

Darwinian Model Upgrades: Model Evolving with Selective Compatibility

Binjie Zhang^{1,2,4§}, Shupeng Su^{1§}, Yixiao Ge^{1¶}, Xuyuan Xu³, Yexin Wang³,
Chun Yuan⁴, Mike Zheng Shou², Ying Shan¹

¹ARC Lab, Tencent PCG

²National University of Singapore

³AI Technology Center of Tencent Video

⁴Tsinghua University

Abstract

The traditional model upgrading paradigm for retrieval requires recomputing all gallery embeddings before deploying the new model (dubbed as “backfilling”), which is quite expensive and time-consuming considering billions of instances in industrial applications. BCT presents the first step towards backward-compatible model upgrades to get rid of backfilling. It is workable but leaves the new model in a dilemma between new feature discriminativeness and new-to-old compatibility due to the undifferentiated compatibility constraints. In this work, we propose Darwinian Model Upgrades (DMU), which disentangle the inheritance and variation in the model evolving with selective backward compatibility and forward adaptation, respectively. The old-to-new heritable knowledge is measured by old feature discriminativeness, and the gallery features, especially those of poor quality, are evolved in a lightweight manner to become more adaptive in the new latent space. We demonstrate the superiority of DMU through comprehensive experiments on large-scale landmark retrieval and face recognition benchmarks. DMU effectively alleviates new-to-new degradation and improves new-to-old compatibility, rendering a more proper model upgrading paradigm in large-scale retrieval systems. Code: <https://github.com/TencentARC/OpenCompatible>.

Introduction

The paradigm of backward-compatible model upgrades was first introduced in (Shen et al. 2020), and attracted increasing attention in the community (Zhang et al. 2022a; Meng et al. 2021; Hu et al. 2022; Zhang et al. 2022b; Su et al. 2022). Such a paradigm enables the backfill-free model upgrades*, that is, the backward-compatible new models can be deployed on system immediately after training. It is especially beneficial for the large-scale retrieval systems in which the conventional offline backfilling may take weeks or even months for a model upgrade.

The essence of backward compatibility is to make old gallery embeddings and new query embeddings interoperable in the same latent space. Despite various efforts (Zhang

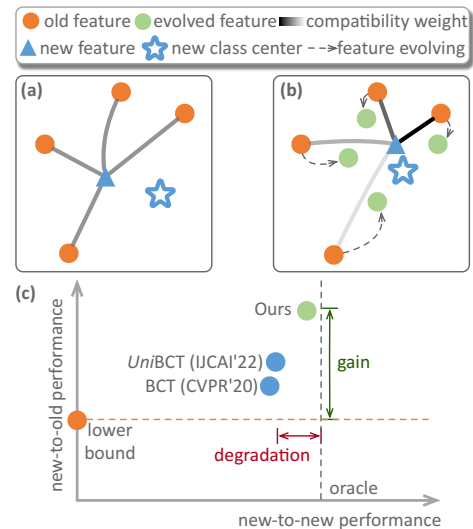


Figure 1: (a) Previous backward-compatible model upgrades: the dilemma between new model capability and backward compatibility caused by undifferentiated inheritance. (b) Darwinian model upgrades: get out of the dilemma with selective backward compatibility/inheritance and forward adaptation/evolution. (c) Experimental verification: our DMU well alleviates the new-to-new retrieval degradation while boosting new-to-old compatibility on Google Landmark v2 dataset.

et al. 2022b; Budnik and Avrithis 2020) to improve cross-model compatibility, a dilemma has yet to be solved that the new models have to sacrifice their internal capabilities (new-to-new) to maintain the compatibility with the older inferior models (new-to-old). Note that the new-to-new accuracy is also important for backward-compatible model upgrades as it affects the feature discriminativeness of the query and new gallery instances. We argue that the dilemma is caused by the *undifferentiated inheritance* in existing compatible methods for retrieval, where new embedding models are required to inherit both beneficial and harmful historical knowledge[†].

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[§]Equal contribution. [¶]Corresponding author.

*The term “backfill” indicates re-extracting all the gallery embeddings with the new model before its deployment.

[†]In embedding applications (e.g., image retrieval), model knowledge is generally assessed by feature discriminativeness.

To this end, we present **Darwinian Model Upgrades[‡] (DMU)**, a new paradigm in which (i) the more discriminative old features that favor retrieval will tend to be inherited and (ii) the inferior old features will be adapted towards the new latent space. The former property well alleviates new-to-new self-discrimination degradation which also improves new-to-old retrieval with better new query features. The latter property further enhances new-to-old compatibility by evolving old gallery features. DMU can properly relieve the dilemma between new model capability and compatibility.

Specifically, DMU is reached by the selective backward compatibility (inheritance) and the forward adaptation (variation). The backward-compatible regularization is re-weighted by the old feature discriminativeness, *i.e.*, the probabilities of correctly identifying certain old features using the old classifier act as coefficients for the new-to-old compatible loss. To further upgrade the old features especially those under quality, a forward-adaptive module is trained to transform old gallery embeddings in a feature-to-feature fashion. The forward-adaptive module is lightweight and requires marginal computational overhead.

