# Appendix A.

This supplementary material contains extra evidences to support our claim that knowledge preservation benefits domain adaptation. It is shown in the main paper that knowledge preservation helps the fine-tuned model to maintain good performance in the original domain. Here, under the zero-shot setting, we measure the similarity between a **target** zero-shot category and the **original** ImageNet categories, and demonstrate how this similarity correlates to the performance improvements brought by SAKE.

## A.1. Setting

We investigate similarities under the zero-shot setting. Different from the main paper as well as all previous work, we create a new held-out set of TU-Berlin, containing 30 categories which are **not** present in ImageNet. This is to make a fair comparison between different target categories. Experiments are performed three times, *i.e.*, we randomly choose three held-out sets, all of which have no category overlap with ImageNet.

We take a vanilla SE-ResNet-50 model, which was pre-trained on ImageNet and fine-tuned on the TU-Berlin reference set with only $\mathcal{L}_{\text{benchmark}}$. Then, we sort these 30 target categories by their mAP@all improvement achieved by SAKE, and equally divide them into 3 groups, "low", "medium" and "high", with each of them containing 10 categories of low, medium and high improvements brought by SAKE, respectively. Similarities between the target categories and the ImageNet categories are investigated separately in these 3 groups.

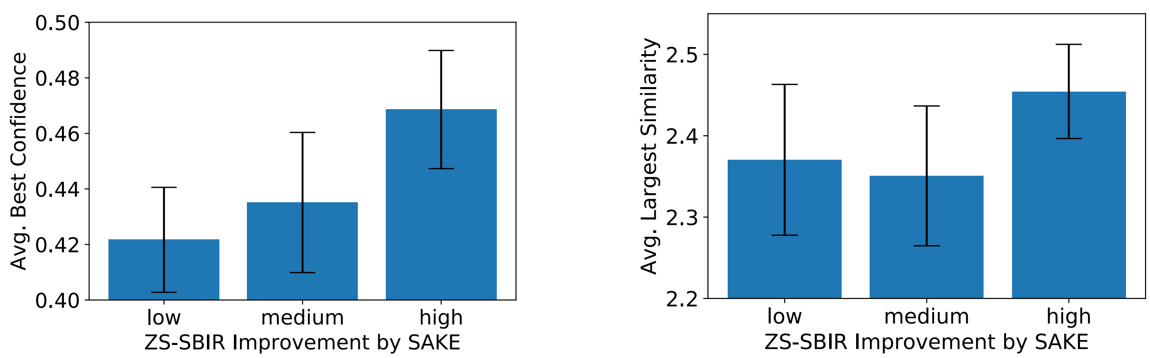


(a) ImageNet classification confidence with respect to the improvement brought by SAKE

(b) ImageNet synset similarity with respect to the improvement brought by SAKE

Figure 5: How the ZS-SBIR mAP@all improvement brought by SAKE on different categories correlates to the category-level similarity to the original ImageNet categories. Error bars and mean values are summarized from three repeated experiments with different held-out category sets, all of which have no overlap with ImageNet.

## A.2. Similarity by Classification Confidence

In Figure 5a, we can see that the improvement of SAKE becomes more significant when the category gets a higher classification confidence in an ImageNet-based classifier. This is to say, knowledge preservation, as expected, helps better to those categories that are closer to ImageNet – in other words, these categories are often heavier impacted by catastrophic forgetting, and knowledge preservation brings more improvement.

We shall point out that this phenomenon does not mean that knowledge preservation is not useful to those categories which are not contained in the original domain. Indeed, in each of these 3 groups, we observe accuracy gain under the zero-shot setting – this is also shown in our main experiments, in which both ImageNet and non-ImageNet categories largely benefit from knowledge preservation.

## A.3. Similarity by Semantic-space Distance

To provide another perspective, we perform the same experiment using the similarity defined by the Leacock-Chodorow Similarity on the WordNet synset, *i.e.*, the semantic tree used to build up ImageNet. In Figure 5b, we once again obtain the same conclusion, *i.e.*, SAKE is more effective on the categories that are better represented to the ImageNet semantic space.