State-Relabeling Adversarial Active Learning

2020.9.2

Beichen Zhang1, Liang Li2*, Shijie Yang1, 2, Shuhui Wang2, Zheng-Jun Zha3, Qingming Huang1, 2, 4
1University of Chinese Academy of Sciences, Beijing, China

CVPR 2020

Outline

- Contribution
  - ➢ leverages both the annotation and the labeled/unlabeled state information
  - ➢ Relabel the unlabeled data’s state information
- How dose SRAAL work
  - ➢ Unified representation generator
  - ➢ State discriminator
  - ➢ online uncertainty indicator
- Experiment

Contribution

Active learning is to select the most valuable unlabeled sample to label

previous works have made full use of the annotation information of labeled data

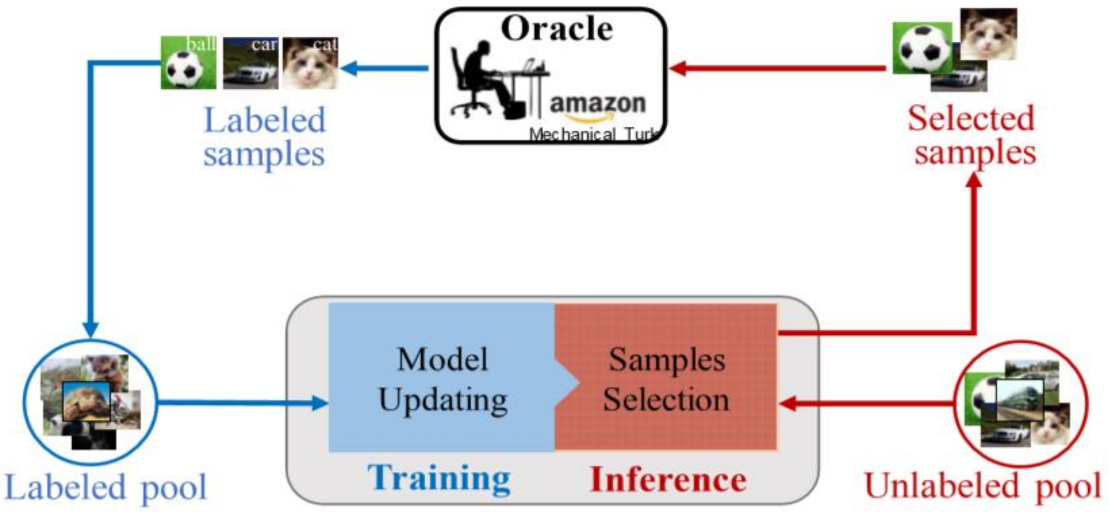

- Shortcoming :in the early iterations of sampling, the labeled pool is usually small, so that it restricts the ability to choose samples with high quality.

Contribution

Recent works focused on utilizing the state information of samples which indicate a sample is labeled (1) or unlabeled (0).