We demonstrate the superiority of DMU against the other model upgrading methods on multiple widely-acknowledged image retrieval benchmarks, including the large-scale landmark retrieval and face recognition. Following state-of-the-art methods (Jeon 2020; Mei et al. 2020), we perform model upgrades on GLDv2-train (Weyand et al. 2020) and MS1Mv3 (Deng et al. 2019), and evaluate retrieval on GLDv2-test, \mathcal{R} Oxford (Radenović et al. 2018b), \mathcal{R} Paris, and IJB-C (Radenović et al. 2018a) respectively. DMU consistently surpasses the competitors in terms of both new-to-new and new-to-old retrieval performances.

Overall, we tackle the dilemma in existing backward-compatible model upgrades that new models have to sacrifice their feature discriminativeness to preserve the compatibility with old models by introducing Darwinian Model Upgrading paradigm with selective backward compatibility and forward adaptation. We extensively verify the effectiveness of our method in terms of both new-to-new and new-to-old retrieval on multiple public benchmarks. We hope our work would inspire the community to think out of the box and discover more proper model upgrading methods.

Related Work

Backward-compatible Model Upgrades. BCT (Shen et al. 2020) is the first work to introduce backward-compatible model upgrades paradigm by adding an influence loss when training the new model. It enforces the new features to approach the corresponding old class centroids, rendering the comparability between the new and old embedding spaces and bypassing the gallery backfilling. RACT (Zhang et al. 2022a) takes a further step by introducing the hot-refresh backward-compatible model upgrades that backfills the gallery on-the-fly after the deployment of the backward-compatible new model. Under this BCT framework, several work (Budnik and Avrithis 2020; Meng

et al. 2021; Zhang et al. 2022b) makes efforts on designing the different training objectives to improve the backward-compatible retrieval performance. Specifically, AML (Budnik and Avrithis 2020) studies the adaptation of the common metric learning into asymmetric retrieval constraints when training the new model towards the old one, and finally introduces a better variant combining the contrastive and regression losses. LCE (Meng et al. 2021) proposes an alignment loss to align the new classifier with the old one and a tighter boundary loss for the new features towards the old class centers. UniBCT (Zhang et al. 2022b) introduces a universal backward-compatible loss which refines the noisy old features with graph transition at first and then performs the backward-compatible training. Although making progress, the previous literature has to sacrifice the discriminativeness of the new model heavily as they put the same weight on the backward-compatible constraint towards each old embedding. It is sensitive to the outliers and accordingly we propose an adaptive re-weight scheme to alleviate this.

Cross-model Compatible Transformation. Compatible feature transformation (Chen et al. 2019; Wang et al. 2020; Ramanujan et al. 2022) also targets the compatible retrieval among different embedding models. However, it concerns training transformation modules to map the features from different models to a common space. Specifically, (Chen et al. 2019) first introduces the embedding transformation problem in the face recognition occasion and tackle it with a R^3AN framework which transforms the source embeddings into the target embeddings as a regression task. CMC (Wang et al. 2020) generalizes the problem in unlimited retrieval occasions and proposes the RBT (Residual Bottleneck Transformation) module with a training scheme combining classification loss, regression loss and KL-divergence loss. FCT (Ramanujan et al. 2022) further introduces the side information feature for each sample which facilitates the upgrade of old embeddings towards the unknown future new embeddings. Although a feasible model upgrade paradigm, the mere cross-model transformation will lose the merits of the immediate deployment of new embedding model and upgrading the gallery features on-the-fly when compared to the backward-compatible paradigm above. Moreover, training transformation alone hardly achieves optimal performance as the feature upgrade module is generally lightweight and difficult to update the old features to a completely incompatible new space.

Focal Distillation. PCT (Yan et al. 2020) proposes a focal distillation loss to relieve the “negative flips” problem when training the new model which gives more distillation loss weights on the samples that are classified correctly by the old model. Considering both are re-weight scheme, our selective backward-compatible loss is different in: (1) PCT targets the classification task while our DMU focuses on the retrieval and constrains the feature compatibility instead of the classification probability similarity; (2) our re-weight scheme is adaptive to the discriminativeness of old features while the PCT one is only an indication function of whether correctly classified by the old model. We will demonstrate that DMU surpasses PCT distinctly in compatible retrieval.

[‡]The name draws from an analogy to “natural selection” in Darwinism, but for model evolving.

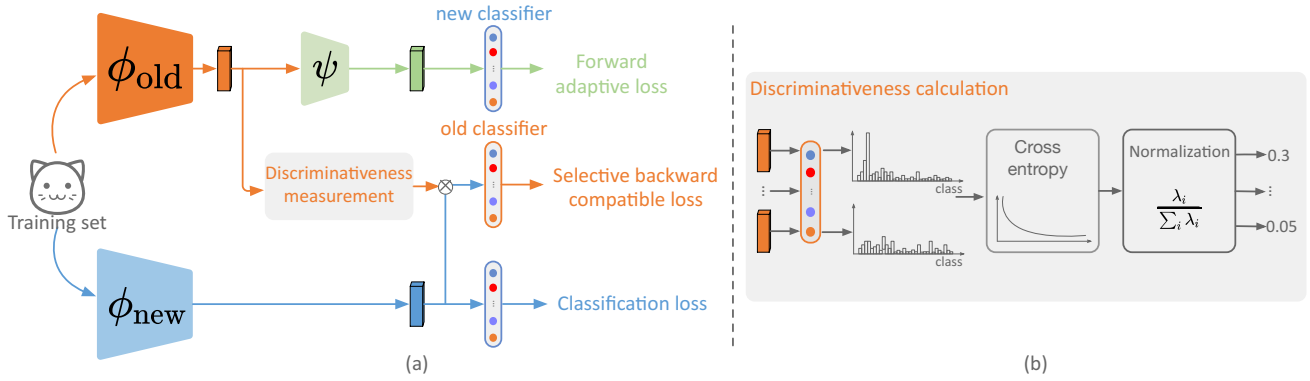


Figure 2: (a) The proposed Darwinian Model Upgrades framework which performs inheritance with selective compatibility and variation with forward adaptation. (b) The details of discriminativeness measurement.

