Supplementary Material

1. Transfer the *p2p* Constraint to the *p2s* Constraint

As mentioned in the paper, the positive pairs and negative pairs under BCT should satisfy the following conditions on distance constraints based on $\forall\{i,j,k\}, y_i = y_j \neq y_k$.

$$\langle \phi_n(x_i), \phi_o(x_j) \rangle < \langle \phi_n(x_i), \phi_o(x_k) \rangle \tag{1}$$

$$\langle \phi_n(x_i), \phi_o(x_j) \rangle < \langle \phi_n(x_i), \phi_n(x_k) \rangle \tag{2}$$

With the triangle inequality, we can transfer this *p2p* constraint into a point-to-set (*p2s*) constraint as followings:

$$\begin{aligned} \|\phi_n(x_i) - \phi_o(x_j)\|_2 &= \|\phi_n(x_i) - E_o(X^c) + E_o(X^c) - \phi_o(x_j)\|_2 \\ &\leq \|\phi_n(x_i) - E_o(X^c)\|_2 + \|E_o(X^c) - \phi_o(x_j)\|_2 \end{aligned} \tag{3}$$

$$\begin{aligned} \|\phi_n(x_i) - E_o(X^c)\|_2 &= \|\phi_n(x_i) - \phi_o(x_j) + \phi_o(x_j) - E_o(X^c)\|_2 \\ &\leq \|\phi_n(x_i) - \phi_o(x_j)\|_2 + \|\phi_o(x_j) - E_o(X^c)\|_2 \end{aligned} \tag{4}$$

where $E_o(X^c)$ is the expectation of $\phi_o(X^c)$ and $X^c = \{x_i\}_{i=1}^n$ is the set of instances of class c, $\forall\{x_i, x_j\} \in X^c$.

Combined Eq. (3) with Eq. (4), we can draw the conclusion as follows:

$$\begin{aligned} \|\phi_n(x_i) - \phi_o(x_j)\|_2 \geq \|\phi_n(x_i) - E_o(X^c)\|_2 - \|\phi_o(x_j) - E_o(X^c)\|_2 \\ \|\phi_n(x_i) - \phi_o(x_j)\|_2 \leq \|\phi_n(x_i) - E_o(X^c)\|_2 + \|\phi_o(x_j) - E_o(X^c)\|_2 \end{aligned} \tag{5}$$

Because the $\|\phi_o(x_j) - E_o(X^c)\|_2$ is a constant, the range of $\|\phi_n(x_i) - \phi_o(x_j)\|_2$ is determined by $\|\phi_n(x_i) - E_o(X^c)\|_2$. Thus, by constraining the distance between $\phi_n(x_i)$ and $E_o(X^c)$, we can constrain the distance between $\phi_n(x_i)$ and $\phi_o(x_j)$. By what mentioned above, the *p2p* constraint can be transferred into a *p2s* constraint.

2. The Upgrade Process of the CMC Method

The method of cross-model compatibility (CMC) [1] aims to learn transformations which can relieve gaps between feature spaces of two models. The transformations T_n^o, T_o^n map features from one space to another. T_n^o and T_o^n represent the new-to-old transformation module and the old-to-new transformation module respectively. This method can be suitable to map two existing models. The Fig. 2 shows the process of updating online systems on the CMC model. For the CMC method, the upgrade still needs to maintain the new embedding model, the transformation module, and two databases.

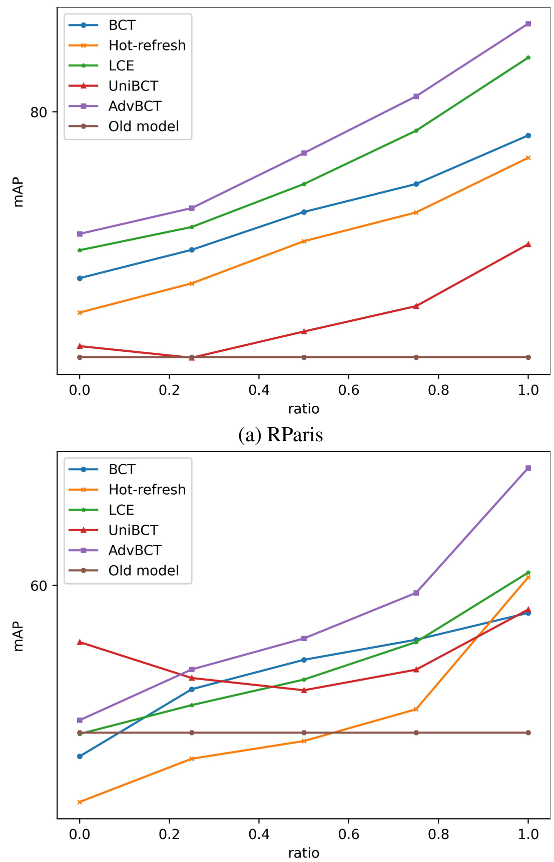

(a) RParis

(b) ROxford

Figure 1. The trend of retrieval performance during the hot-refresh upgrades on RParis and ROxford datasets. All results are tested on models trained on Extended-class setting.

3. Hot-refresh Model Upgrades

With BCT, query embeddings encoded by the new embedding model can be directly compared to the old embeddings in the old database. During the system upgrade, the old gallery embeddings will be replaced by the new embeddings progressively which is termed a hot-refresh upgrade. Referred to [2], we conducted experiments under the setting that the percentage of old embeddings re-indexed by the new embeddings increase from 0% to 100% to simulate the upgrade process. Fig. 1 shows the trend of retrieval

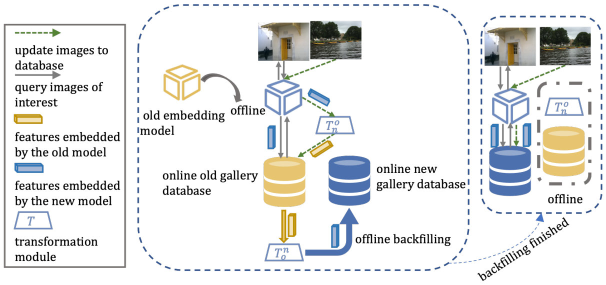

Figure 2. The process of updating online systems on the CMC model. T_n^o and T_o^n represent the new-to-old transformation module and the old-to-new transformation module respectively.

performance during the hot-refresh upgrades.

From the results, we can find that our *AdvBCT* performs best among all methods. As the ratio of new embeddings gradually increases, the retrieval performance becomes higher. In addition, the retrieval effect of *AdvBCT* is better than that of the old model during the entire database upgrade process, which shows that the iteration of our new model is effective. Furthermore, on ROxford dataset, *BCT* and *Hot-refresh* performs worse than the old model in some ratios when the hot-refresh upgrades proceed, and the retrieval performance of *UniBCT* degrades as Ref [2] call it model regression.

References

[1] Qiang Meng, Chixiang Zhang, Xiaoqiang Xu, and Feng Zhou. Learning compatible embeddings. In *CVPR*, pages 9939–9948, 2021. 1

[2] Binjie Zhang, Yixiao Ge, Yantao Shen, Yu Li, Chun Yuan, Xuyuan Xu, Yexin Wang, and Ying Shan. Hot-refresh model upgrades with regression-free compatible training in image retrieval. In *ICLR*, 2021. 1, 2