

CoReS: Compatible Representations via Stationarity

Niccolò Biondi, Federico Pernici, Matteo Bruni, and Alberto Del Bimbo, *Senior Member, IEEE*

Abstract—Compatible features enable the direct comparison of old and new learned features allowing to use them interchangeably over time. In visual search systems, this eliminates the need to extract new features from the gallery-set when the representation model is upgraded with novel data. This has a big value in real applications as re-indexing the gallery-set can be computationally expensive when the gallery-set is large, or even infeasible due to privacy or other concerns of the application. In this paper, we propose CoReS, a new training procedure to learn representations that are *compatible* with those previously learned, grounding on the stationarity of the features as provided by fixed classifiers based on polytopes. With this solution, classes are maximally separated in the representation space and maintain their spatial configuration stationary as new classes are added, so that there is no need to learn any mappings between representations nor to impose pairwise training with the previously learned model. We demonstrate that our training procedure largely outperforms the current state of the art and is particularly effective in the case of multiple upgrades of the training-set, which is the typical case in real applications.

Index Terms—Deep Convolutional Neural Network, Representation Learning, Compatible Learning, Fixed Classifiers.



1 INTRODUCTION

NATURAL intelligent systems learn from visual experience and seamlessly exploit such learned knowledge to identify similar entities. Modern artificial intelligence systems, on their turn, typically require distinct phases to perform such visual search. An internal representation is first learned from a set of images (the *training-set*) using Deep Convolutional Neural Network models (DCNNs) [1], [2], [3], [4] and then used to index a large corpus of images (the *gallery-set*). Finally, visual search is obtained by identifying the closest images in the gallery-set to an input *query-set* by comparing their representations. Successful applications of learning feature representations are: face-recognition [5], [6], [7], [8], [9], person re-identification [10], [11], [12], [13], image retrieval [14], [15], [16], and car re-identification [17] among the others.

In the case in which novel data for the training-set and/or more recent or powerful network architectures become available, the representation model may require to be *upgraded* to improve its search capabilities. In this case, not only the query-set but also all the images in the gallery-set should be re-processed by the upgraded model to generate new features and replace the old ones to benefit from such upgrading. The re-processing of the gallery-set is referred to as *re-indexing* (Fig. 1).

For visual search systems with a large gallery-set, such as in surveillance systems, social networks or in autonomous robotics, re-indexing is clearly computationally expensive [18] or has critical deployment, especially when the working system requires multiple upgrades or there are real-time constraints. Re-indexing all the images in the gallery-set can be also infeasible when, due to privacy or ethical concerns,

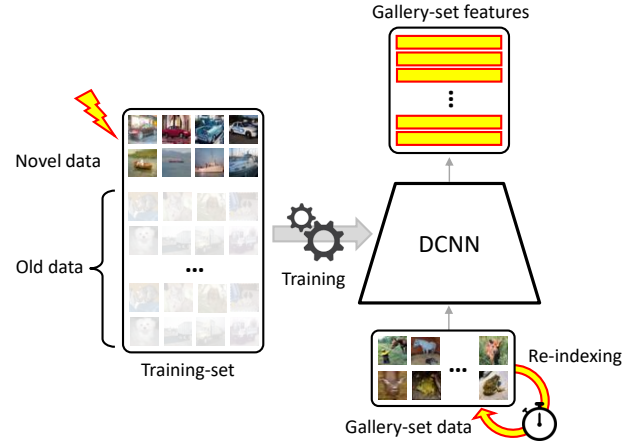


Fig. 1. Upgrading the DCNN representation model with novel data, typically requires the gallery-set to be re-indexed. Learning compatible representations allows to compare the newly learned representation of an input query-set with the old representation of the gallery-set, thus eliminating its computationally intensive re-indexing.

the original gallery images cannot be permanently stored [19] and the only viable solution is to continue using the feature vectors previously computed. In all these cases, it should be possible to directly compare the upgraded features of the query with the previously learned features of the gallery, i.e., the new representation should be *compatible* with the previously learned representation.