Methodology

We study the task of efficient and effective model upgrades in large-scale embedding applications and take the widely-used image retrieval as an example, introduced in Sec. . (Shen et al. 2020) breaks the traditional way of model upgrades and comes up with a new paradigm, backward-compatible model upgrades, that enables immediate deployment of the new model without the time-consuming backfilling. Despite a feasible solution, it suffers from the undifferentiated inheritance, discussed in Sec. . Consequently, we introduce a new paradigm, namely, Darwinian Model Upgrades, with selective backward compatibility and forward adaptation (see Sec.). We elaborate each part as follows.

Image Retrieval

Let the gallery be denoted as \mathcal{G} and query as \mathcal{Q} . Given an embedding model $\phi : x \mapsto \mathbb{R}^D$, the task of image retrieval is to identify the object/content of interest in gallery resorting to the query-gallery similarities, *i.e.*, $\|\phi(\mathcal{Q}) - \phi(\mathcal{G})\|$. Following state-of-the-art metric learning methods (Jeon 2020; Mei et al. 2020) in image retrieval, we employ classification as a pretext task in the form of ArcFace loss (Deng et al. 2019). Given an image x of label y , the loss is formulated as

$$\ell_{\text{arc}}(x; \phi, \omega) = -\log \frac{e^{s \cdot k(\phi(x), \omega(y), m)}}{e^{s \cdot k(\phi(x), \omega(y), m)} + \sum_{j \neq y} e^{s \cdot k(\phi(x), \omega(j), 0)}}, \quad (1)$$

where m is the margin, s is a scalar, and ω is a classifier with fully-connected layers. The kernel function is defined as $k(\phi(x), \omega(j), m) = \cos(\arccos(\langle \phi(x), \omega(j) \rangle) + m)$, where $\langle \cdot, \cdot \rangle$ represents the inner product.

Backfill-free Model Upgrades: A Revisit

Backward Compatible Training (BCT). BCT was first introduced by (Shen et al. 2020) to make the new and old features interchangeable in the same latent space. The training objectives are formulated as,

$$\arg \min_{\phi_n} (\mathcal{L}_{\text{new}} + \mathcal{L}_{\text{BC}}), \quad (2)$$

where ϕ_n is the new embedding model, \mathcal{L}_{new} is the self-discriminative loss for the new model, and \mathcal{L}_{BC} is the backward-compatible constraint.

In a mini-batch \mathcal{B} , the self-discriminative loss for the new model is formulated as an ArcFace loss, such as

$$\mathcal{L}_{\text{new}} = \frac{1}{|\mathcal{B}|} \sum_{x \in \mathcal{B}} \ell_{\text{arc}}(x; \phi_n, \omega_n), \quad (3)$$

where ω_n indicates the new classifier. To make new features and old ones directly comparable with each other, backward compatibility is achieved by enforcing the new features to approach their corresponding old class centers,

$$\mathcal{L}_{\text{BC}} = \frac{1}{|\mathcal{B}|} \sum_{x \in \mathcal{B}} \ell_{\text{arc}}(x; \phi_n, \omega_o), \quad (4)$$

where ω_o is the old and fixed classifier.

Backfill-free Model Upgrades. Given the interoperable old and new features, backfill-free model upgrades can be achieved where the new embedding model can be immediately deployed without re-extracting the gallery features. The retrieval accuracy of the new system is improved by more discriminative query features, *i.e.*,

$$\mathcal{M}(\phi_o(\mathcal{Q}), \phi_o(\mathcal{G})) < \mathcal{M}(\phi_n(\mathcal{Q}), \phi_o(\mathcal{G})), \quad (5)$$

where $\mathcal{M}(\cdot, \cdot)$ denotes the evaluation metric for retrieval (*e.g.*, mAP used in landmark retrieval or TAR@FAR used in face recognition). For simplicity, we will omit \mathcal{Q}, \mathcal{G} and abbreviate $\mathcal{M}(\phi_o(\mathcal{Q}), \phi_o(\mathcal{G}))$ as $\mathcal{M}(\phi_o, \phi_o)$. We denote the system-level model upgrading gain (Δ_{\uparrow}) as

$$\Delta_{\uparrow} = \frac{\mathcal{M}(\phi_n, \phi_o) - \mathcal{M}(\phi_o, \phi_o)}{\mathcal{M}(\phi_o, \phi_o)}. \quad (6)$$

Limitation. Despite the model upgrading gains achieved by new queries, we observe significant degradation of new-to-new retrieval accuracy compared to the model trained without backward compatibility (denoted as ‘‘oracle’’). It is actually due to the fact that new models have to sacrifice




Figure 3: Illustration of the backfill-free paradigm and Darwinian paradigm in sequential model upgrades.

