Generative-discriminative Feature Representations for Open-set Recognition

Pramuditha Perera*
Johns Hopkins University
pperera3@jhu.edu

Vlad I. Morariu
Adobe Research
morariu@adobe.com

Rajiv Jain
Adobe Research
rajijain@adobe.com

Varun Manjunatha
Adobe Research
vmanjuna@adobe.com

Curtis Wigington
Adobe Research
wigingto@adobe.com

Vicente Ordonez
University of Virginia
vicente@virginia.edu

Vishal M. Patel
Johns Hopkins University
vpatel36@jhu.edu

α_1	α_2	**Accuracy**	**AUC**
0.2	0.8	90.77	88.76
0.4	0.6	91.90	88.44
0.5	0.5	91.75	87.35
0.6	0.4	92.53	89.09
0.8	0.2	92.72	89.85

Table 1. Sensitivity of the algorithm to different hyper-parameters.

1. Impact of hyper-parameters

The proposed method optimizes a deep network with the objective of minimizing the loss $l_t = \alpha_1 l_c + \alpha_2 l_{ss}$. We carried out a study investigating performance of the proposed method for different values of α_1 and α_2. For this study, we used animal classes of CIFAR10 as known classes and considered remaining classes as open-set classes. It should be noted that this known-openset class split is different from all splits considered in the CIFAR10 experiment in the main paper. Parameters α_1 and α_2 were changed into different combinations between $0.2 - 0.8$ and performance of the system was analyzed. Seed was fixed in all these experiments to make sure comparisons are fair. In Table 1, closed-set classification accuracy and open-set detection performance(in AUC of ROC curve) are tabulated.

According to Table 1, AUC doesn't vary a lot for different combinations of α_1 and α_2. However, the highest performance is obtained when $\alpha_1 = 0.8$ and $\alpha_2 = 0.2$. General trend in Table 1 suggests that a better performance is obtained when a lower weight is assigned to the self-supervision loss. On the other hand, as shown in the ablation study of the main paper, removing self-supervision loss reduces open-set recognition performance. Based on these observations, we set hyper-parameters $\alpha_1 = 0.8$ and $\alpha_2 = 0.2$ for all experiments.

*This work was completed while the author was working as an intern at Adobe Research.