For example,
for Variational Adversarial Active Learning

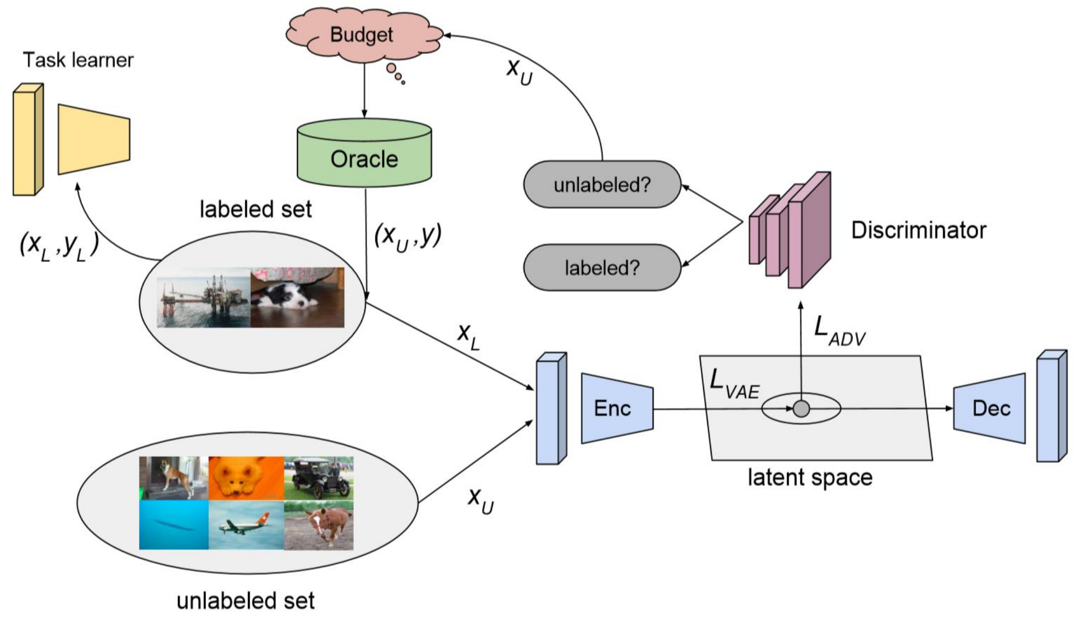

Model Framework

- VAE learns a latent space
- Discriminator is trained to discriminate between labeled and unlabeled

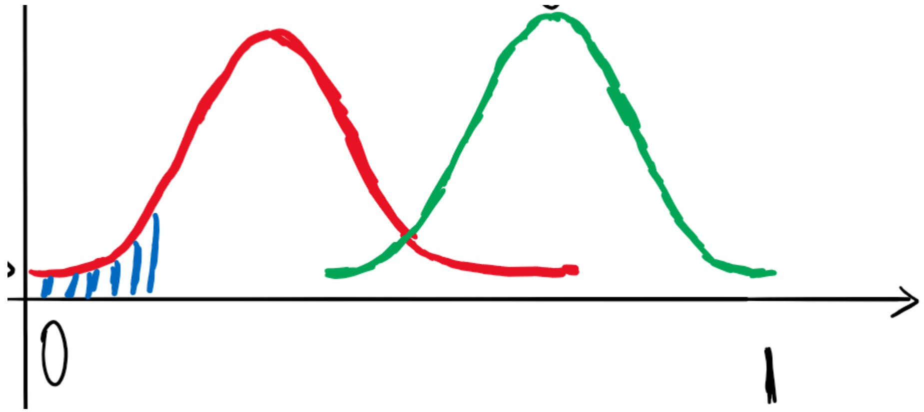

How dose SRAAL work?

- leverages both the annotation and the labeled/unlabeled state information
- The unlabeled state is a continuous value which can measure the importance of unlabeled sample