self-discriminateness to maintain compatibility with the older and inferior models under the current *undifferentiated inheritance* constraints, rendering BCT a sub-optimal solution for model upgrades. We evaluate the degree of such discriminateness degradation (Δ_{\downarrow}) with the following metric,

$$\Delta_{\downarrow} = \frac{\mathcal{M}(\phi_n^{\text{oracle}}, \phi_n^{\text{oracle}}) - \mathcal{M}(\phi_n, \phi_n)}{\mathcal{M}(\phi_n^{\text{oracle}}, \phi_n^{\text{oracle}})}. \quad (7)$$

Darwinian Model Upgrades (DMU)

Overview. Our DMU consists of a backward-compatible embedding model ϕ_n and a lightweight forward-adaptive module ψ (e.g., MLPs), as illustrated in Fig. 2. To resolve the dilemma between upgrading gain (Δ_{\uparrow}) and discriminateness degradation (Δ_{\downarrow}), DMU inherits the discriminative old knowledge with selective backward-compatible training and in the meanwhile evolves old embeddings towards new space with forward-adaptive training. The overall training objectives of DMU can be formulated as,

$$\arg \min_{\phi_n, \psi} (\mathcal{L}_{\text{new}} + \mathcal{L}_{\text{SBC}} + \mathcal{L}_{\text{FA}}), \quad (8)$$

where \mathcal{L}_{SBC} is the selective backward-compatible loss, \mathcal{L}_{FA} is the forward-adaptive loss, and \mathcal{L}_{new} is the same as Eq. 3.

Inheritance with Selective Backward Compatibility.

The key factor causing the discriminateness degradation lies in the undifferentiated compatibility constraint where the new model is required to inherit both beneficial and harmful historical knowledge. We tackle the challenge with re-weighted backward-compatible objectives based on the old feature discriminateness. We measure the feature discriminateness with uncertainty metric from the information entropy theory,

$$\Lambda(x; \phi, \omega) = \sum_{i \in |\mathcal{C}|} -p_i(x) \log p_i(x), \quad \text{where} \quad (9)$$

$$p_i(x) = \frac{e^{\langle \phi(x), \omega(i) \rangle}}{\sum_j e^{\langle \phi(x), \omega(j) \rangle}}, \quad \left(\sum_{i \in |\mathcal{C}|} p_i(x) = 1 \right),$$

where $|\mathcal{C}|$ is the number of classes and ω is a classifier. $\Lambda(x; \phi, \omega)$ is inversely proportional to discriminateness.

For alleviating the negative effects of poor historical knowledge, we should punish the inheritance of samples with poor discriminateness while stressing the good ones.




Figure 4: Variation of entropy-based discriminateness. We randomly select 1k gallery images from GLDv2-test and extract features with ϕ_{old}, ψ . The horizontal axes is defined as $\Lambda_{\text{old}} = \Lambda(\mathcal{G}; \phi_{\text{old}}, \omega_{\text{old}})$, and Λ_{FA} in the vertical axes as $\Lambda_{\text{FA}} = \Lambda(\mathcal{G}; \psi(\phi_{\text{old}}), \omega_{\text{new}})$

Specifically, we use the old classifier ω_o to evaluate the discriminateness of old features, with which we formulate our selective backward-compatible loss as follows,

$$\mathcal{L}_{\text{SBC}} = \sum_{x \in \mathcal{B}} \lambda(x) \ell_{\text{arc}}(x; \phi_n, \omega_o), \quad \text{where} \quad (10)$$

$$\lambda(x) = \frac{1 - \text{softmax}[\Lambda(x; \phi_o, \omega_o)]}{|\mathcal{B}| - 1}, \quad \left(\sum_{x \in \mathcal{B}} \lambda(x) = 1 \right),$$

where the softmax function is operated on a batch samples $(x_1, \dots, x_{|\mathcal{B}|})$. The old features with better discriminateness will be allocated with larger backward-compatible weights.

Variation with Forward Adaptation. The feature forward-adaptation module ψ can be a simple Multi-Layer Perceptron (MLP), i.e., stack of fully-connected and non-linear activation layers (the detailed structure is listed in Experiments section). ψ upgrades the old features towards the new space with the following objective,

$$\mathcal{L}_{\text{FA}} = \frac{1}{|\mathcal{B}|} \sum_{x \in \mathcal{B}} \ell_{\text{arc}}(x; \psi(\phi_o), \omega_n). \quad (11)$$

We do not specifically design a selective module for the forward adaptation, however, since the new model already inherits good old knowledge, it automatically selects poorer features for evolution. As a demonstration, we exhibit in Fig. 4 that the old features with poor discriminateness receive a large adaption degree than the good ones which is exactly attributed to our proposed selective backward-compatible constraints.

Darwinian Model Upgrades. As depicted in Fig. 3, our proposed new paradigm, Darwinian Model Upgrades, improves the query embedding quality by the selective backward-compatible new model in the meanwhile improves the gallery embeddings by the feature forward-adaptation head. We formulate the objective of DMU formally as follows,

$$\mathcal{M}(\phi_o, \phi_o) < \mathcal{M}(\phi_n, \phi_o) < \mathcal{M}(\phi_n, \psi \circ \phi_o), \quad (12)$$

