

# Towards Universal Backward-Compatible Representation Learning

Binjie Zhang<sup>1,2†</sup>, Yixiao Ge<sup>2†</sup>, Yantao Shen<sup>4‡</sup>, Shupeng Su<sup>2</sup>, Fanzi Wu<sup>4</sup>,  
 Chun Yuan<sup>1†</sup>, Xuyuan Xu<sup>3</sup>, Yexin Wang<sup>3</sup>, Ying Shan<sup>2</sup>

<sup>1</sup>Tsinghua University

<sup>2</sup>ARC Lab, Tencent PCG

<sup>3</sup>AI Technology Center of Tencent Video

<sup>4</sup>AWS/Amazon AI

{zbj19@mails,yuanc@sz}.tsinghua.edu.cn, {yixiao,yingsshan}@tencent.com

## Abstract

Conventional model upgrades for visual search systems require offline refreshment of gallery features by feeding gallery images into new models (dubbed as “backfill”), which is time-consuming and expensive, especially in large-scale applications. The task of backward-compatible representation learning [Shen *et al.*, 2020] is therefore introduced to support backfill-free model upgrades, where the new query features are interoperable with the old gallery features. Despite the success, previous works only investigated a close-set training scenario (*i.e.*, the new training set shares the same classes as the old one), and are limited by more realistic and challenging open-set scenarios. To this end, we first introduce a new problem of universal backward-compatible representation learning, covering all possible data split in model upgrades. We further propose a simple yet effective method, dubbed as **Universal Backward-Compatible Training (UniBCT)** with a novel structural prototype refinement algorithm, to learn compatible representations in all kinds of model upgrading benchmarks in a unified manner. Comprehensive experiments on the large-scale face recognition datasets MS1Mv3 and IJB-C fully demonstrate the effectiveness of our method. Source code is available at <https://github.com/TencentARC/OpenCompatible>.

## 1 Introduction

Deep learning-based methods [Ghifary *et al.*, 2016; He *et al.*, 2016; Li and Hoiem, 2017] have achieved great success in visual search tasks, such as face recognition [Liu *et al.*, 2017; Yang *et al.*, 2017; Wang *et al.*, 2017; Wang *et al.*, 2018; Deng *et al.*, 2019; Zhang *et al.*, 2020] and landmark retrieval [Philbin *et al.*, 2007; Philbin *et al.*, 2008; Weyand *et al.*, 2020; Radenović *et al.*, 2018; Ge *et al.*, 2020].

The task of visual search requires to retrieve the same objects’ images from a large-scale database (dubbed as *gallery*),

<sup>†</sup>Corresponding authors.

<sup>‡</sup>Work done when Binjie and Yantao are at ARC Lab.

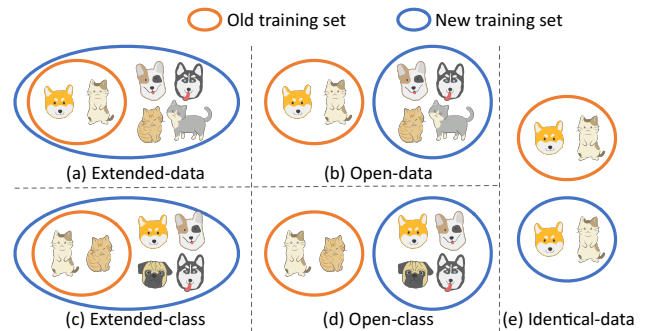


Figure 1: Illustration of different training data distributions for universal backward-compatible training. According to the data and category differences between old and new training sets, we summarize the data splits into five types from (a) to (e), covering most of the compatible training scenarios for backfill-free model upgrades.

given an image of interest (dubbed as *query*). The process of offline “backfilling” the gallery is always necessary for conventional model upgrades in retrieval systems, which is computationally expensive and time-consuming. Moreover, it is infeasible when the raw images are inaccessible due to privacy issues or storage limitations.

Thanks to the introduction of backward-compatible representation learning [Shen *et al.*, 2020; Zhang *et al.*, 2021], new models that are trained with compatibility constraints can be immediately deployed in a backfill-free manner, where the encoded new features for queries are interoperable with the old gallery features. The follow-up works make efforts to further improve the feature compatibility by designing advanced training constraints [Budnik and Avrithis, 2020; Meng *et al.*, 2021] or transformation architectures [Wang *et al.*, 2020]. Positive as the results are, they only focused on a single close-set model upgrading scenario (dubbed as extended-data in Figure 1 (a)), where the new training data share the identical class set as the old one. It is notable that the data split for model upgrades in real-world applications is complex and unpredictable, that is, both close-set and open-set scenarios should be considered. Existing methods [Shen *et al.*, 2020;

\*Since the upgraded (new) model is not directly comparable with the old gallery features, the gallery needs to be re-extracted via feeding all the raw images into the new model.

Wang *et al.*, 2020; Budnik and Avrithis, 2020; Meng *et al.*, 2021; Su *et al.*, 2022] did not investigate the open-set data split and are even inapplicable in such a scenario.

