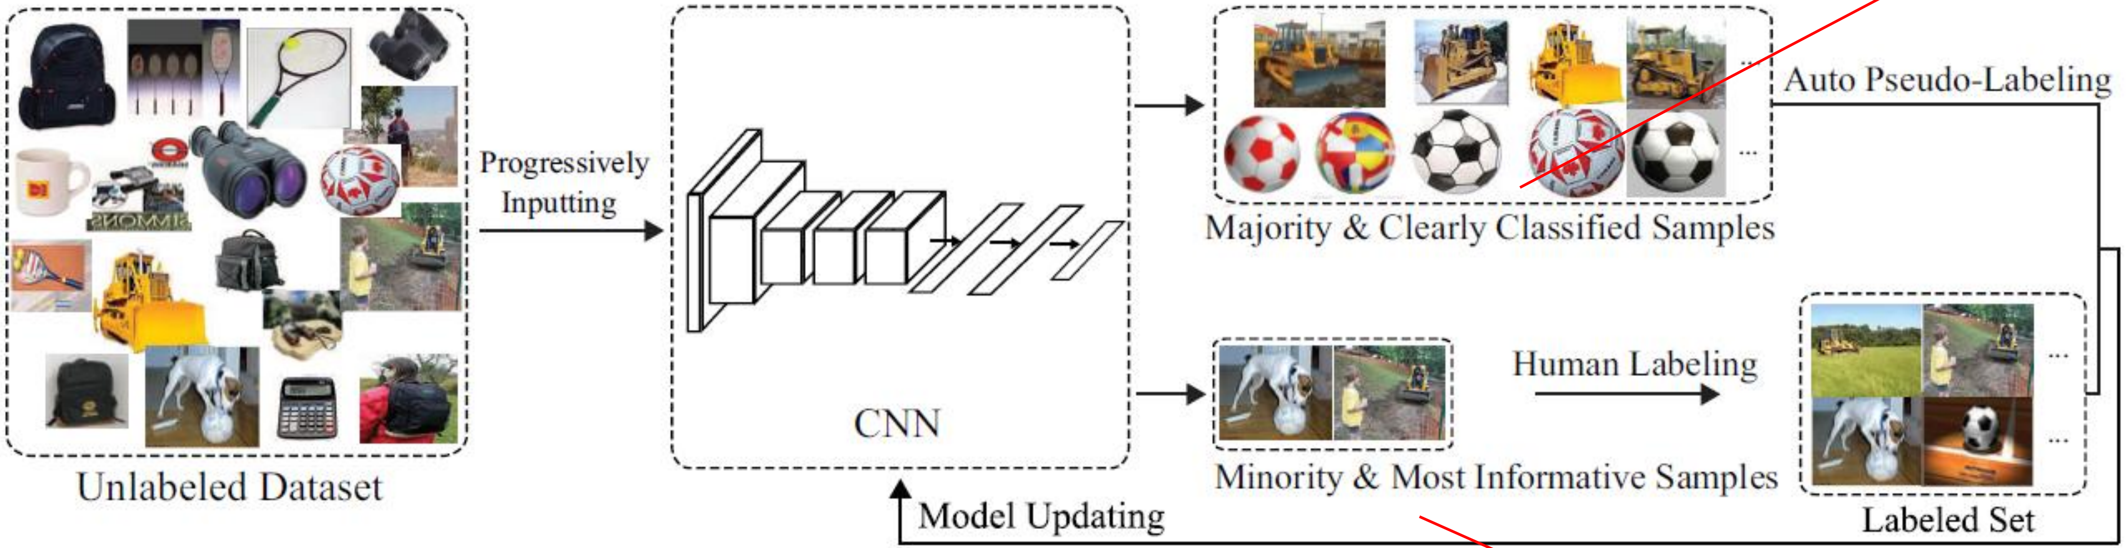

The minority informative samples contribute to train more powerful classifier

Cost-Effective Active Learning

- Objective function

$$\min_{\{\mathcal{W}, y_i, i \in D^U\}} -\frac{1}{n} \sum_{i=1}^{n} \sum_{j=1}^{m} 1\{y_i = j\} \log p(y_i = j | x_i; \mathcal{W})$$

Alternatively update the pseudo-labeled sample $y_i \in D^U$ and the network parameters $\mathcal{W}$

Method

- Initialization

  Randomly select few training samples from D^U and manually annotate them as the start point to initialize the CNN parameters

- Complementary sample selection

  - Informative Sample Annotating
  - High Confidence Sample Pseudo-labeling

Method

Informative Sample Annotating

D^L

Select K most informative/
uncertain samples left in D^U

Least confidence :

$$lc_i = \max_j p(y_i = j | x_i; \mathcal{W})$$

Margin sampling :

$$ms_i = p(y_i = j_1 | x_i; \mathcal{W}) - p(y_i = j_2 | x_i; \mathcal{W})$$

Entropy :

$$en_i = -\sum_{j=1}^{m} p(y_i = j | x_i; \mathcal{W}) \log p(y_i = j | x_i; \mathcal{W})$$

Method

High Confidence Sample Pseudo-Labeling

Majority & Clearly Classified Samples

Auto Pseudo-Labeling → High confidence samples D^H

Select high confidence samples from D^U, whose entropy is smaller than the threshold δ

$$j^* = \arg\max_j p(y_i = j | x_i; \mathcal{W}),$$

$$y_i = \begin{cases} j^*, & en_i < \delta, \\ 0, & \text{otherwise.} \end{cases}$$

Threshold Updating

$$\delta = \begin{cases} \delta_0, & t = 0 \\ \delta - dr * t, & t > 0 \end{cases}$$

Method

CNN fine-tune

Fixing the labels of samples D^H and D^L, the objective function can be simplified as

$$\min_{\mathcal{W}} -\frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{m}\mathbf{1}\{y_i = j\}\log p(y_i = j|x_i;\mathcal{W})$$

After fine-tuning, the high confidence samples D^H would be put back to D^U and erased pseudo-label.

