Learning Memory-Augmented Unidirectional Metrics for Cross-modality Person Re-identification

1. The impact of unidirectional metric loss

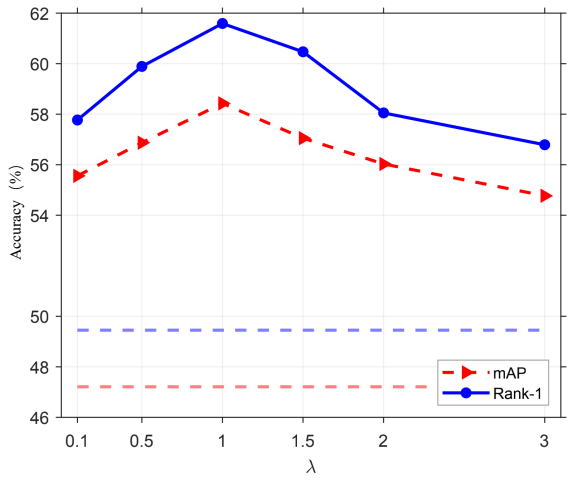

Figure 1. The results of parameter analysis for λ of Eq. 3 in manuscript. We evaluate the proposed MAUM on SYSU-MM01 under *all-search* test protocol. Rank-1 and mAP accuracy are reported.

The unidirectional metric loss is a key component of MAUM. λ (in Eq. 3) controls the strength of unidirectional metric loss in $\mathcal{L}_{\text{Total}}$ (Eq. 3). Fig. 1 experimentally analyzes the impact of λ on SYSU-MM01. We vary the λ from 0.1 to 3. We make two observations as follows.

First, when λ is very small, for example $\lambda = 0.1$, the unidirectional metric loss makes little contribution to the embedding optimization. The modality-specific loss plays the dominant role in embedding learning. Even so MAUM also obtains the better performance than baseline. It indicates that cross-modality re-ID benefits from the unidirectional metric loss.

Second, as λ increases, the performance of MAUM gradually increases to the maximum (when λ is 1.0) and then decreases. When λ is too large, the unidirectional metric loss makes more contribution than modality-specific loss in embedding leaning and shrinks the performance of MAUM.

Based on the above observations, we select $\lambda = 1$ in all our experiments, which balances the modality-specific loss and cross-modality unidirectional metric loss in Eq. 3.