Learning compatible representation has recently received increasing attention and novel methods have been proposed in [18], [20], [21], [22], [23], [24]. Differently from these works, in this paper we address *compatibility* leveraging the *stationarity* of the learned internal representation. Stationarity allows to maintain the same distribution of the

The authors are with Media Integration and Communication Center (MICC), Dipartimento di Ingegneria dell'Informazione, Università degli Studi di Firenze, 50139, Firenze, Italy. E-mail: name.surname@unifi.it

features over time so that it is possible to compare the features of the upgraded representation with those previously learned. In particular, we enforce stationarity by leveraging the properties of a family of classifiers whose parameters are not subject to learning, namely *fixed classifiers* based on regular polytopes [25] [26] [27], that allow to reserve regions of the representation space to future classes while classes already learned remain in the same spatial configuration.

The main contributions of our research are the following:

- 1) We identify stationarity as a key property for compatibility and propose a novel training procedure for learning compatible feature representations via stationarity, without the need of learning any mappings between representations nor to impose pairwise training with the previously learned model. We called our method: Compatible Representations via Stationarity (CoReS).
- 2) We introduce new criteria for comparing and evaluating compatible representations in the case of sequential multi-model upgrading.
- 3) We demonstrate through extensive evaluation on large scale verification, re-identification and retrieval benchmarks that CoReS improves the current state-of-the-art in learning compatible features for both single and sequential multi-model upgrading.

In the following, in Sec. 2, we discuss the main literature on compatible representation learning and highlight the distinguishing features of our solution. In Sec. 3, we present in detail the problem of learning compatible representations and define new criteria and metrics for compatibility evaluation in sequential multi-model upgrading. In Sec. 4, we present our solution for learning compatible representations by exploiting feature stationarity. In Sec. 5, we evaluate our solution against state-of-the-art methods on different benchmark datasets and network architectures and demonstrate its superior performance in learning compatible representations. Finally, in Sec. 6, we perform an extensive ablation study.

2 RELATED WORKS

Compatible Representation Learning. The term *backward compatibility* was first introduced in [28] for the classification task. They noted that although machine learning models can increase on average their performance with the availability of more training data, upgrading the model could result into incorrect classification of data correctly classified with the previous model. As a consequence, the trust in machine learning systems is severely harmed. Compatibility in classification has been further investigated in [29], [30], [31], [32].

However, learning compatible representations is substantially different from learning compatible classifier models, although both follow the same general principle. As a distinct hallmark, learning compatible representations directly imposes constraints in the semantic distance of the feature representation. In [21], [22], [23], [24] the problem of feature compatibility was addressed by learning a *mapping* between two representation models so that the new and old feature vectors can be directly compared. The mapping in [21] was learned through a three-step procedure: adversarial

learning for reconstruction, feature extraction and regression to jointly optimize the whole model. In [22], the mapping was learned through an autoencoder by minimizing the distance between the two representation spaces and the reconstruction error. In [23], the mapping was learned from a residual bottleneck transformation module trained by three different losses: classification loss, similarity loss between feature spaces, and KL-divergence loss between the prototypes of the classifiers. In [24], the estimated mapping aligns the class prototypes between the models. To further encourage compatibility, the method also reduces intra-class variations for the new model. All these methods do not completely avoid the cost of re-indexing of the gallery-set as, at each upgrade, the old feature vectors must be re-processed with the learned mapping. Therefore, they are not suited for sequential multi-model learning and large gallery-sets. Differently from these works, we avoid learning specific space-to-space mappings for each previous upgraded representation model and completely avoid the cost of re-indexing also in the case of multiple upgrades.