Towards this end, we for the first introduce the task of *universal* backward-compatible representation learning, where five kinds of data split covering both close-set and open-set scenarios are considered, as demonstrated in Figure 1. The open-set data split (including extended-class, open-data and open-class) poses a great challenge for learning compatible representations due to the potential domain gaps among different data and categories.

To tackle the challenge, we introduce a simple yet effective method, namely **Universal Backward-Compatible Training (UniBCT)**, to encode compatible representations in all kinds of data splits in a unified manner. Specifically, inspired by [Shen *et al.*, 2020], we utilize the old classifier (in the form of a fully-connected layer) to provide valuable supervision from the old latent space, *i.e.*, enforcing the new features to be closer to their corresponding old class centers. As for the novel categories in the open-class and extended-class scenarios, we extract the features of the new categories’ images and leverage their class centroids to construct pseudo prototypes. Due to the category gaps [You *et al.*, 2019], the pseudo prototype inevitably carries some noise that may affect the representation learning of backward compatibility. Therefore we propose to improve the class centroids of the pseudo prototype via a novel structural prototype refinement algorithm, *i.e.*, the “old” features of the new classes’ images are refined by propagating their neighbors’ knowledge via a fully-connected graph. The graph works under the assumption that visually similar images (measured by the new model which has the stronger capability) should have close-by old features.

In a nutshell, our contributions are three-fold. (1) We introduce a new task, namely universal backward-compatible representation learning, which aims at investigating all possible data splits in practical model upgrading scenarios. (2) We propose a novel method, dubbed as universal backward-compatible training (*UniBCT*), to tackle the challenge of different kinds of data splits in a unified manner. Our method is simple yet effective to refine the noisy pseudo prototype and improve the feature compatibility on both close-set and open-set scenarios. (3) We conduct comprehensive experiments on the large-scale face recognition datasets MS1Mv3 [Deng *et al.*, 2019] under five different model upgrading benchmarks, and investigate different compatibility constraints via the evaluations on IJB-C [Maze *et al.*, 2018]. Our *UniBCT* consistently outperforms the baseline and other advanced regularizations, fully indicating the effectiveness of our method.

## 2 Related Work

**Backward-Compatible Learning.** Backward-Compatible Learning aims to make new features and the old ones interoperable and realize backfill-free model upgrades. [Shen *et al.*, 2020] first formulated the problem by deriving influence loss from an empirical criterion, and solved it by utilizing the old classifier to regularize the optimization process. [Wang *et al.*, 2020] proposed Residual Bottleneck Transformation (RBT) blocks for feature embedding transferring. In [Budnik

and Avrithis, 2020], authors investigated the problem of the asymmetric test, where the database images are encoded by a teacher model and query images are encoded by a student model. A pair-based metric for instance-level image retrieval was proposed to achieve the goal. [Meng *et al.*, 2021] extended RBT blocks and designed advanced boundary loss to obtain more compact intra-class distributions. Though the above works could properly improve the compatible performance, they severely rely on the old training data or class. The open-set compatible scenarios are never investigated before.

**Universal Domain Adaptation.** While it is true that universal domain adaptation (UDA) [You *et al.*, 2019; Saito *et al.*, 2020] and our universal backward-compatible representation learning both take data/category domain gaps between old and new training data into consideration, they have entirely different purposes. UDA focuses on transferring the model knowledge from the old domain to the new one and only requires the model to perform well on the new domain without any cross-domain operations. Universal backward-compatible learning requires the new model to encode backward-compatible features that can be directly compared with the old features.

## 3 Universal Backward-Compatible Representation Learning

In this section, we first investigate the problem settings of universal backward-compatible representation learning in Sec. 3.1. Then we introduce our universal backward-compatible training (*UniBCT*) method in Sec. 3.2.

### 3.1 Problem Settings

Given the gallery features extracted by the old model, backward-compatible representation learning task requires the trained new model to encode query features that can be directly indexed by the old gallery features. In real-world applications, the new training set may differ from the old one in the aspects of data or classes, raising a universal backward-compatible representation learning problem.

**Symbol Definition.** We denote the training set, gallery set, and query set as  $\mathcal{D}$ ,  $\mathcal{G}$ ,  $\mathcal{Q}$ . An old model  $\phi_o$  trained on old training set  $\mathcal{D}_o$  embeds an image  $x$  to a feature vector  $v_o = \phi_o(x)$ . For model upgrades, a new model  $\phi_n$  trained on  $\mathcal{D}_n$  is obtained. The new model  $\phi_n$  embeds the image  $x$  into a new feature vector  $v_n$ .

**Benchmarks.** Taking both close-set and open-set model upgrading scenarios into consideration, we discuss five kinds of dataset settings as depicted in Table 1: (1) **Extended-data:** The old training set  $\mathcal{D}_o^{30\%-d}$  composes of 30% images which are randomly sampled from the whole dataset, and the new training set  $\mathcal{D}_n^{100\%}$  is made up of 100% data. The old and new training sets share the same classes. (2) **Open-data:** The new training data  $\mathcal{D}_n^{70\%-d}$  and the old data  $\mathcal{D}_o^{30\%-d}$  are exclusive from each other but they share the same classes. (3) **Extended-class:** We randomly pick 30% classes for the old training set  $\mathcal{D}_o^{30\%-c}$  and 100% classes for the new one  $\mathcal{D}_n^{100\%}$ . (4) **Open-class:** Both the data and the class are different between  $\mathcal{D}_o^{30\%-c}$  and  $\mathcal{D}_n^{70\%-c}$ . (5) **Identical-data:** The new training set  $\mathcal{D}_n^{30\%-d}$  and the old one  $\mathcal{D}_o^{30\%-d}$  are identical.