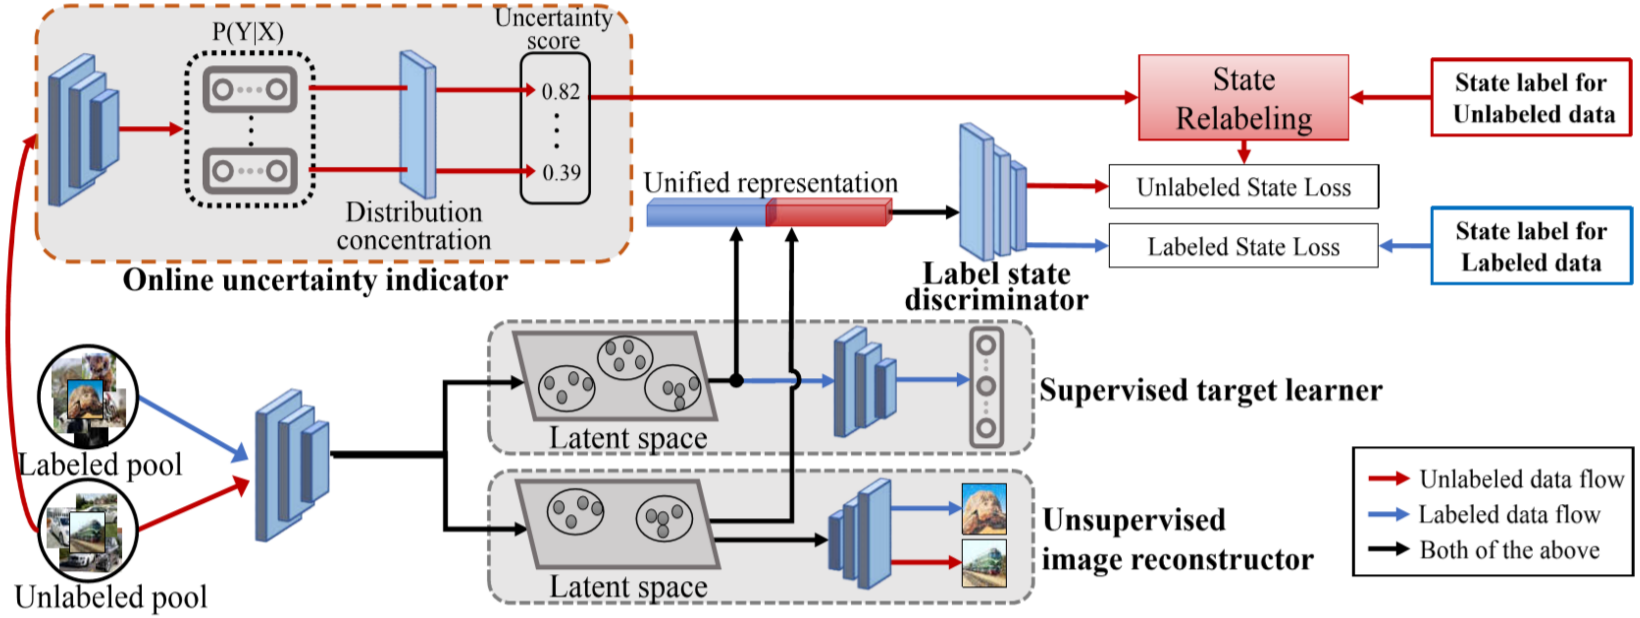

Figure 2. Network architecture of our proposed SRAAL. It consists of a unified representation generator and a labeled/unlabeled state discriminator. The generator embeds the annotation information into the final image features via the supervised target learner and unsupervised image reconstructor. Online uncertainty indicator is introduced to relabel the state of unlabeled samples and endues them with different importance. Finally, the state discriminator is updated through the labeled and unlabeled state losses, and helps select the more informative samples.

Generator

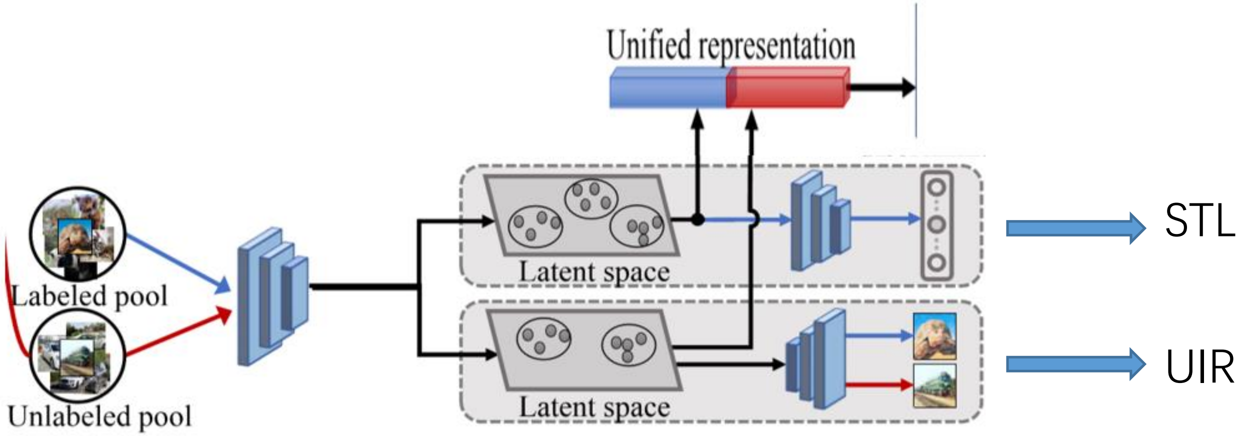

Part1: unsupervised image reconstructor (UIR)

