

Adversarially Occluded Samples for Person Re-identification (Supplementary Material)

1. Improvement of Ranking Results

In this section we provide more illustrations of the improvement over the original model. The examples are shown in Figure 1, 2, 3 for dataset Market1501, CUHK03, DukeMTMC-reID respectively. In most cases the retrained model activates on larger regions on the body compared with the original model. Since more regions and patterns are considered by the re-trained model, we hypothesize that the feature vectors now contain more information for discriminating between different persons.




Figure 1: Some ranking examples on dataset Market1501 that the re-trained model improves over the original model. The four rows of each query case corresponds to (1) ranklist by the original model, (2) activation maps from the original model for the 1st row, (3) ranklist by the re-trained model, (4) activation maps from the re-trained model for the 3rd row. Images with green and red boundary denote true positive and distractors respectively.




Figure 2: Some ranking examples on dataset CUHK03 that the re-trained model improves over the original model. The four rows of each query case corresponds to (1) ranklist by the original model, (2) activation maps from the original model for the 1st row, (3) ranklist by the re-trained model, (4) activation maps from the re-trained model for the 3rd row. Images with green and red boundary denote true positive and distractors respectively.




Figure 3: Some ranking examples on dataset DukeMTMC-reID that the re-trained model improves over the original model. The four rows of each query case corresponds to (1) ranklist by the original model, (2) activation maps from the original model for the 1st row, (3) ranklist by the re-trained model, (4) activation maps from the re-trained model for the 3rd row. Images with green and red boundary denote true positive and distractors respectively.