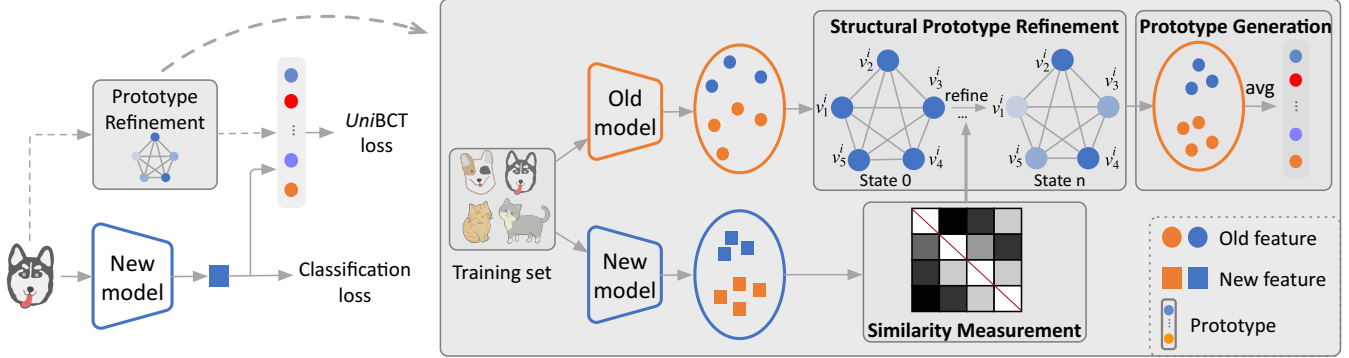


Figure 2: Pipeline of our Universal Backward-Compatible Training (*UniBCT*). The new model is supervised by a classification loss to learn discriminative features, and an additional universal backward-compatible loss to make sure the new features be interchangeable with the old ones. To alleviate the negative effects of data and category gaps between old and new training data in open-set model upgrading scenarios, we introduce a novel module named structural prototype refinement. It improves the old feature quality by propagating their neighbors’ knowledge via a fully-connected graph. Note that during the training process, the pseudo prototypes will be not updated by the loss backpropagation.

Allocation type	Old train-set		New train-set	
	# images	# classes	# images	# classes
Extended-data	1,511,514	93,431	5,179,510	93,431
Open-data	1,511,514	93,431	3,667,996	93,431
Extended-class	1,549,785	28,029	5,179,510	93,431
Open-class	1,549,785	28,029	3,629,725	65,402
Identical-data	1,511,514	93,431	1,511,514	93,431

Table 1: Five different allocations for the training data, where all the images are sampled from MS1Mv3. The “extended-data”, “open-data” and “identical-data” setups share the same old training set.

**Compatibility Evaluation.** Cross-model compatibility means that the gallery features produced by  $\phi_o$  can be directly comparable with the query features extracted by  $\phi_n$ . Following [Shen *et al.*, 2020], we claim that the feature compatibility is achieved if the following empirical criterion is satisfied,

$$\mathcal{M}(\phi_n, \phi_o; \mathcal{Q}, \mathcal{G}) > \mathcal{M}(\phi_o, \phi_o; \mathcal{Q}, \mathcal{G}), \quad (1)$$

where  $\mathcal{M}$  is an evaluation metric for the corresponding test set. **Cross Test**, denoted as  $\mathcal{M}(\phi_n, \phi_o; \mathcal{Q}, \mathcal{G})$ , is the query-to-gallery retrieval performance, where query features are extracted by new model  $\phi_n$  and gallery ones with old model  $\phi_o$ . **Self Test** reflects the performance where query and gallery features are extracted by the same model (*e.g.*, the old one).

### 3.2 Universal Backward-Compatible Training

To achieve compatibility in new model training process, two universal objectives are essential: (1) obtaining discriminative feature representations for better performance, and (2) making old and new representation features interoperable. The overall training objective of our universal backward-compatible training can be therefore formulated as

$$\mathcal{L} = \mathcal{L}_{\text{cls}} + \eta \mathcal{L}_{\text{uniBCT}}, \quad (2)$$

where  $\mathcal{L}_{\text{cls}}$  is the classification loss to achieve the first goal,  $\mathcal{L}_{\text{uniBCT}}$  is the universal backward-compatible loss to achieve the second goal and  $\eta$  is the loss weight.