Scenario	Method	GLDv2-test				\mathcal{R} Oxford				\mathcal{R} Paris			
		$\mathcal{M}(\phi_n, \phi_n)$	$\mathcal{M}(\phi_n, \phi_o)$	Δ_{\uparrow}	Δ_{\downarrow}	$\mathcal{M}(\phi_n, \phi_n)$	$\mathcal{M}(\phi_n, \phi_o)$	Δ_{\uparrow}	Δ_{\downarrow}	$\mathcal{M}(\phi_n, \phi_n)$	$\mathcal{M}(\phi_n, \phi_o)$	Δ_{\uparrow}	Δ_{\downarrow}
30%data →	ϕ_o^{r50}	18.93	-	-	-	64.33	-	-	-	82.94	-	-	-
	ϕ_n^{r50} (oracle)	25.58	-	-	-	75.88	-	-	-	86.09	-	-	-
100%data	ϕ_n^{r50} (BCT)	24.55	21.12	11.57	4.03	73.74	64.75	0.65	2.82	85.73	83.53	0.71	0.42
	ϕ_n^{r50} (ours)	25.00	22.56	19.18	2.27	74.07	66.68	3.65	2.39	86.59	84.48	1.86	-0.58
30%data →	ϕ_o^{r50}	18.93	-	-	-	64.33	-	-	-	82.94	-	-	-
	ϕ_n^{r50}	23.66	-	-	-	72.20	-	-	-	86.64	-	-	-
70%data	ϕ_n^{r50} (BCT)	23.47	20.74	9.56	0.80	72.91	63.73	-0.93	-0.98	86.44	83.4	0.55	0.23
	ϕ_n^{r50} (ours)	23.76	21.70	14.63	-0.42	73.01	64.14	-0.30	-1.12	85.90	84.64	2.05	0.85
30%class →	ϕ_o^{r50}	17.53	-	-	-	70.07	-	-	-	83.62	-	-	-
	ϕ_n^{r50}	25.58	-	-	-	75.88	-	-	-	86.09	-	-	-
100%class	ϕ_n^{r50} (BCT)	25.10	24.39	39.13	1.88	73.99	73.34	4.67	2.49	85.64	83.97	0.42	0.52
	ϕ_n^{r50} (ours)	25.45	25.01	42.67	0.51	74.13	73.78	5.29	2.31	86.64	84.52	1.08	-0.64
resnet50 →	ϕ_o^{r50}	18.93	-	-	-	64.33	-	-	-	82.94	-	-	-
	ϕ_n^{r101}	27.42	-	-	-	75.89	-	-	-	86.62	-	-	-
resnet101	ϕ_n^{r101} (BCT)	25.56	21.68	14.53	6.78	75.7	66.12	2.78	0.25	87.35	85.59	3.20	-0.84
	ϕ_n^{r101} (ours)	25.92	22.38	18.23	5.47	76.49	66.35	3.14	-0.79	88.08	85.93	3.61	-1.69

Table 1: Comparison of baselines (BCT) and our method (DMU) in various compatible scenarios on Landmark Retrieval datasets (*i.e.*, GLDv2-test, \mathcal{R} Oxford and \mathcal{R} Paris). All models are trained on GLDv2-train dataset. The evaluation protocols are mAP ($\mathcal{M}(\cdot, \cdot)$, larger is better), upgrade gain (Δ_{\uparrow} , larger is better) and discriminativeness degradation (Δ_{\downarrow} , smaller is better). The architectures are ResNet50 (r50) and ResNet101 (r101). A negative value for degradation means that the degradation has been fully resolved, and the new model even benefits from backward-compatible regularization.

where ϕ_n, ψ are the selective backward-compatible new model and feature forward-adaptation module. As will be shown later, performing DMU can bring consistent improvements for the retrieval performance and especially benefits sequential upgrades as it relieves the shackle of the new model for keeping compatibility to the oldest gallery.

Experiments

In this section, we will validate the proposed Darwinian Model Upgrades on the widely-acknowledged image retrieval datasets including the Google Landmark v2 (Weyand et al. 2020), \mathcal{R} Oxford (Radenović et al. 2018b), \mathcal{R} Paris (Radenović et al. 2018b), MS1Mv3 (Deng et al. 2019) and IJB-C (Maze et al. 2018).

Experiment Settings

Datasets. (1) **Landmark Retrieval:** We adopt the GLDv2-train-clean version (Weyand et al. 2020) as training set which is comprised of 1,580,470 images in 81,313 landmarks. After the training, the retrieval performance is evaluated on GLDv2-test (Weyand et al. 2020), Revisited Oxford (\mathcal{R} Oxford) (Radenović et al. 2018b), and Revisited Paris (\mathcal{R} Paris) (Radenović et al. 2018b). For GLDv2-test, there are 750 query images and 761,757 gallery images. As to \mathcal{R} Oxford, it contains 70 query images and 4,993 gallery images. \mathcal{R} Paris is comprised of 70 query images and

6,322 gallery images. (2) **Face Recognition:** All models are trained on MS1Mv3 (Deng et al. 2019) dataset, which contains 5,179,510 training images with 93,431 labels. To evaluate the model generalization, we conduct face verification on IJB-C test set, which provides 469,376 template pairs.

Metric. (1) **Landmark Retrieval:** We adopt the cosine similarity for ranking and the retrieval results are measured with the mean Average Precision (mAP). Specifically, following the test protocols proposed in (Weyand et al. 2020; Radenović et al. 2018b), we use mAP@100 for GLDv2-test, and mAP@all in the **Medium** evaluation setup (*Easy* and *Hard* images are treated as positive, while *Unclear* are ignored) for \mathcal{R} Oxford and \mathcal{R} Paris. (2) **Face Recognition:** We calculate the true acceptance rate (TAR) at different false acceptance rates (FAR) for template pairs. In cross-model retrieval, we extract the first template with the new model, and the second with the old model.