By avoiding to learn space-to-space mappings, our work has some affinity with the Backward Compatible Training (BCT) [18] method for compatible learning that represents the current state-of-the-art. BCT grounds on pairwise compatibility learning to obtain compatible features. It takes advantage of an influence loss that biases the new representation in a way that it can be used by both the new and the old classifier. During learning with novel data, the old classifier is *fixed* and the prototypes of the new classifier align with the prototypes of the old classifier. In the case of multi-model upgrading, such pairwise cooperation supports compatibility only indirectly (i.e., through transitive compatibility). In fact, for a two-model upgrading, i.e., when the model ϕ_1 is upgraded to ϕ_2 and ϕ_2 to ϕ_3 , ϕ_3 is compatible with ϕ_1 thanks to the compatibility of ϕ_3 with ϕ_2 and of ϕ_2 with ϕ_1 . BCT has been extended in [20], where small and large representation models are taken into account for the query and the gallery-set, respectively. Differently from BCT our method is not based on pairwise compatibility learning and does not use previous classifiers which might be incorrectly learned. Instead, we learn a representation that is directly compatible to *all* the previous representations by following a training strategy that only leverages feature stationarity.

A different compatible representation learning scenario, referred to as asymmetric metric learning, was addressed in [33]. Large network architectures are used for the gallery-set and smaller architectures for queries without considering new data for the training-set.

Compatibility was implicitly studied also in [34], [35] in which representation similarity between two networks with identical architecture but trained from different initializations was evaluated.

Neural Collapse. Our method is based on the concept of learning stationary and maximally separated features using the d -Simplex *fixed* classifier introduced in [26] and [27]. In this classifier, weights are not trainable and are determined from the coordinates of the vertices of the d -Simplex regular polytope. The goal of learning stationary and maximally separated features has similarities to the Neural Collapse phenomenon described in [36]. This phenomenon, which can be further explored in [37], shows that

the final *learnable* classifier and the learned feature representation tend to orient towards a d -Simplex structure. Along this line of research, the works [38] and [39] investigated fixing the final classifier in the scenario of balanced datasets. In contrast, in [40] it is shown that training on imbalanced datasets does not lead to Neural Collapse. However, [41] shows that Neural Collapse can occur in imbalanced and long-tail scenarios as long as the classifier is *fixed* to a d -Simplex, and this evidence has been recently supported in [42] also for out-of-distribution detection. Our experimental results on large long-tailed datasets with distribution shifts, such as [43] and [44], demonstrate that learning stationary and maximally separated features is effective even in scenarios with increased complexity due to the extra challenge of finding compatible representations.

Class-incremental Learning. Class-incremental Learning (CiL) is the process of sequentially increasing the number of classes to learn over time [45], [46], [47]. Although apparently similar to sequential learning of compatible representations, the main focus of CiL is to reduce catastrophic forgetting [48]. CiL differs from compatible representations learning in two important aspects: (1) in CiL the new model is typically initialized with the old model, while in compatible representation learning this is optional (2) in CiL the new model has not access to the whole data during the model upgrade, while in compatible representation learning all training data is available at each upgrade step. According to this, in compatible representation learning, the learned representation is not affected by catastrophic forgetting.

3 COMPATIBILITY EVALUATION

We indicate with $I_G = \{\mathbf{x}_i\}_{i=1}^N$ and $F_G = \{\mathbf{f}_i\}_{i=1}^N$ respectively the set of images and their features in the gallery-set \mathcal{G} . The gallery-set \mathcal{G} might be grouped into a number of classes or identities L according to a set of labels $\mathcal{Y} = \{y_i\}_{i=1}^L$. We assume that the features F_G are extracted using the representation model $\phi_{\text{old}} : \mathbb{R}^D \rightarrow \mathbb{R}^d$ that transforms each image $\mathbf{x} \in \mathbb{R}^D$ into a feature vector $\mathbf{f} \in \mathbb{R}^d$, where d and D are the dimensions of the feature and the image space, respectively. Analogously, we will refer to I_Q and F_Q respectively as the set of images and their features in the query-set \mathcal{Q} . The model ϕ_{old} is trained on a training-set \mathcal{T}_{old} and used to perform search tasks using a distance $\text{dist} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ to identify the closest images to the query images I_Q . As novel images \mathcal{X} become available, a new training-set $\mathcal{T}_{\text{new}} = \mathcal{T}_{\text{old}} \cup \mathcal{X}$ is created and exploited to learn a new model $\phi_{\text{new}} : \mathbb{R}^D \rightarrow \mathbb{R}^d$ that improves (i.e., upgrades) the ϕ_{old} model. Our goal is to design a training procedure to learn a compatible model ϕ_{new} so that any query image transformed with it can be used to perform search tasks against the gallery-set directly without re-indexing, i.e., without computing $F_G = \{\mathbf{f} \in \mathbb{R}^d \mid \mathbf{f} = \phi_{\text{new}}(\mathbf{x}) \forall \mathbf{x} \in I_G\}$.