- Standard VAE

$$\mathcal{L}^{UIR} = \mathcal{L}_U^{UIR} + \mathcal{L}_L^{UIR}$$
$$\mathcal{L}_U^{UIR} = E[log[p_\phi(x_U \mid z_U)] - D_{KL}(q_\theta(z_U \mid x_U) \parallel p(z))]$$
$$\mathcal{L}_L^{UIR} = E[log[p_\phi(x_L \mid z_L)] - D_{KL}(q_\theta(z_L \mid x_L) \parallel p(z))] \quad (1)$$

Part2: supervised target learner (STL)

- The decoder of STL varies with different tasks

$$\mathcal{L}_L^{STL} = E[log[p_\phi(y_L \mid z_L)]] - D_{KL}(q_\theta(z_L \mid x_L) \parallel p(z))] \quad (2)$$

Discriminator

The state information of unlabeled sample is Indicator(Xu)

Play the mini-max game between Generator and Discriminator

- Generator tries to trick the Discriminator into predicting that all data points are from the unlabeled pool(the state information is 1)

$$\mathcal{L}^{G}_{adv} = -E[log(D(q_\theta(z_L \mid x_L)))] - E[log(D(q_\theta(z_U \mid x_U)))]$$

- Discriminator learns how to discriminate the truth

$$\mathcal{L}^{D} = -E[log(D(q_\theta(z_L \mid x_L)))] - E[log(Indicator(x_U) - D(q_\theta(z_U \mid x_U)))]$$

$$\mathcal{L}^{G} = \lambda_1 \mathcal{L}^{UIR} + \lambda_2 \mathcal{L}^{STL}_{L} + \lambda_3 \mathcal{L}^{G}_{adv}$$

online uncertainty indicator

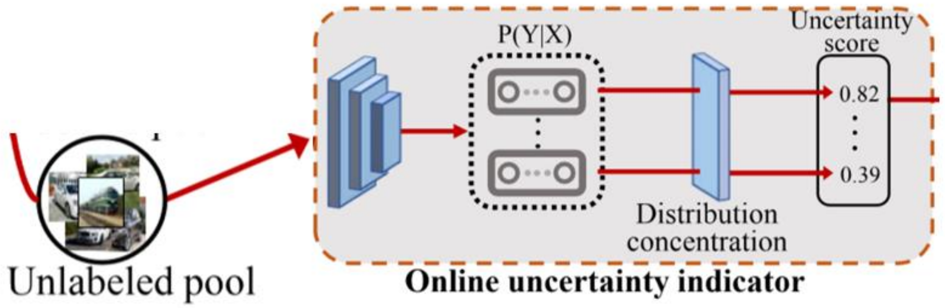

The OUI calculates the uncertainty score based on the prediction vector of the target model

- For image classification, the prediction is a probability vector for each category.
- For segmentation, each pixel has a probability vector.

online uncertainty indicator

$$Indicator(x_U) = 1 - \frac{MINVar(V)}{Var(V)} \times max(V)$$

MINVar(V) is the variance of the vector V', whose maximum element is the same with the V 's and other elements have the same value

For example, for classification Task

- Assume the V of Xu1 is [0.1, 0.2, 0.7], the V' is [0.15, 0.15, 0.7], the Indicator(Xu1) is 0.316
- Assume the V of Xu2 is [0.2, 0.3, 0.5], the V' is [0.25, 0.25, 0.5], the Indicator(Xu1) is 0.553

The properties of uncertainty score

(1) it has a boundary of [0,1);

(2) it has a negative correlation with the value of maximum probability;

(3) It has a positive correlation with the concentration of the probabilities distribution