Experiment

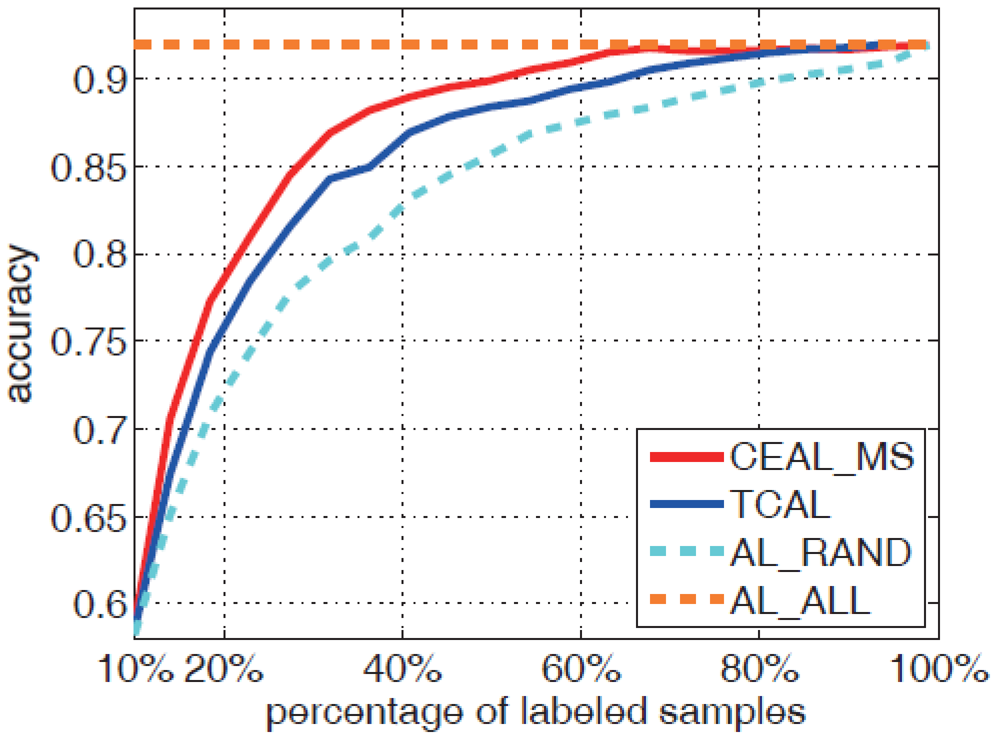

(a) CACD dataset

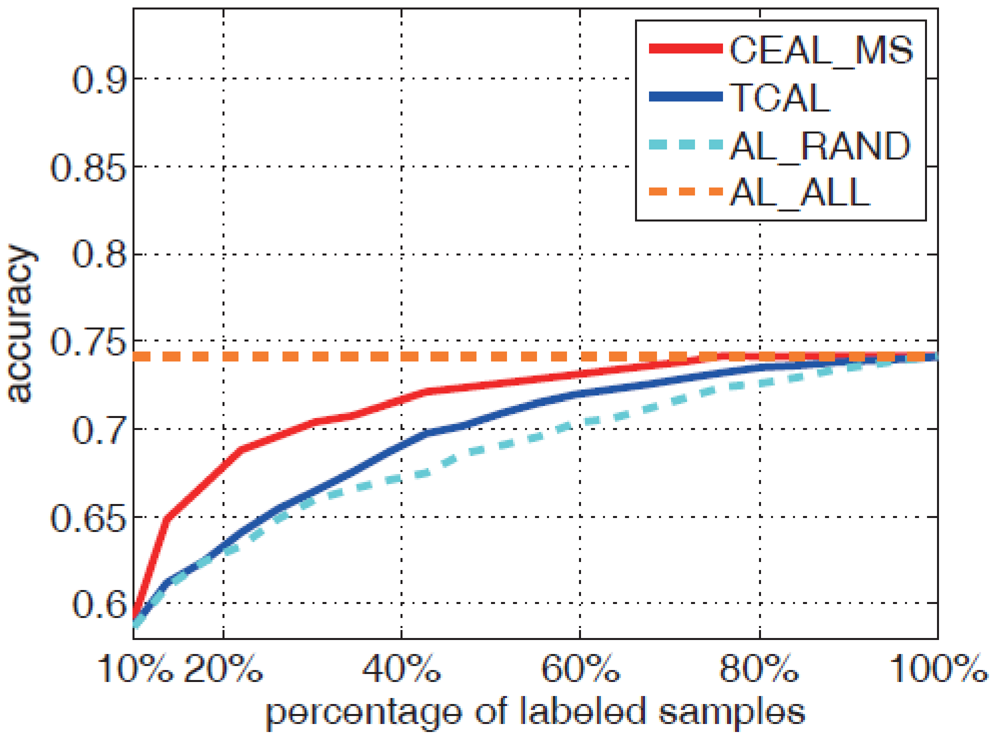

(b) Caltech-256 dataset

Experiment

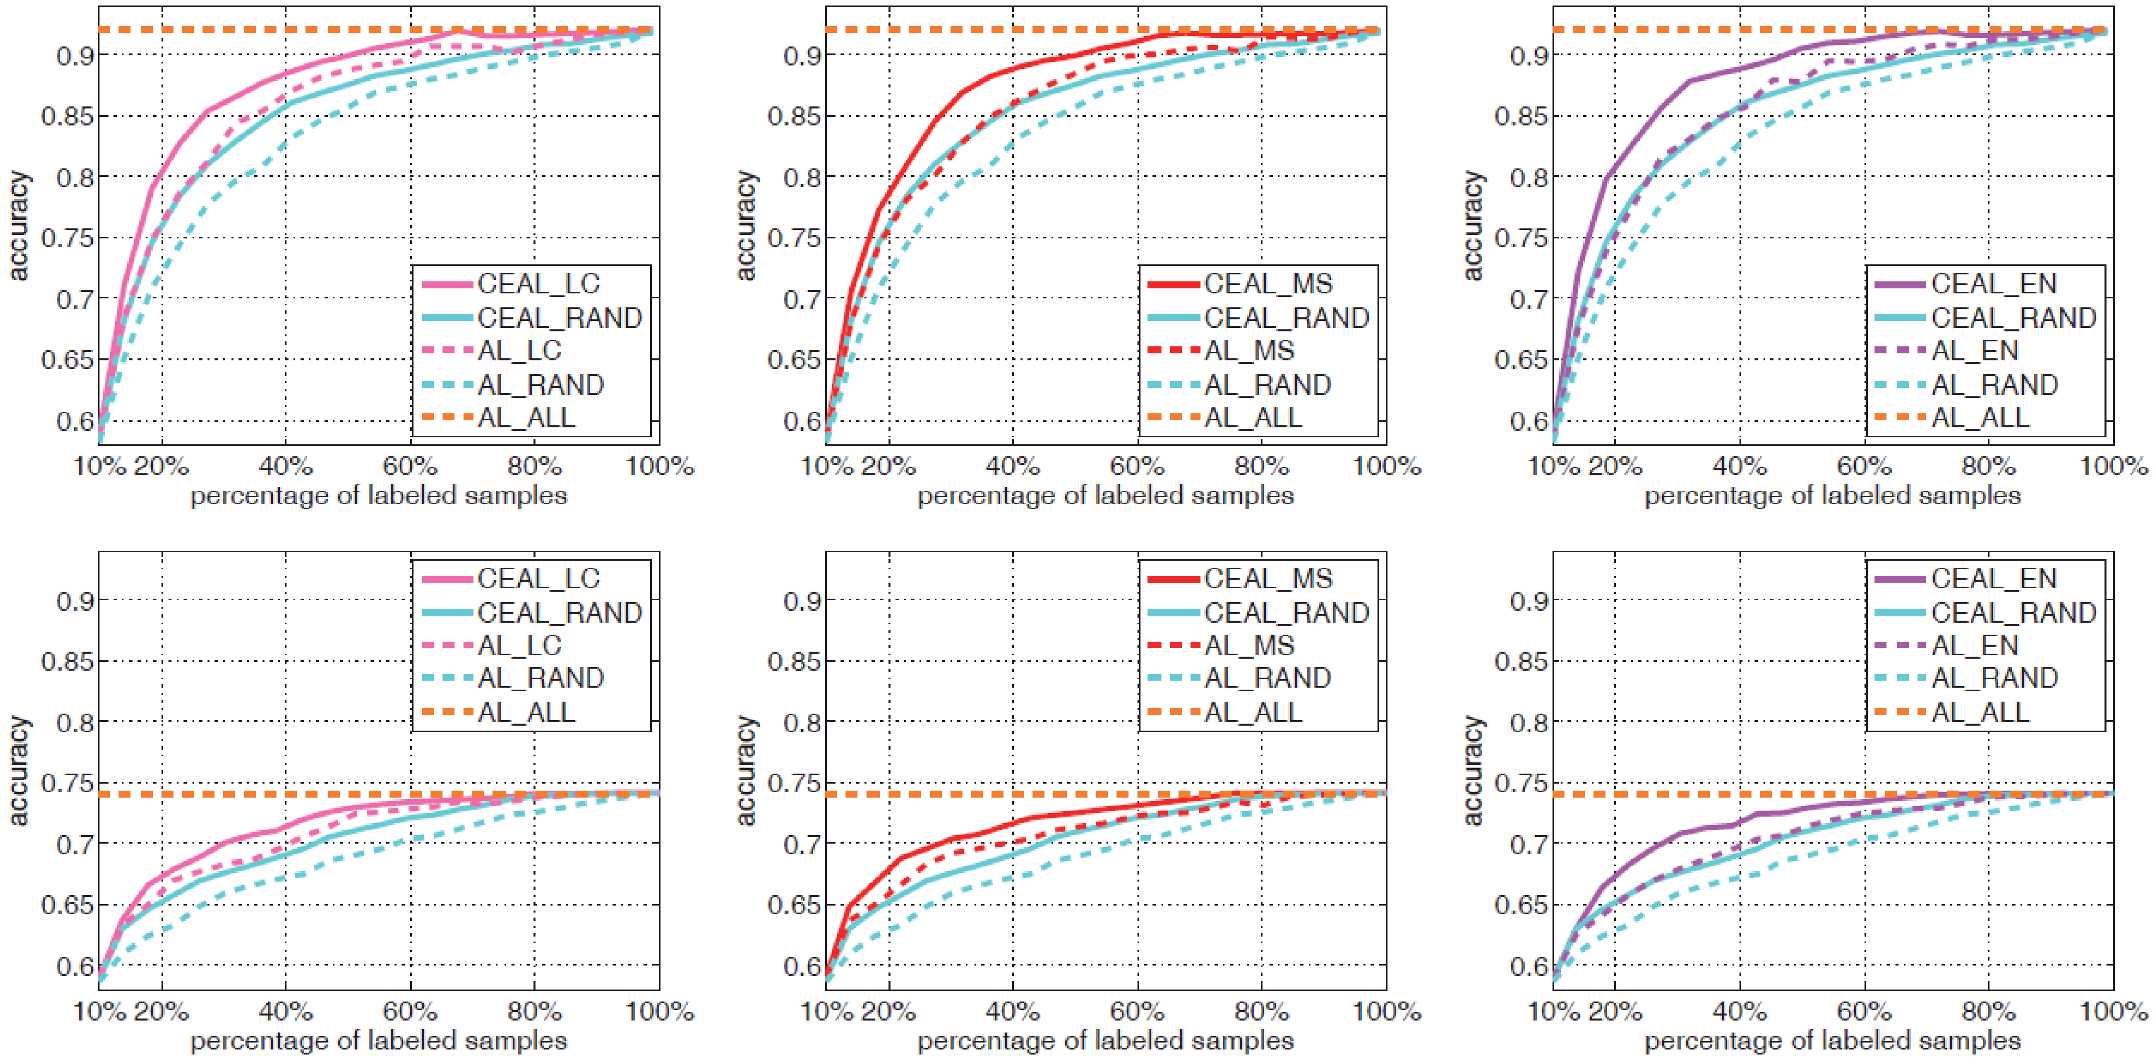

accuracy
0.9
0.85
0.8
0.75
0.7
0.65
0.6
10% 20% 40% 60% 80% 100%
percentage of labeled samples
CEAL_LC
CEAL_RAND
AL_LC
AL_RAND
AL_ALL
CEAL_MS
AL_MS
CEAL_EN
AL_EN