3.1 Compatibility Criterion

In [18], a general criterion to evaluate compatibility was defined. According to this, a new and compatible representation model must be at least as good as its previous version in clustering images from the same class and separating those from different classes. A new representation model

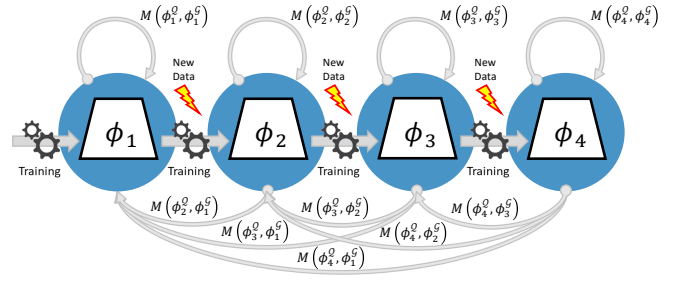


Fig. 2. Multi-model Empirical Compatibility Criterion (Eq. 5): representation models ϕ_i with $i = 1, 2, \dots, T$ are sequentially trained. Gray arrows represent self and cross-tests (example with $T = 4$).

ϕ_{new} is therefore compatible with an old representation model ϕ_{old} if:

$$\begin{aligned} \text{dist}(\phi_{\text{new}}(\mathbf{x}_u), \phi_{\text{old}}(\mathbf{x}_v)) &\leq \text{dist}(\phi_{\text{old}}(\mathbf{x}_u), \phi_{\text{old}}(\mathbf{x}_v)) \\ &\quad \forall (u, v) \in \{(u, v) \mid y_u = y_v\} \end{aligned} \quad \text{and} \quad (1)$$

$$\begin{aligned} \text{dist}(\phi_{\text{new}}(\mathbf{x}_u), \phi_{\text{old}}(\mathbf{x}_v)) &\geq \text{dist}(\phi_{\text{old}}(\mathbf{x}_u), \phi_{\text{old}}(\mathbf{x}_v)) \\ &\quad \forall (u, v) \in \{(u, v) \mid y_u \neq y_v\}, \end{aligned}$$

where \mathbf{x}_u and \mathbf{x}_v are two input samples and $\text{dist}(\cdot, \cdot)$ is a distance in feature space. Since it constrains all pairs of samples, Eq. 1 is relaxed to the following *Empirical Compatibility Criterion*:

$$M(\phi_{\text{new}}^{\mathcal{Q}}, \phi_{\text{old}}^{\mathcal{G}}) > M(\phi_{\text{old}}^{\mathcal{Q}}, \phi_{\text{old}}^{\mathcal{G}}), \quad (2)$$

where M is a metric used to evaluate the performance based on $\text{dist}(\cdot, \cdot)$. The notation $M(\phi_{\text{new}}^{\mathcal{Q}}, \phi_{\text{old}}^{\mathcal{G}})$ underlines that the upgraded model ϕ_{new} is used to extract feature vectors F_Q from query images I_Q , while the old model ϕ_{old} is used to extract features F_G from gallery images I_G . This performance value is referred to as *cross-test*. Correspondingly, $M(\phi_{\text{old}}^{\mathcal{Q}}, \phi_{\text{old}}^{\mathcal{G}})$ evaluates the case in which both query and gallery features are extracted with ϕ_{old} and is referred to as *self-test*. The underlying intuition of Eq. 2 is that model ϕ_{new} is compatible with ϕ_{old} when the cross-test is greater than the self-test, i.e., by using the upgraded representation for the query-set and the old representation for the gallery-set the system improves its performance with respect to the previous condition.