Experiment

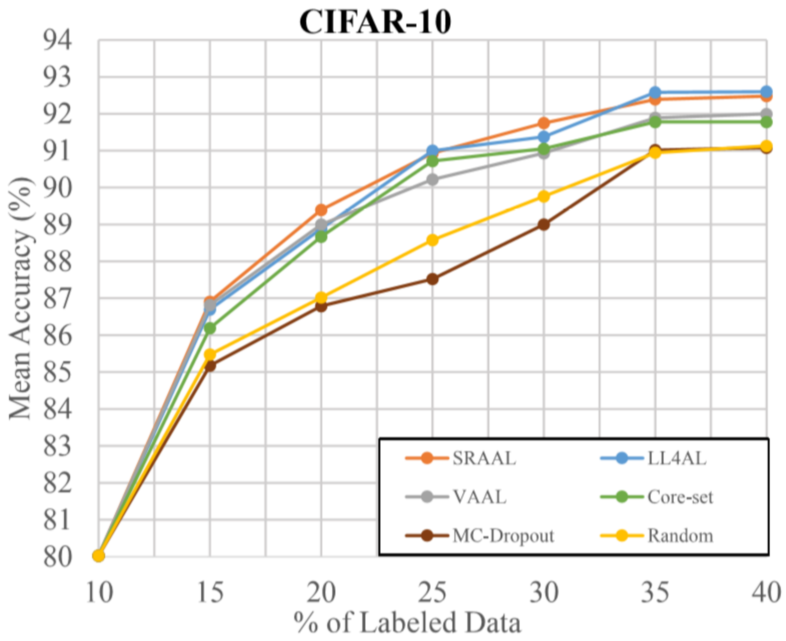
CIFAR-10
Mean Accuracy (%)
94
93
92
91
90
89
88
87
86
85
84
83
82
81
80
10
15
20
25
30
35
40
% of Labeled Data
SRAAL
LL4AL
VAAL
Core-set
MC-Dropout
Random

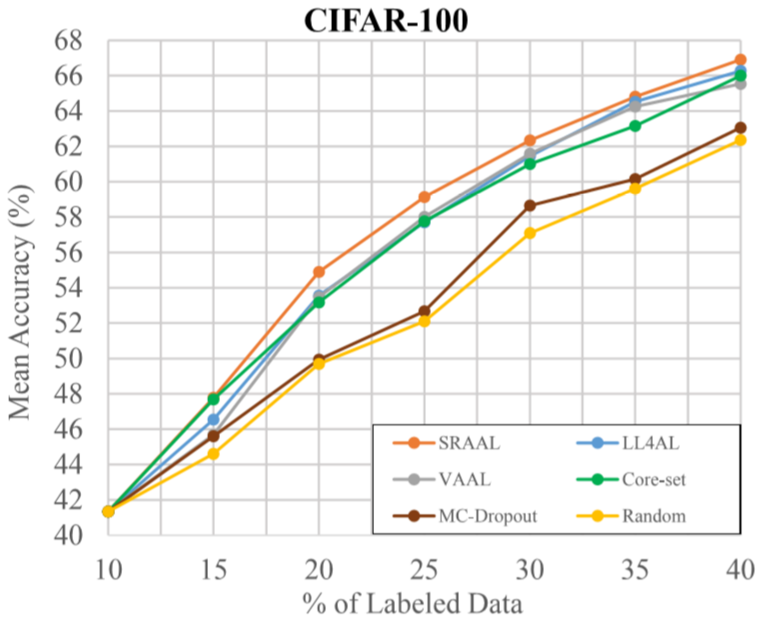
CIFAR-100
Mean Accuracy (%)
68
66
64
62
60
58
56
54
52
50
48
46
44
42
40
10
15
20
25
30
35
40
% of Labeled Data
SRAAL
LL4AL
VAAL
Core-set
MC-Dropout
Random