Specifically, following the state-of-the-art method in metric learning, we use the form of ArcFace loss [Deng *et al.*, 2019] to regularize the pretext task of classification, that is,

$$\mathcal{L}_{\text{cls}} = \ell_{\text{arc}}(\omega_n, \phi_n), \quad (3)$$

where  $\omega_n$  and  $\phi_n$  denote the classifier and backbone of the new model. The formulation of ArcFace loss is

$$\ell_{\text{arc}}(\omega, \phi) = -\frac{1}{|\mathcal{D}_n|} \sum_{x \in \mathcal{D}_n} \log \frac{e^{s(\cos(\theta_y + m))}}{e^{s(\cos(\theta_y + m))} + \sum_{j \neq y} e^{s \cos \theta_j}}, \quad (4)$$

where  $y$  is the label of the training image  $x$ .  $s$  is a scale factor,  $m$  is the margin, and  $\theta_y = \arccos \langle \omega^y, \phi(x) \rangle$  is the angle between the weight  $\omega^y$  ( $y$ -th prototype of the classifier  $\omega$ ) and the feature  $\phi(x)$ . With  $\mathcal{L}_{\text{cls}}$ , the new model can be properly trained to encode discriminative representations for self test.

According to [Shen *et al.*, 2020], the old classifier (on top of the old backbone model) embeds the characteristic (*i.e.*, class prototypes) of the old latent space, which can be directly leveraged as the valuable supervision in close-set compatible training. However, in the open-set benchmarks of our universal backward-compatible representation learning task, the off-the-shelf old classifier is inapplicable due to the novel new classes. Intuitively, to overcome this limitation, we can modify the off-the-shelf old classifier into a pseudo classifier via (1) extracting the features of the new training set by the old model, and (2) using their class centroids as the pseudo classifier weights. We denote the pseudo old classifier as  $\hat{\omega}_o$ , and the backward-compatible loss can be formulated as

$$\mathcal{L}_{\text{uniBCT}} = \ell_{\text{arc}}(\hat{\omega}_o, \phi_n), \quad (5)$$

where  $\ell_{\text{arc}}$  is the form of ArcFace loss.  $\mathcal{L}_{\text{uniBCT}}$  regularizes to push the new features be closer to their corresponding old class centroids in order to align the old and new latent spaces.

It is notable that the quality of pseudo old prototypes is essential to the training of feature backward compatibility. Due to the domain gap (including data gap and category gap) between old and new training sets in open-set model upgrading scenarios, the pseudo old prototypes generated by the simple average operation inevitably carry some noise, affecting the representation learning of backward compatibility. To tackle the challenge, we introduce a novel structural algorithm to refine the prototypes via a fully-connected graph.

**Structural Prototype Refinement.** As illustrated in Figure 2, we improve the old prototypes through knowledge propagation under the assumption that visually similar samples of the same class should have close-by old features. We use the training new model to measure their similarities since the new model is expected to have the stronger model capability and could encode more discriminative representations for more accurate similarity measurement.

Specifically, we construct a fully-connected undirected graph  $G = (V, E)$  for each class, where  $V$  and  $E$  represent its vertices and edges. In our context, each *old* feature  $v_o \in \mathbb{R}^d$  serves as a vertex, and features of the same class can be denoted as a matrix  $V \in \mathbb{R}^{m \times d}$ , where  $d$  is the feature dimension and  $m$  is the sample number for a certain class. The edges among vertices are the similarity scores between pairwise samples, which are measured by cosine similarity, *i.e.*,  $\langle v_n^i, v_n^j \rangle$ . Note that we use *new* model features to measure the similarity. All the edges of a graph  $G$  can be denoted as a symmetric matrix  $E$ . And we further normalize it by row,

$$\tilde{E}(i, j) = \begin{cases} \frac{\exp(\langle v_n^i, v_n^j \rangle / \tau)}{\sum_{j \neq i} \exp(\langle v_n^i, v_n^j \rangle / \tau)}, & i \neq j \\ 0, & i = j \end{cases} \quad (6)$$

where  $\tau$  is the temperature hyper-parameter and a lower temperature leads to a sharper probability distribution.

Each node in the graph randomly visits neighbor images driven by transition probabilities (*i.e.*, similarity scores). Similar nodes (neighbors) are enhanced by each other and closer to the real center of the current class. The outlier features would also be rectified by other nodes. Such a propagation process can be formulated as,

$$V^{(t)} = \tilde{E}V^{(t-1)}, \quad (7)$$

where  $t$  is the iteration times.

The initial feature matrix  $V^{(0)}$  is aggregated to avoid potential collapse in the propagation process, that is,

$$V^{(t)} = \lambda \tilde{E}V^{(t-1)} + (1 - \lambda)V^{(0)}, \quad (8)$$

where  $\lambda \in [0, 1]$  is the aggregation weight. When  $t$  tends to infinity, Eq. (8) has a converged close form,

$$V^{(\infty)} = (1 - \lambda)(I - \lambda \tilde{E})^{-1}V^{(0)}, \quad (9)$$

where  $I$  is an identity matrix and  $(\cdot)^{-1}$  denotes matrix inverse operation. Once  $V^\infty$  is obtained, the class prototype  $\hat{\omega}_o$  could be computed by column-wise average pooling of  $V^\infty$ ,

$$\hat{\omega}_o(j) = \frac{1}{m} \sum_{i=1}^m V^{(\infty)}(i, :), \quad (10)$$

where  $\hat{\omega}_o(j) \in \mathbb{R}^d$  is the  $j$ -th pseudo prototype, and  $m$  is the number of vertices belonging to the  $j$ -th class. The refined class prototype  $\hat{\omega}_o$  are used as supervision signals in the universal backward compatible loss  $\mathcal{L}_{\text{uniBCT}}$  (Eq. (5)).