Experiment

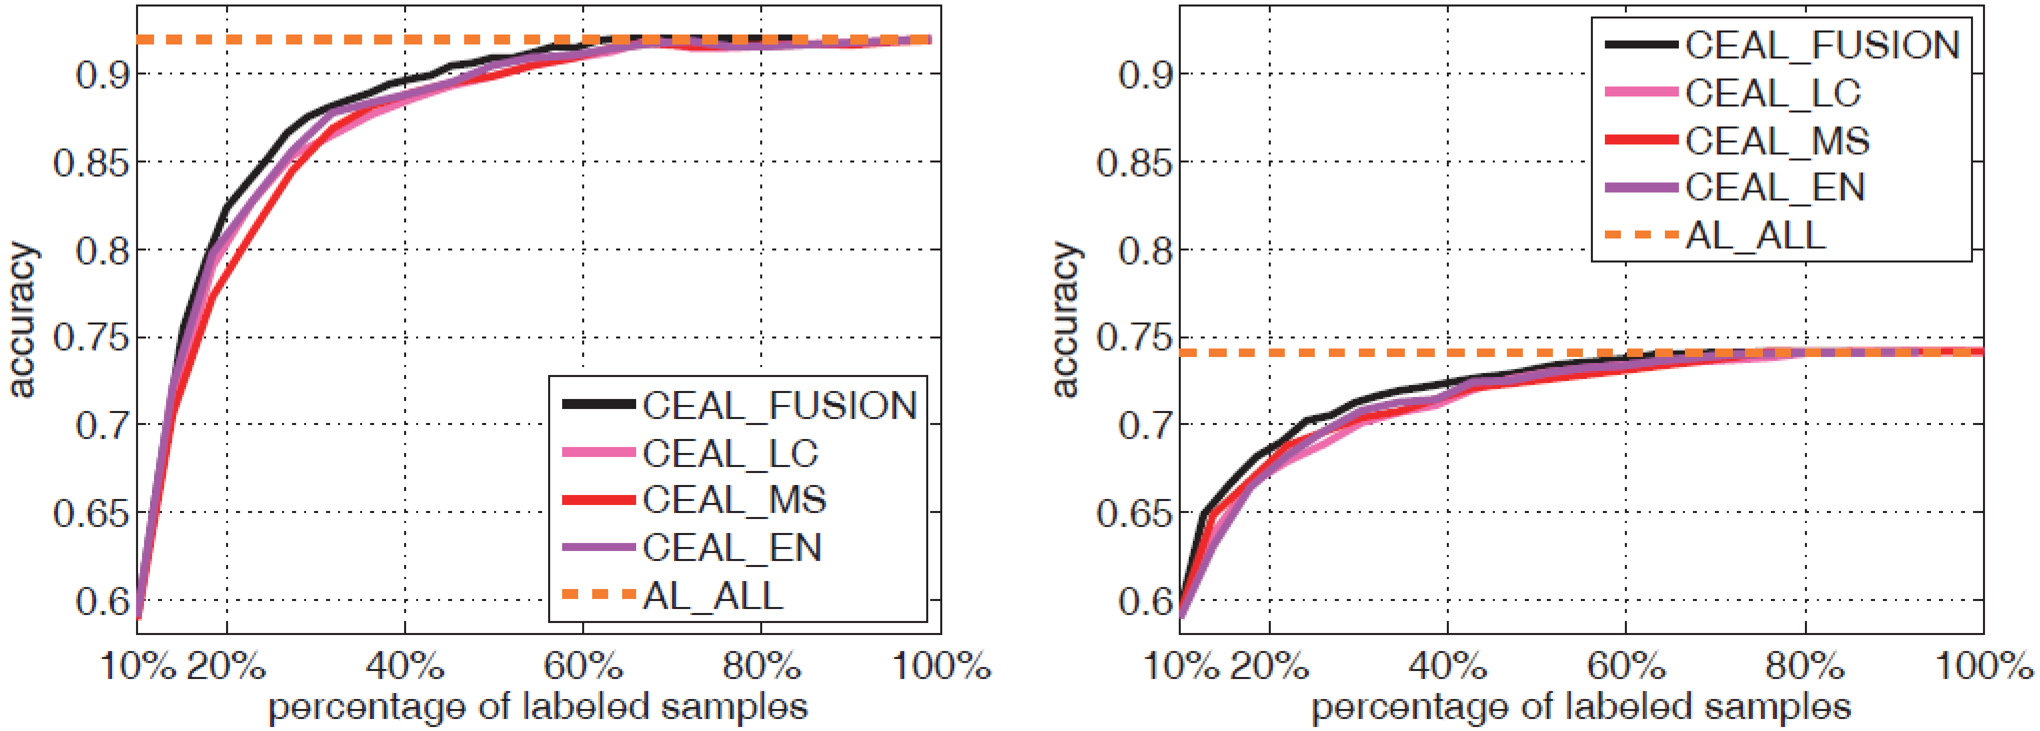

accuracy
0.9
0.85
0.8
0.75
0.7
0.65
0.6
10% 20% 40% 60% 80% 100%
percentage of labeled samples
CEAL_FUSION
CEAL_LC
CEAL_MS
CEAL_EN
AL_ALL