Compatible Scenarios. To simulate the different model upgrade occasions, we randomly split the training set with different percentages (30%, 70% and 100%). For extended-data (30%data → 100%data) scenario, the old training set and the new one share all landmarks; for extended-class (30%class → 100%class) scenario, the old training set only covers part of all labels. Besides, we verify our algorithm in different model architectures (resnet50→resnet101). Details about data allocation are in supplemental materials.

Scenario	Method	IJB-C verification			
		$\mathcal{M}(\phi_n, \phi_n)$	$\mathcal{M}(\phi_n, \phi_o)$	Δ_{\uparrow}	Δ_{\downarrow}
30%data	ϕ_o^{ir18}	91.67	-	-	-
	ϕ_n^{ir18} (oracle)	93.89	-	-	-
100%data	ϕ_n^{ir18} (BCT)	93.96	92.82	1.25	-0.07
	ϕ_n^{r18} (ours)	94.17	92.96	1.41	-0.30
30%data	ϕ_o^{ir18}	91.67	-	-	-
	ϕ_n^{ir18} (oracle)	93.79	-	-	-
70%data	ϕ_n^{ir18} (BCT)	93.69	92.61	1.03	0.11
	ϕ_n^{ir18} (ours)	93.96	92.77	1.20	-0.18

Table 2: Comparison of baseline (BCT) and our method (DMU) on Face Recognition dataset (IJB-C). All models are trained on MS1M-v3 dataset. The evaluation protocol is TAR@FAR= $1e^{-4}$ ($\mathcal{M}(\cdot, \cdot)$, larger is better), upgrade gain (Δ_{\uparrow} , larger is better) and discriminativeness degradation (Δ_{\downarrow} , smaller is better). The architecture is iResNet18 (ir18).

Λ type	GLDv2-test			
	$\mathcal{M}(\phi_n, \phi_n)$	$\mathcal{M}(\phi_n, \phi_o)$	Δ_{\uparrow}	Δ_{\downarrow}
uniform (BCT)	24.55	21.12	11.57	4.03
least conf.	24.88	22.25	17.54	2.74
entropy	25.00	22.56	19.18	2.27

Table 3: Comparison of different discriminativeness measurement in 30%data \rightarrow 100%data on GLDv2-test. The backbone is ResNet50.

Training Details. Most of our basic settings follows the experiments in (Ozaki and Yokoo 2019; Henkel and Singer 2020). Specifically, we adopt the ResNet (resnet50 by default) as the backbone of embedding model and substitute the global average pooling layer by Generalized-Mean (GeM) pooling (Radenović, Tolias, and Chum 2018) with hyper-parameter $p=3$. A fully-connected (fc) layer is appended in the end which transforms the pooling features into the output embedding (the dimension is set as 1024). For the feature upgrade module ψ , we adopt [fc-bn-relu] as basic block and stack 3 of this with a output fc layer as the final MLP construction. The input image is resized to 224×224 for training and inference. Random image augmentation is applied which includes the random resized cropping and horizontal flipping. With 6 Tesla V100 for training, the batch size per GPU is set as 256 for the embedding model and 512 for the feature upgrade module. SGD optimizer with 0.9 momentum and 10^{-4} weight decay is adopted. Besides, we uniformly use the cosine lr scheduler with 1 warm-up epoch in the total running of 30 epochs. The initial learning rate is set as 0.1. For the training objective, we set hyper-parameters in ArcFace loss as $s = 30$, $m = 0.3$ in landmark retrieval, $s = 64$, $m = 0.5$ in face recognition.

Method	\mathcal{L}_{SBC}	\mathcal{L}_{FA}	GLDv2-test			
			$\mathcal{M}(\phi_n, \phi_n)$	$\mathcal{M}(\phi_n, \phi_o)$	Δ_{\uparrow}	Δ_{\downarrow}
ϕ_n^{r50} (BCT)	-	-	24.55	21.12	11.57	4.03
ϕ_n^{r50} (ours)	✓	-	24.75	21.60	14.10	3.24
ϕ_n^{r50} (ours)	-	✓	24.72	22.42	18.44	3.36
ϕ_n^{r50} (ours)	✓	✓	25.00	22.56	19.18	2.27

Table 4: Ablation of our Darwinian Model Upgrades paradigm in 30%data \rightarrow 100%data on GLDv2-test. The backbone is ResNet50.

Dimension	GLDv2-test			
	$\mathcal{M}(\phi_n, \phi_n)$	$\mathcal{M}(\phi_n, \phi_o)$	Δ_{\uparrow}	Δ_{\downarrow}
512	24.78	22.12	16.85	3.13
1024	25.00	22.56	19.18	2.27
2048	24.96	22.77	20.29	2.42

Table 5: Comparison of different dimension in forward-adaptive head in 30%data \rightarrow 100%data on GLDv2-test. The backbone is ResNet50.

Analysis of DMU

Different Model Upgrade Occasions with DMU. As shown in Table 1, our DMU surpasses baselines on various practical model upgrade occasions, which comprehensively demonstrates the effectiveness of our approach.

Effectiveness of Selective Backward Compatibility. To demonstrate the proposed selective backward compatibility can alleviate the issue of discriminativeness degradation, we only train the backbone without the forward-adaptive head. Results in Table 4 show that inheriting good old knowledge will benefit the new model training.