Compared with the vanilla average-based prototype, our introduced structural prototype refinement effectively alleviates the outlier effects by propagating and aggregating the knowledge from neighbor features of the same class.

## 4 Experiments

To perform a thorough evaluation, we estimate our method (*UniBCT*) under all compatible settings on the large-scale face recognition dataset. Satisfying results indicates the effectiveness and robustness of our approach.

### 4.1 Experimental Setup

**Datasets.** **MS-Celeb-1M (MS1M)** [Guo *et al.*, 2016] is a large-scale face recognition training dataset, which consists of about 10 million images with 1 million identities. Since the original MS1M dataset includes abundant noisy images, we adopt MS1Mv3 [Deng *et al.*, 2019] as the training set, which is made up of 5,179,510 training images with 93,431 labels. **IJB-C** [Maze *et al.*, 2018], a challenging benchmark, is utilized as the open-set evaluation dataset, which has around 1.3 million images. For verification task, there are 469,376 templates pairs. For identification task, the query set contains 19,593 images and the gallery set consists of 3,531 images.

**Metric.** We employ two standard test protocols in face recognition: (1) 1:1 verification calculates the true acceptance rate (TAR) at different false acceptance rates (FAR) for template pairs. In Cross-Test, we extract the first template with the new model, and the second with the old model. (2) 1:N identification evaluates the retrieval accuracy at top-k. In Cross-Test, we process the query set (prob images) and the gallery set (template images) with the new and old model, respectively.

**Training Details.** We use 4 NVIDIA V100 GPUs for training. The training index file is split with fixed random seed 666. We adopt ResNet18 and ResNet50 [He *et al.*, 2016] architectures as the backbones of the old and new models; one Fully Connected layer is followed to project the output dimension into 512. We adopt standard stochastic gradient descent (SGD) to optimize the model parameters. The learning rate is set to 0.1 and decreases 10 times at the 20th, 26th and 32th epoch. The training stops after 35 epochs. The weight decay is set to  $10^{-4}$  and momentum is 0.9. Batch size is set to 256. The scale factor  $s$  and margin  $m$  in Eq. 4 are 64, 0.5 following the default setting<sup>†</sup>. In graph-based prototype refinement, we set  $\lambda$  to 0.9,  $T$  to 0.05.

### 4.2 Analysis of UniBCT

**Effectiveness of Structural Prototype Refinement.** Since the quality of pseudo prototypes has essential impact on the backward-compatible learning, we introduce a structural prototype refinery mechanism to improve old features by allocating knowledge from their neighbors. As illustrated in Table 2, our method ( $\mathcal{L}_{\text{uniBCT}}$ ) not only fulfills the requirement of feature