Introduction

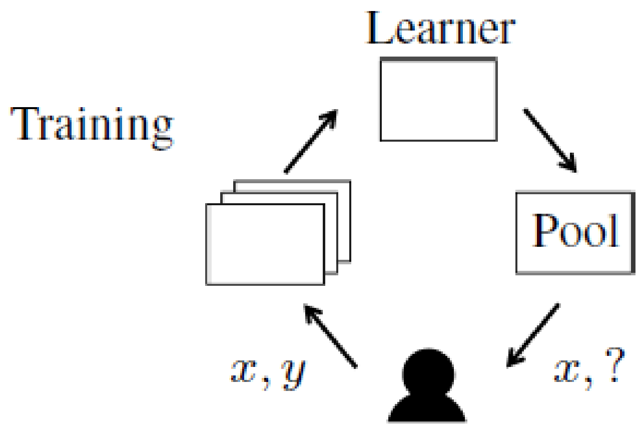

(a) Pool-based

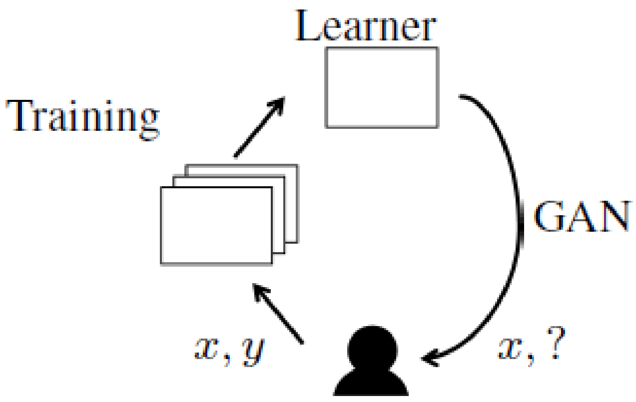

(b) GAAL

The learner selects samples for querying from a given unlabeled pool

The learner synthesis samples for querying using GAN

Generative Adversarial Net

A two-player minimax game between the generator G and the discriminator D

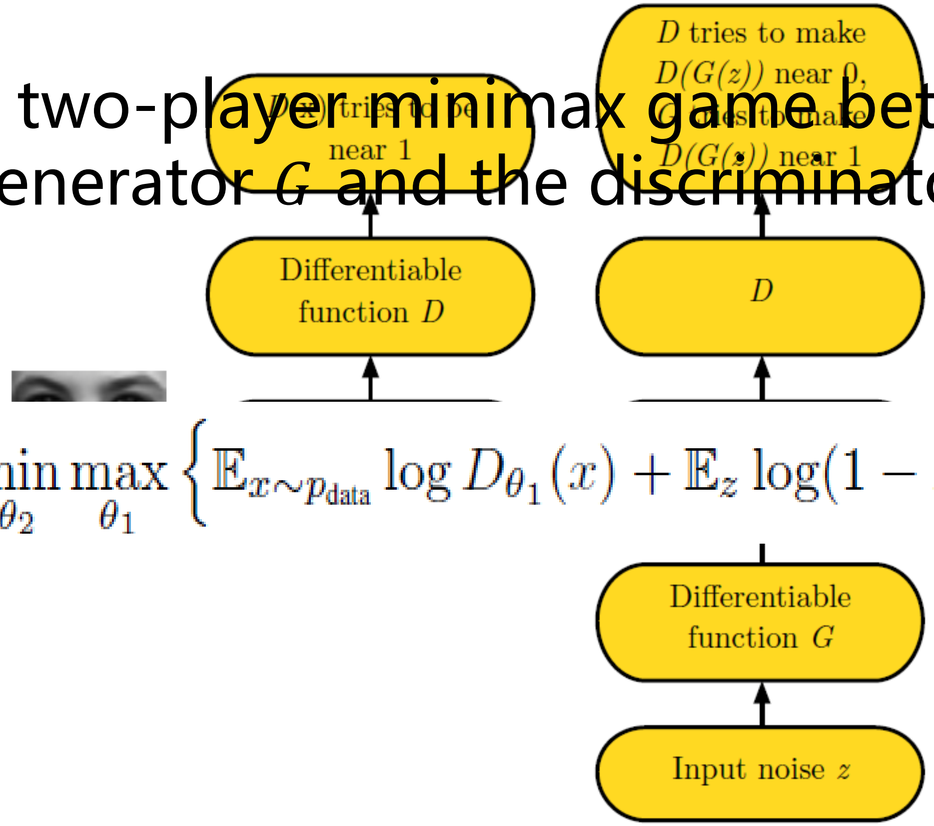

$$\min_{\theta_2}\max_{\theta_1}\left\{\mathbb{E}_{x\sim p_{\text{data}}}\log D_{\theta_1}(x)+\mathbb{E}_z\log(1-D_{\theta_1}(G_{\theta_2}(z)))\right\}$$

(NIPS 2016 Tutorial: Generative Adversarial Networks, Ian Goodfellow)