Except from the entropy-based discriminativeness measurement, we verify another approach based on least confidence, defined as the following,

$$\Lambda(x; \phi, \omega) = \sum_i \mathbb{I}(i = y) \cdot p_i(x), \quad (13)$$

$$\lambda(x) = \text{softmax}(\Lambda(x; \phi, \omega)),$$

where $p_i(x)$ is calculated by the same formulation in Eq. 9. As illustrated in Table 3, different forms of discriminativeness measurement alleviate new-to-new degradation and boost new-to-old retrieval performance, which indicates the generalization of our method.

Effectiveness of Forward Adaptation. Results in the third row of Table 4 shows that mere forward-adaptive head can extremely boost the new-to-old retrieval performance compared with baseline. Evolving old features (especially poor features) towards to the new embedding space contributes to the improvement.

Scenario	Method	1st generation				2nd generation			
		$\mathcal{M}(\phi_n, \phi_n)$	$\mathcal{M}(\phi_n, \phi_o)$	Δ_{\uparrow}	Δ_{\downarrow}	$\mathcal{M}(\phi_n, \phi_n)$	$\mathcal{M}(\phi_n, \phi_o)$	Δ_{\uparrow}	Δ_{\downarrow}
Extended data: 25%→50%→75%	$\phi_{\text{new}}^{\text{oracle}}/\phi_{\text{old}}$	22.39/17.89	-	-	-	24.04/17.89	-	-	-
	BCT	21.47	18.88	5.53	4.11	23.52	19.82	10.79	2.16
	ours	22.66	19.84	10.90	-1.21	23.88	21.16	18.28	0.67

Table 6: Sequential model upgrades with DMU on GLDv2-test. The backbone is ResNet50.

Comp. loss	GLDv2-test			
	$\mathcal{M}(\phi_n, \phi_n)$	$\mathcal{M}(\phi_n, \phi_o)$	Δ_{\uparrow}	Δ_{\downarrow}
BCT (ArcFace)	24.55	21.12	11.57	4.03
DMU (ArcFace)	25.00	22.56	19.18	2.27
BCT (Regression)	25.11	19.96	5.44	1.84
DMU (Regression)	25.27	21.23	12.15	1.21

Table 7: Different regularizers in 30%data→100%data on GLDv2-test. The backbone is ResNet50.

Dimension	GLDv2-test			
	$\mathcal{M}(\phi_n, \phi_n)$	$\mathcal{M}(\phi_n, \phi_o)$	Δ_{\uparrow}	Δ_{\downarrow}
BCT(CVPR’20)	24.55	21.12	11.57	4.03
FD(CVPR’21)	25.24	20.02	5.76	1.3
UniBCT(IJCAI’22)	24.60	21.77	15.00	3.83
ours	25.00	22.56	19.18	2.27

Table 8: Comparison with SOTAs in 30%data→100%data on GLDv2-test. The backbone is ResNet50.

We ablate the hidden dimension of the forward-adaptive head ψ and report results in Table 5. Considering the trade-off between the accuracy and the refresh efficiency, we adopt the median hidden dimension 1024 as the default setting in the following experiments unless stated otherwise.

The MACs (Multiply-Accumulate Operations) of a Forward-adaptive head with 1024-dim is about 8M, which is marginal to refresh gallery database. Specifically, refreshing a billion features with FA head with 1024 (running with TensorRT on NVIDIA Tesla T4 GPU) only consumes less than 8 hours. At a negligible cost, user experience can be gradually improved with DMU framework.

Different Compatible Regularizers. To elaborate our proposed selective backward compatibility is independent with compatible regularizers, we apply our DMU framework on another regression version (Budnik and Avrithis 2020) which minimizes the cosine embedding distance between new and old features.

$$\ell_{\text{reg}}(x; \phi_n, \phi_o) = 1 - \cos\langle \phi_n(x), \phi_o(x) \rangle. \quad (14)$$

We replace $\ell_{\text{arc}}(x; \phi_n, \phi_o)$ with $\ell_{\text{reg}}(x; \phi_n, \phi_o)$ in Eq. 10 and 11. Table 7 displays the comparison using different types of compatible regularizers. We observe that our method can also alleviate the new-to-new degradation, which indicates dynamically inheriting discriminative old features is effective. To summarize, our proposed paradigm, Darwinian Model Upgrades, is general and effective.

Sequential Model Upgrades. To further explore the advantage of DMU, we conduct sequential model upgrades experiments and the results are shown in Table 6. As can be seen, our proposed DMU paradigm consistently surpasses the previous BCT paradigm in both new-to-old and new-to-new retrieval performance. It demonstrates again that our

novel selective backward compatibility loss can facilitate the better query features and the introduced forward-adaptation head can achieve the better gallery features.

Comparison with Other Methods. We compare with different SoTA compatible learning works in Table 8. When compared with the Focal Distillation method, our DMU can achieve comparable performance in new-to-new metric and surpass it largely in new-to-old which exactly demonstrates the superiority of our re-weight scheme than the Focal Distillation one. Besides, when compared with UniBCT, DMU achieves both better new-to-old and new-to-new performances. It demonstrates that our selective backward-compatible loss works better in noisy samples resistance than the graph transition proposed in UniBCT.