<sup>†</sup><https://github.com/deepsight/insightface>

Scenarios	Model <sub>old</sub>	Model <sub>new</sub>	Training Set	Comp. Loss	1:1 Verification		1:N Identification			
					Cross Test	Self Test	Cross Test		Self Test	
					TAR@FAR	TAR@FAR	Top1	Top5	Top1	Top5
Extended-data	$\phi_o^{r18}$		$D_o^{30\%-d}$	-	-	93.36	-	-	69.90	75.88
		$\phi_n^{r18}$	$D_n^{100\%}$	-	-	96.35	-	-	80.67	85.14
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{100\%}$	$\mathcal{L}_{\text{regress}}$	0.12	94.78	8.12	10.43	76.34	80.88
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{100\%}$	$\mathcal{L}_{\text{contra}}$	92.26	94.58	<b>73.36</b>	81.35	<b>80.90</b>	<b>85.99</b>
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{100\%}$	$\mathcal{L}_{\text{uniBCT}}^*$	93.88	94.62	72.46	81.25	80.51	84.78
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{100\%}$	$\mathcal{L}_{\text{uniBCT}}$	<b>94.13</b>	<b>94.85</b>	72.89	<b>81.77</b>	80.83	85.95
Open-data	$\phi_o^{r18}$		$D_o^{30\%-d}$	-	-	93.36	-	-	69.90	75.88
		$\phi_n^{r18}$	$D_n^{70\%-d}$	-	-	94.28	-	-	75.55	80.24
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{70\%-d}$	$\mathcal{L}_{\text{regress}}$	0.02	94.51	7.36	9.12	73.21	78.84
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{70\%-d}$	$\mathcal{L}_{\text{contra}}$	92.23	94.42	70.34	78.20	76.69	81.75
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{70\%-d}$	$\mathcal{L}_{\text{uniBCT}}^*$	93.75	94.37	70.35	77.68	76.54	81.69
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{70\%-d}$	$\mathcal{L}_{\text{uniBCT}}$	<b>94.18</b>	<b>94.52</b>	<b>71.42</b>	<b>79.14</b>	<b>76.88</b>	<b>81.92</b>
Extended-class	$\phi_o^{r18}$		$D_o^{30\%-c}$	-	-	92.95	-	-	68.84	74.72
		$\phi_n^{r18}$	$D_n^{100\%}$	-	-	96.35	-	-	80.67	85.14
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{100\%}$	$\mathcal{L}_{\text{regress}}$	0.08	93.21	7.55	9.67	74.15	78.72
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{100\%}$	$\mathcal{L}_{\text{contra}}$	92.70	94.53	71.83	<b>79.26</b>	78.43	83.76
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{100\%}$	$\mathcal{L}_{\text{uniBCT}}^*$	93.54	94.32	71.67	79.33	78.51	84.14
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{100\%}$	$\mathcal{L}_{\text{uniBCT}}$	<b>93.75</b>	<b>94.55</b>	<b>72.02</b>	79.13	<b>78.84</b>	<b>84.33</b>
Open-class	$\phi_o^{r18}$		$D_o^{30\%-c}$	-	-	92.95	-	-	68.84	74.72
		$\phi_n^{r18}$	$D_n^{70\%-c}$	-	-	94.28	-	-	75.55	80.24
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{70\%-c}$	$\mathcal{L}_{\text{regress}}$	0.01	92.78	6.88	8.12	70.26	75.95
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{70\%-c}$	$\mathcal{L}_{\text{contra}}$	92.51	<b>94.24</b>	66.51	75.82	73.63	79.96
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{70\%-c}$	$\mathcal{L}_{\text{uniBCT}}^*$	93.35	93.96	67.14	76.38	74.21	80.28
	$\phi_o^{r18}$	$\phi_n^{r18}$	$D_n^{70\%-c}$	$\mathcal{L}_{\text{uniBCT}}$	<b>93.46</b>	94.10	<b>67.47</b>	<b>77.01</b>	<b>74.79</b>	<b>81.22</b>
Identical-data	$\phi_o^{r18}$		$D_o^{30\%-d}$	-	-	93.36	-	-	69.90	75.88
		$\phi_n^{r50}$	$D_n^{30\%-d}$	-	-	94.97	-	-	70.21	76.34
	$\phi_o^{r18}$	$\phi_n^{r50}$	$D_o^{30\%-d}$	$\mathcal{L}_{\text{regress}}$	0.11	93.78	7.73	9.35	67.41	73.43
	$\phi_o^{r18}$	$\phi_n^{r50}$	$D_o^{30\%-d}$	$\mathcal{L}_{\text{contra}}$	92.53	95.58	64.67	<b>74.69</b>	70.49	77.13
	$\phi_o^{r18}$	$\phi_n^{r50}$	$D_o^{30\%-d}$	$\mathcal{L}_{\text{uniBCT}}^*$	94.40	95.42	67.38	73.25	70.57	78.34
	$\phi_o^{r18}$	$\phi_n^{r50}$	$D_o^{30\%-d}$	$\mathcal{L}_{\text{uniBCT}}$	<b>94.59</b>	<b>95.63</b>	<b>67.71</b>	73.81	<b>70.66</b>	<b>78.76</b>

Table 2: Comparison of baselines and our proposed approach on IJB-C dataset in universal backward-compatible scenarios, including five different benchmarks. The architectures are ResNet18 (r18) and ResNet50 (r50).  $\mathcal{L}_{\text{uniBCT}}^*$  denotes the *vanilla* version of universal backward-compatible loss where the pseudo prototypes are simply averaged over the raw old features.  $\mathcal{L}_{\text{uniBCT}}$  uses our introduced structural prototype refinement algorithm to improve the pseudo classifier and achieves the optimal performance. We evaluate all models in two aspects: (1) For 1:1 verification, the first and second templates are extracted by the new and old model in Cross-Test (CT), and they are processed by the same new model in Self-Test (ST). TAR@FAR= $1e^{-4}$  is adopted as the compatible metric. (2) For 1:N Identification, the query and gallery set are extracted by the new and old models respectively in CT. We report the retrieval accuracy in terms of top1 and top5.

compatible training, but also well boosts the baseline method ( $\mathcal{L}_{\text{uniBCT}}^*$ ), which adopts vanilla prototypes for training.

In addition, an alternative approach for refining pseudo prototype is to discard outlier samples which are far away from the class centroids. Specifically, we filter out the top-10% data that is away from the mean feature vector in each class, and utilize the rest features to generate the class prototype. As

shown in Table 3, it (denoted as “drop avg.”) performs worse than the proposed graph-based refinement (denoted as “refined avg.”). That is because the distribution of the old features is noisy and unreliable, the drop strategy only refers to the old distribution while our structural refinement utilizes the sample similarities in the new latent space as propagation guidance.

Method	Prototype	1:1 Verification		1:N Identification		
		TAR@FAR	Comp.?	Top1	Top5	Comp.?
$\phi_o^{r18}$	-	93.36	-	69.90	75.88	-
Ours	vanilla avg.	93.88	✓	72.46	81.25	✓
Ours	drop avg.	94.03	✓	72.35	80.97	✓
Ours	refined avg.	94.13	✓	72.89	81.77	✓

Table 3: The comparison of different prototypes for the old pseudo classifier. ‘‘Refined avg.’’ denotes our optimal solution of structural prototype refinement. The results are reported on IJB-C (extended-data) in terms of 1:1 verification (TAR@FAR=1e<sup>-4</sup>).