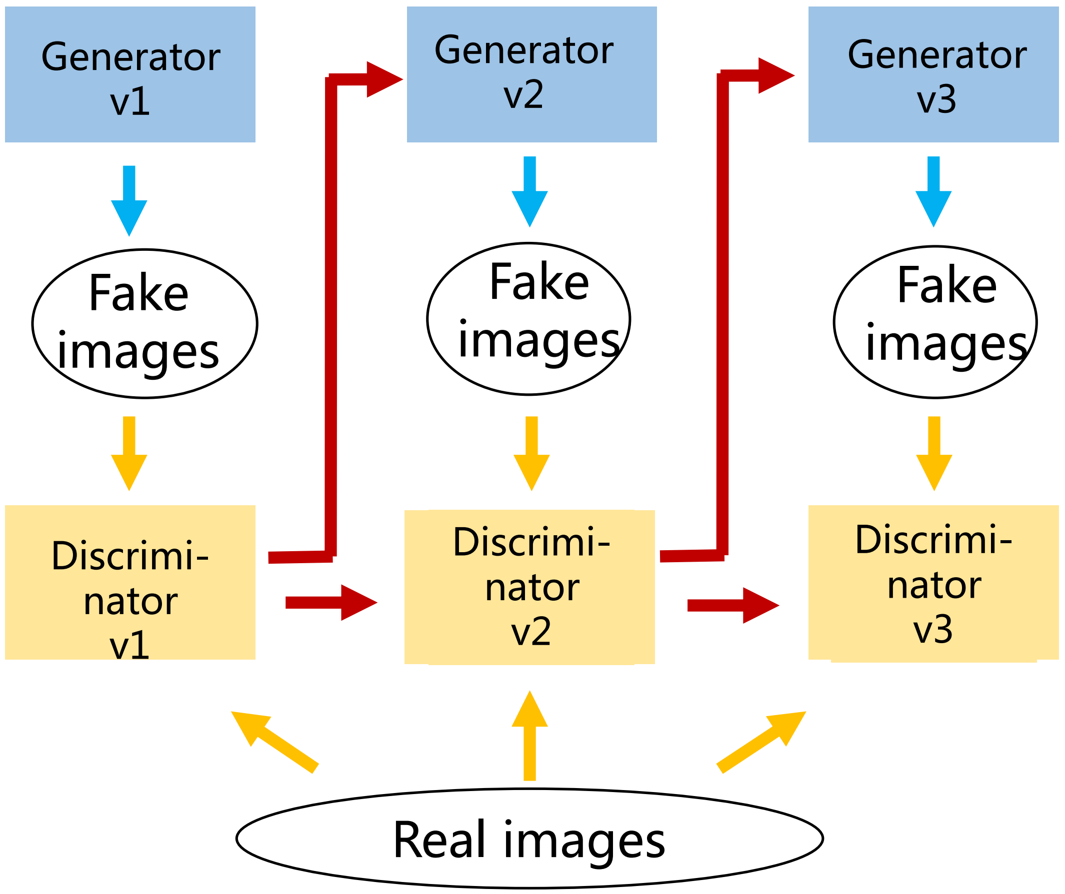

Generative Adversarial Active Learning

$$\|W\phi(x)+b\|$$

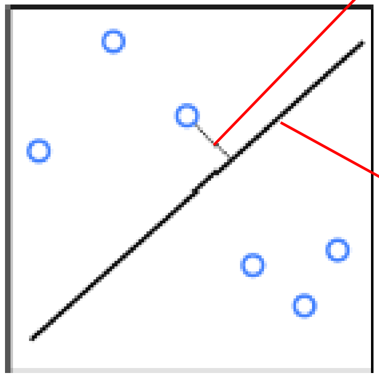

$$f(x)=W\phi(x)+b$$

(a) SVM_{active}

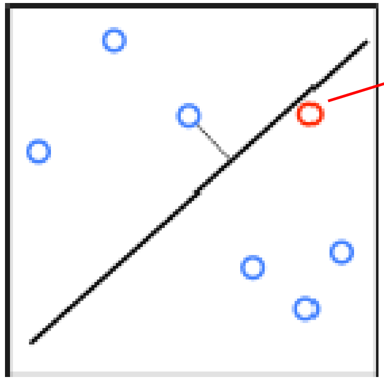

$G(z)$

(b) GAAL

$$\min_{z} \|W^{\top}\phi(G(z))+b\|$$

$\{z_1, z_2, \ldots\}$ $\xrightarrow{\text{GANs}}$

$\xrightarrow{\text{Labeling}}$

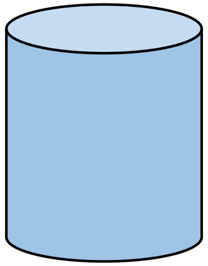

Experiment

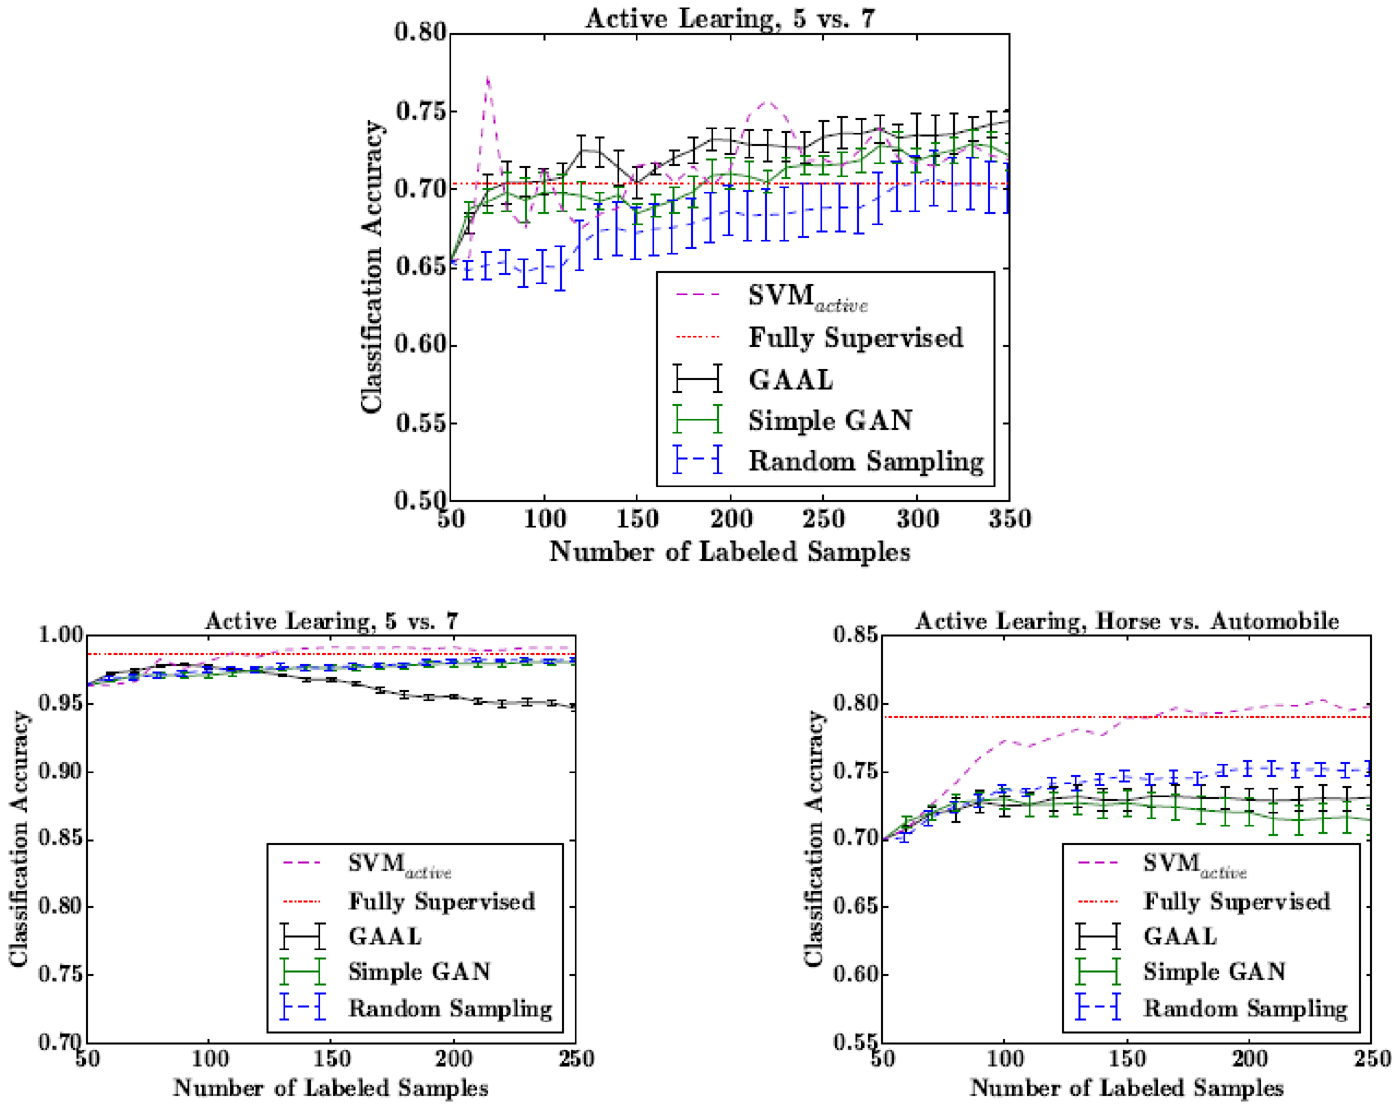
Active Learing, 5 vs. 7
Classification Accuracy
Number of Labeled Samples
SVM$_{active}$
Fully Supervised
GAAL
Simple GAN
Random Sampling
Active Learing, 5 vs. 7
Active Learing, Horse vs. Automobile