Experiment

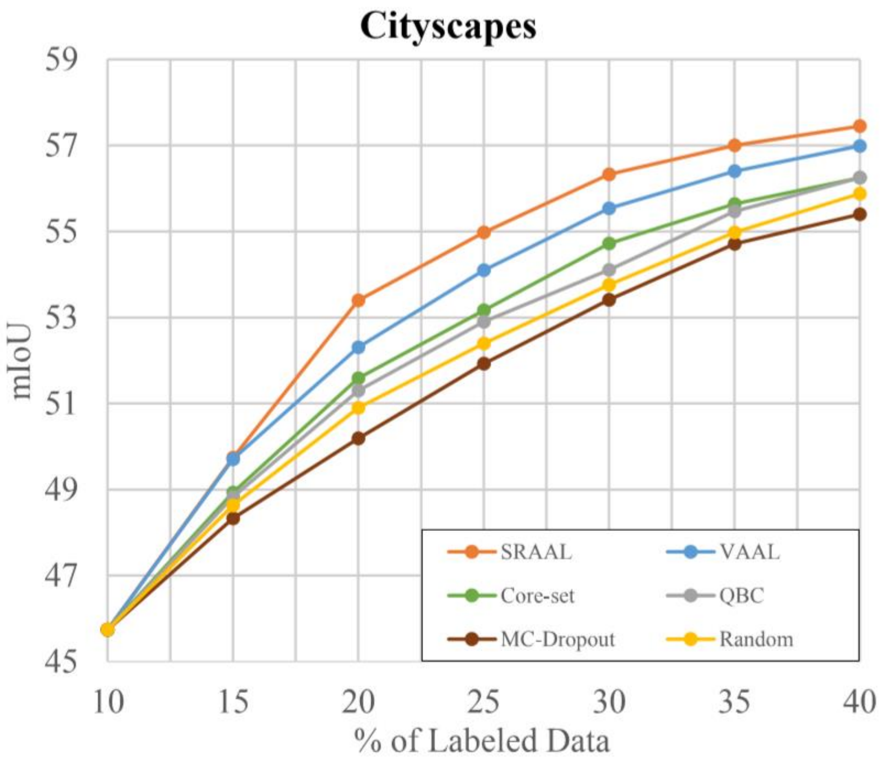
Cityscapes
mIoU
59
57
55
53
51
49
47
45
10
15
20
25
30
35
40
% of Labeled Data
SRAAL
VAAL
Core-set
QBC
MC-Dropout
Random

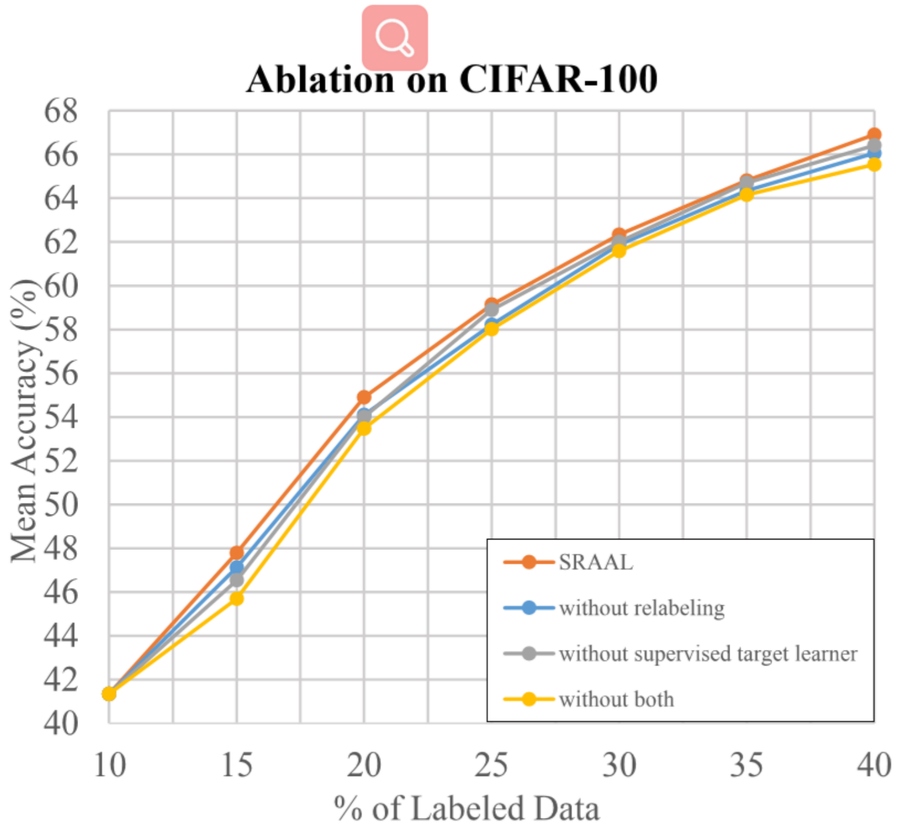
Ablation on CIFAR-100
Mean Accuracy (%)
68
66
64
62
60
58
56
54
52
50
48
46
44
42
40
10
15
20
25
30
35
40
% of Labeled Data
SRAAL
without relabeling
without supervised target learner
without both

Thanks