**Compare to Other Form of Constraints.** The old prototype represents the global contents of the old model, in the meanwhile, each old feature indicates local details. Therefore, directly maximizing the similarity between the new feature and the corresponding old feature is an alternative choice to achieve compatibility. Specifically, one direct way is to minimize the Euclidean distance between the old and new features extracted from the same image:

$$\mathcal{L}_{\text{regress}}(\phi_n, \phi_o) = -\frac{1}{|\mathcal{D}_n|} \sum_{x \in \mathcal{D}_n} \|\phi_n(x) - \phi_o(x)\|^2. \quad (11)$$

As demonstrated in Table 2, we notice that feature regression fails in all settings. The reason might be that simply minimizing the distance between positive pairs is not enough. Thus we turn to another solution, *i.e.*, pulling the new-old positive pairs close and pushing away the negative pairs in the form of contrastive learning. Considering each new feature  $(\phi_n(x^i), y_i)$  as the anchor, the positive set consists of old features with the same class  $\mathcal{P}(i) = \{\phi_o(x^j) | y_j = y_i\}$ , and the negative set is comprised of the other old features  $\mathcal{N}(i) = \{\phi_o(x^j) | y_j \neq y_i\}$ . To simplify the training process, we only consider one positive pair  $(\phi_n(x^i), \phi_o(x^i))$ . The compatible loss is formulated as,

$$\begin{aligned} \mathcal{L}_{\text{contra}}(\phi_n, \phi_o) \\ = -\frac{1}{|\mathcal{D}_n|} \sum_{x^i \in \mathcal{D}_n} \log \frac{e^{(\phi_n(x^i) \cdot \phi_o(x^i) / \tau)}}{\sum_{k \in \{x^i, \mathcal{N}(i)\}} e^{(\phi_n(x^i) \cdot \phi_o(k) / \tau)}}, \end{aligned} \quad (12)$$

where  $\tau$  is a temperature hyper-parameter. As shown in Table 2, the performance of *UniBCT* surpasses the other losses in terms of Cross Test and Self Test. Our *UniBCT* loss adopts the classification-like form following [Shen *et al.*, 2020] and it considers global intra-class and inter-class relations. In contrast, the contrastive loss (Eq. (12)) only considers the classes in the current mini-batch, neglecting the global information.

**Close-set v.s. Open-set.** For 1:1 verification task, our method achieves remarkable performance in all close-set and open-set scenarios. For 1:N identification task, the empirical criterion (Eq. (1)) is satisfied in most practical settings, except for the most challenging scenario (*i.e.*, open-class), demonstrating that *UniBCT* can properly alleviate the category gap but cannot entirely solve it. Even though, we still outperforms other competing methods, indicating the effectiveness of *UniBCT*.

Method	Comp. Loss	1:1 Verification		1:N Identification		
		TAR@FAR	Comp.?	Top1	Top5	Comp.?
$\phi_o^{r18}$	-	93.36	-	69.90	75.88	-
AML	$\mathcal{L}_{\text{regress}}$	0.12	✗	8.12	10.43	✗
BCT	$\mathcal{L}_{\text{BCT}}$	94.01	✓	72.64	81.49	✓
Ours	$\mathcal{L}_{\text{uniBCT}}^*$	93.88	✓	72.46	81.25	✓
Ours	$\mathcal{L}_{\text{uniBCT}}$	94.13	✓	72.89	81.77	✓

Table 4: Compare to state-of-the-art backward-compatible training methods on IJB-C (extended-data). Only extended-data is evaluated here since BCT is inapplicable for other open-set benchmarks. The results are reported in terms of 1:1 verification (TAR@FAR=1e<sup>-4</sup>).

### 4.3 Comparison with State-of-the-arts

To indicate our approach *UniBCT* can consistently surpasses previous compatible training methods in the conventional close-set benchmarks, we conduct the comparison experiments on the ‘‘extended-data’’ setup. As shown in Table 4, we compare with BCT [Shen *et al.*, 2020] and AML [Budnik and Avrithis, 2020]. Note that [Wang *et al.*, 2020] and [Meng *et al.*, 2021] are not listed as they require extra network parameters which is not fair.

AML aims to enlarge the similarity of positive pairs, which is the same as the regression loss in Eq. (11). AML fails to achieve compatibility in face recognition task though it works well in landmark retrieval in its original paper. Regression loss only focuses on decreasing the distance between positive pairs while ignoring the distance restriction between negative pairs, leading to unsatisfactory performance in fine-grained retrieval tasks, like face recognition.

As we introduced in the method section, in the close-set setup, the off-the-shelf old classifier can directly serve as the old prototypes according to [Shen *et al.*, 2020]. To first investigate the difference between the off-the-shelf old classifier ( $\omega_o$ ) and the pseudo classifier ( $\hat{\omega}_o$ ), we compare  $\mathcal{L}_{\text{BCT}}$  and  $\mathcal{L}_{\text{uniBCT}}^*$  (vanilla avg.) in Table 4. It is notable that the vanilla prototype achieves comparable performance with minor sacrifice, indicating that the class centers of the pseudo classifier may not be as real as those of the trained classifier. However, with our structural prototype refinement method, *UniBCT* well surpasses the original BCT, which further demonstrates the effectiveness of our method.