Adversarial Sampling for Active Learning

Pool-based Uncertainty Sampling

$$x_{\text{new}} := \underset{x \in \mathcal{P}}{\arg\min} \, \|W_{\theta(k)} \phi(x) + b\|_2$$

The information entropy directly measures the information of a given sample

$$H(x) := -\sum_{c \in \mathcal{C}} p(c|x) \log\left(p(c|x)\right)$$

Maximum entropy criteria

$$\begin{aligned} x_{\text{new}} &:= \underset{x \in \mathcal{P}}{\arg\max} \, H(x, \theta(k)) \\ &= \underset{x \in \mathcal{P}}{\arg\min} \sum_{c \in \mathcal{C}} p_{\theta(k)}(c|x) \log\left(p_{\theta(k)}(c|x)\right) \end{aligned}$$

Adversarial Sampling for Active Learning

Adversarial Sample Generation using GANs

Minimum distance: $$z^* := \arg\min_z \|W_{\theta(k)}\phi(G(z)) + b\|_2$$

Maximum entropy $$z^* := \arg\min_z \sum_{c \in C} p_{\theta(k)}(c|G(z)) \log\left(p_{\theta(k)}(c|G(z))\right)$$

A human annotator needs reliably label the samples.

Many similar data points sampled at the transition of different classes may lead to sample bias.

Labeling and training on real samples is advantageous

Adversarial Sampling for Active Learning

Sample Matching

- Methods

  Use nearest neighbor and LSH-forest for closest matches based on Euclidean distance.

- Features

  PCA

  Autoencoder

  Discriminator

Experiment

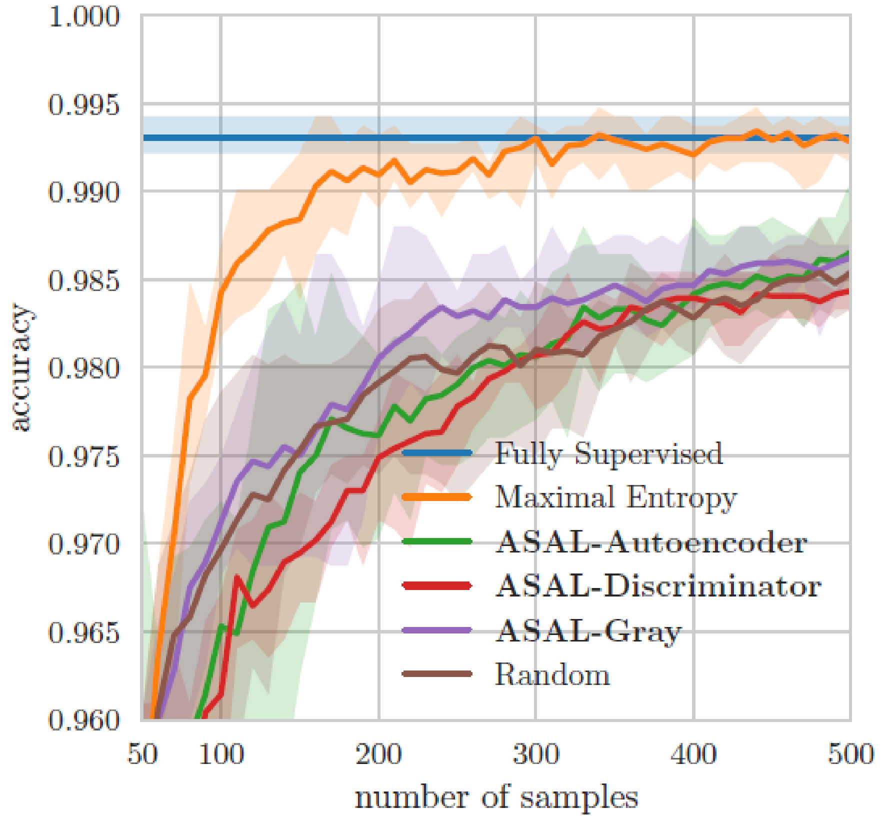

(a) DCGAN.