Conclusions

Due to the undifferentiated compatibility constraints, existing backward-compatible works are stuck in dilemma between new feature discriminativeness and new-to-old compatibility. To tackle the issue, we introduce a novel paradigm, Darwinian Model Upgrade with selective compatibility, which tends to inherit discriminative old features with selective backward compatible training and evolves old features to become increasingly well-suited to the better latent space with a light-weight forward-adaptive head. With DMU, the query embedding quality is improved by the new backbone in the meanwhile the gallery is quickly refreshed by the forward-adaptive head, rendering a continual cross-model retrieval boost. We demonstrate the effectiveness of DMU in many mainstreaming datasets including Google Landmark, ROxford, RParis, MS1Mv3 and IJB-C. We hope our work can inspire the community to discover proper model upgrading approaches.

Acknowledgments

This work is partially inspired by the previous tech report BiCT (Su et al. 2022).

Mike Zheng Shou is only supported by the National Research Foundation, Singapore under its NRFF Award NRF-NRFF13-2021-0008. This work is also supported by the National Key R&D Program of China (2022YFB4701400/4701402), SZSTC Grant (JCYJ20190809172201639, WDZC20200820200655001), Shenzhen Key Laboratory (ZDSYS20210623092001004), and Joint Research Center of Tencent and Tsinghua.

References

- Budnik, M.; and Avrithis, Y. 2020. Asymmetric metric learning for knowledge transfer. *arXiv preprint arXiv:2006.16331*.
- Chen, K.; Wu, Y.; Qin, H.; Liang, D.; Liu, X.; and Yan, J. 2019. R3 adversarial network for cross model face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9868–9876.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*.
- Henkel, C.; and Singer, P. 2020. Supporting large-scale image recognition with out-of-domain samples. *arXiv preprint arXiv:2010.01650*.
- Hu, W.; Bansal, R.; Cao, K.; Rao, N.; Subbian, K.; and Leskovec, J. 2022. Learning Backward Compatible Embeddings. *arXiv preprint arXiv:2206.03040*.
- Jeon, S. 2020. 1st Place Solution to Google Landmark Retrieval 2020. *arXiv preprint arXiv:2009.05132*.
- Maze, B.; Adams, J.; Duncan, J. A.; Kalka, N.; Miller, T.; Otto, C.; Jain, A. K.; Niggel, W. T.; Anderson, J.; Cheney, J.; et al. 2018. Iarpa janus benchmark-c: Face dataset and protocol. In *ICB*.
- Mei, K.; Xu, J.; Cheng, Y.; Lin, Y.; et al. 2020. 3rd Place Solution to Google Landmark Retrieval 2020. *arXiv preprint arXiv:2008.10480*.
- Meng, Q.; Zhang, C.; Xu, X.; and Zhou, F. 2021. Learning Compatible Embeddings. In *ICCV*.
- Ozaki, K.; and Yokoo, S. 2019. Large-scale landmark retrieval/recognition under a noisy and diverse dataset. *arXiv preprint arXiv:1906.04087*.
- Radenović, F.; Iscen, A.; Tolias, G.; Avrithis, Y.; and Chum, O. 2018a. Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5706–5715.
- Radenović, F.; Iscen, A.; Tolias, G.; and Chum, O. 2018b. Revisiting oxford and paris: Large-scale image retrieval benchmarking. In *CVPR*.
- Radenović, F.; Tolias, G.; and Chum, O. 2018. Fine-tuning CNN image retrieval with no human annotation. *IEEE TPAMI*.
- Ramanujan, V.; Vasu, P. K. A.; Farhadi, A.; Tuzel, O.; and Pouransari, H. 2022. Forward Compatible Training for Large-Scale Embedding Retrieval Systems. In *CVPR*, 19386–19395.
- Shen, Y.; Xiong, Y.; Xia, W.; and Soatto, S. 2020. Towards Backward-Compatible Representation Learning. In *CVPR*.
- Su, S.; Zhang, B.; Ge, Y.; Xu, X.; Wang, Y.; Yuan, C.; and Shan, Y. 2022. Privacy-Preserving Model Upgrades with Bidirectional Compatible Training in Image Retrieval. *arXiv preprint arXiv:2204.13919*.
- Wang, C.-Y.; Chang, Y.-L.; Yang, S.-T.; Chen, D.; and Lai, S.-H. 2020. Unified Representation Learning for Cross Model Compatibility. *BMVC*.
- Weyand, T.; Araujo, A.; Cao, B.; and Sim, J. 2020. Google landmarks dataset v2-a large-scale benchmark for instance-level recognition and retrieval. In *CVPR*.
- Yan, S.; Xiong, Y.; Kundu, K.; Yang, S.; Deng, S.; Wang, M.; Xia, W.; and Soatto, S. 2020. Positive-Congruent Training: Towards Regression-Free Model Updates. *arXiv preprint arXiv:2011.09161*.
- Zhang, B.; Ge, Y.; Shen, Y.; Li, Y.; Yuan, C.; Xu, X.; Wang, Y.; and Shan, Y. 2022a. Hot-Refresh Model Upgrades with Regression-Alleviating Compatible Training in Image Retrieval. In *ICLR*.
- Zhang, B.; Ge, Y.; Shen, Y.; Su, S.; Yuan, C.; Xu, X.; Wang, Y.; and Shan, Y. 2022b. Towards Universal Backward-Compatible Representation Learning. *arXiv preprint arXiv:2203.01583*.