## 5 Conclusion

We for the first time introduce the task of universal backward-compatible representation learning, which covers both close-set and open-set compatible training scenarios for real-world model upgrades. To tackle the challenge of noisy old prototype features, we propose a simple yet effective method, namely *UniBCT*, to properly refine the prototypes by propagating and aggregating their neighbors’ knowledge. *UniBCT* trains the new models to encode discriminative and compatible representations in five different benchmarks in a unified manner. It is the first step towards universal compatible feature learning, and there’s still a long way to go for totally solving this problem. Further studies are called for.

## Acknowledgements

This work was supported by NSFC project Grant No. U1833101, SZSTI Grant No. JCYJ20190809172201639, ZDSYS20210623092001004 and WDZC20200820200655001, the Joint Research Center of Tencent and Tsinghua.

## References

- [Budnik and Avrithis, 2020] Mateusz Budnik and Yannis Avrithis. Asymmetric metric learning for knowledge transfer. *arXiv preprint arXiv:2006.16331*, 2020.
- [Deng *et al.*, 2019] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019.
- [Ge *et al.*, 2020] Yixiao Ge, Haibo Wang, Feng Zhu, Rui Zhao, and Hongsheng Li. Self-supervising fine-grained region similarities for large-scale image localization. In *ECCV*, 2020.
- [Ghifary *et al.*, 2016] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, David Balduzzi, and Wen Li. Deep reconstruction-classification networks for unsupervised domain adaptation. In *ECCV*. Springer, 2016.
- [Guo *et al.*, 2016] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *ECCV*, 2016.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [Li and Hoiem, 2017] Zhizhong Li and Derek Hoiem. Learning without forgetting. *PAMI*, 2017.
- [Liu *et al.*, 2017] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphreface: Deep hypersphere embedding for face recognition. In *CVPR*, 2017.
- [Maze *et al.*, 2018] Brianna Maze, Jocelyn Adams, James A Duncan, Nathan Kalka, Tim Miller, Charles Otto, Anil K Jain, W Tyler Niggel, Janet Anderson, Jordan Cheney, et al. Iarpa janus benchmark-c: Face dataset and protocol. In *ICB*, 2018.
- [Meng *et al.*, 2021] Qiang Meng, Chixiang Zhang, Xiaoqiang Xu, and Feng Zhou. Learning compatible embeddings. *ICCV*, 2021.
- [Philbin *et al.*, 2007] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *2007 IEEE conference on computer vision and pattern recognition*. IEEE, 2007.
- [Philbin *et al.*, 2008] James Philbin, Ondrej Chum, Michael Isard, Josef Sivic, and Andrew Zisserman. Lost in quantization: Improving particular object retrieval in large scale image databases. In *2008 IEEE conference on computer vision and pattern recognition*. IEEE, 2008.
- [Radenović *et al.*, 2018] Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *CVPR*, 2018.
- [Saito *et al.*, 2020] Kuniaki Saito, Donghyun Kim, Stan Sclaroff, and Kate Saenko. Universal domain adaptation through self supervision. In *NeurIPS*, 2020.
- [Shen *et al.*, 2020] Yantao Shen, Yuanjun Xiong, Wei Xia, and Stefano Soatto. Towards backward-compatible representation learning. In *CVPR*, 2020.
- [Su *et al.*, 2022] Shupeng Su, Binjie Zhang, Yixiao Ge, Xuyuan Xu, Yexin Wang, Chun Yuan, and Ying Shan. Privacy-preserving model upgrades with bidirectional compatible training in image retrieval. *arXiv preprint arXiv:2204.13919*, 2022.
- [Wang *et al.*, 2017] Feng Wang, Xiang Xiang, Jian Cheng, and Alan Loddon Yuille. Normface: 12 hypersphere embedding for face verification. In *ACMMM*. ACM, 2017.
- [Wang *et al.*, 2018] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *CVPR*, 2018.
- [Wang *et al.*, 2020] Chien-Yi Wang, Ya-Liang Chang, Shang-Ta Yang, Dong Chen, and Shang-Hong Lai. Unified representation learning for cross model compatibility. *arXiv preprint arXiv:2008.04821*, 2020.
- [Weyand *et al.*, 2020] Tobias Weyand, Andre Araujo, Bingyi Cao, and Jack Sim. Google landmarks dataset v2-a large-scale benchmark for instance-level recognition and retrieval. In *CVPR*, 2020.
- [Yang *et al.*, 2017] Jiaolong Yang, Peiran Ren, Dongqing Zhang, Dong Chen, Fang Wen, Hongdong Li, and Gang Hua. Neural aggregation network for video face recognition. In *CVPR*, 2017.
- [You *et al.*, 2019] Kaichao You, Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Universal domain adaptation. In *CVPR*, 2019.
- [Zhang *et al.*, 2020] Xiao Zhang, Rui Zhao, Yu Qiao, and Hongsheng Li. Rbf-softmax: Learning deep representative prototypes with radial basis function softmax. pages 296–311. Springer, 2020.
- [Zhang *et al.*, 2021] 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 *International Conference on Learning Representations*, 2021.