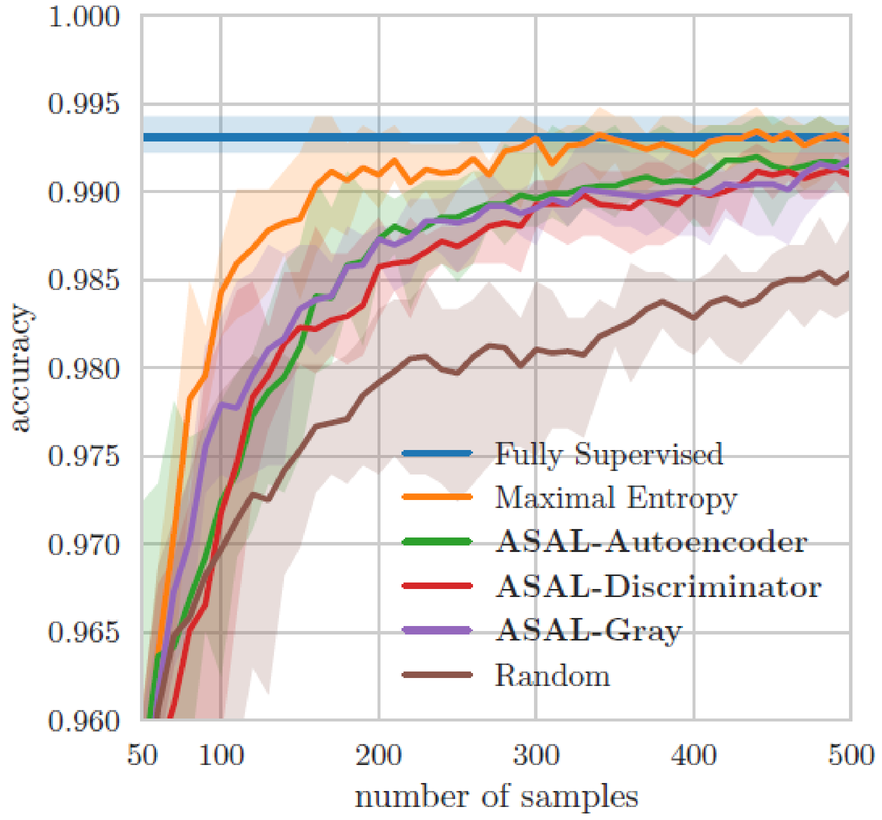

(b) WGAN-GP.

Experiment

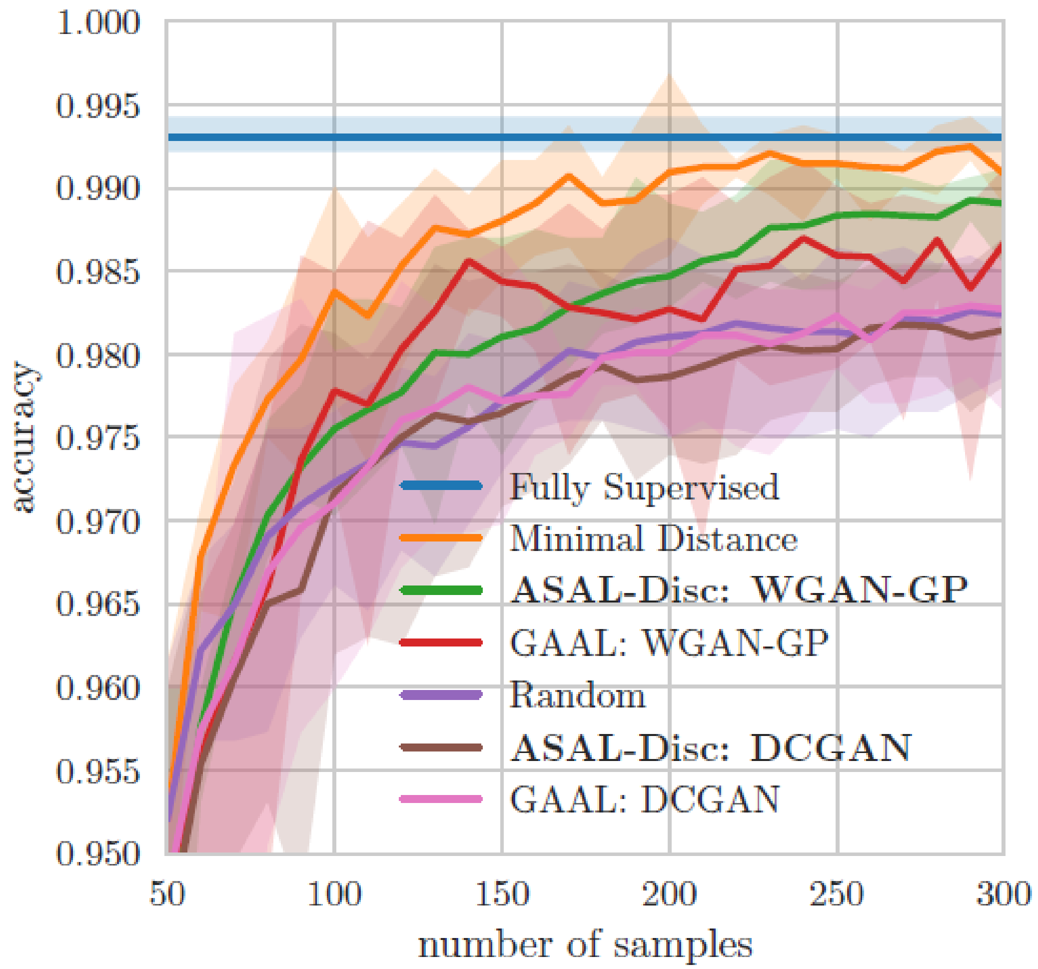

1.000
0.995
0.990
0.985
0.980
0.975
0.970
0.965
0.960
0.955
0.950
accuracy
50
100
150
200
250
300
number of samples
Fully Supervised
Minimal Distance
ASAL-Disc: WGAN-GP
GAAL: WGAN-GP
Random
ASAL-Disc: DCGAN
GAAL: DCGAN

Thanks