To evaluate the relative improvement gained by a new learned compatible representation, the following *Update Gain* has been defined:

$$\Gamma(\phi_{\text{new}}^{\mathcal{Q}}, \phi_{\text{old}}^{\mathcal{G}}) = \frac{M(\phi_{\text{new}}^{\mathcal{Q}}, \phi_{\text{old}}^{\mathcal{G}}) - M(\phi_{\text{old}}^{\mathcal{Q}}, \phi_{\text{old}}^{\mathcal{G}})}{M(\tilde{\phi}_{\text{new}}^{\mathcal{Q}}, \tilde{\phi}_{\text{new}}^{\mathcal{G}}) - M(\phi_{\text{old}}^{\mathcal{Q}}, \phi_{\text{old}}^{\mathcal{G}})}, \quad (3)$$

where $M(\tilde{\phi}_{\text{new}}^{\mathcal{Q}}, \tilde{\phi}_{\text{new}}^{\mathcal{G}})$ stands for the best accuracy level we can achieve by re-indexing the gallery-set with the new representation [18] and can be considered as the upper bound of the best achievable performance.

3.2 Multi-step Compatibility Criterion

In real world applications, multi-step upgrading is often required, i.e., different representation models must be sequentially learned through time, in multiple upgrade steps. At each step t , the training-set is upgraded as:

$$\mathcal{T}_t = \mathcal{T}_{t-1} \cup \mathcal{X}_t \quad (4)$$

being \mathcal{X}_t the new data and \mathcal{T}_{t-1} the training-set at step $t-1$. In the multi-step upgrading case, we define the following *Multi-model Empirical Compatibility Criterion* as follows:

$$M(\phi_{t'}^{\mathcal{Q}}, \phi_t^{\mathcal{G}}) > M(\phi_t^{\mathcal{Q}}, \phi_t^{\mathcal{G}}) \quad \forall t' > t \quad (5)$$

with $t' \in \{2, 3, \dots, T\}$ and $t \in \{1, 2, \dots, T-1\}$

where $\phi_{t'}$ and ϕ_t are two different models such that ϕ_t is upgraded before $\phi_{t'}$, T is the number of upgrade steps and M the metric used to evaluate the performance. Model $\phi_{t'}$ is compatible with ϕ_t when their cross-test is greater than the self-test of ϕ_t for each pair of upgrade steps. Fig. 2 illustrates the Multi-model Empirical Compatibility Criterion, where $\{\phi_1, \phi_2, \dots, \phi_T\}$ are the representation models, black arrows indicate the model upgrades and gray arrows represent self and cross-tests.

In order to assess multi-model compatibility of Eq. 5 for a sequence of T upgrade steps, we define the following square triangular *Compatibility Matrix* C :

$$C = \begin{pmatrix} M(\phi_1^{\mathcal{Q}}, \phi_1^{\mathcal{G}}) & & & & \\ M(\phi_2^{\mathcal{Q}}, \phi_1^{\mathcal{G}}) & M(\phi_2^{\mathcal{Q}}, \phi_2^{\mathcal{G}}) & & & \\ \vdots & \vdots & \ddots & & \\ M(\phi_T^{\mathcal{Q}}, \phi_1^{\mathcal{G}}) & M(\phi_T^{\mathcal{Q}}, \phi_2^{\mathcal{G}}) & \cdots & M(\phi_T^{\mathcal{Q}}, \phi_T^{\mathcal{G}}) & \end{pmatrix} \quad (6)$$

where each entry C_{ij} is the performance value according to metric M , taking model ϕ_i for the query-set \mathcal{Q} and model ϕ_j for the gallery-set \mathcal{G} . Entries on the main diagonal, $i = j$, represent the self-tests, while the entries off-diagonal, $i > j$, represent the cross-tests. While showing compatibility performance across multiple upgrade steps, matrix C can be used to provide a scalar metric to quantify the global multi-model compatibility in a sequence of upgrade steps. In particular, we define the *Average Multi-model Compatibility (AC)* as the number of times that Eq. 5 is verified with respect to all its possible occurrences, independently of the number of the learning steps:

$$AC = \frac{2}{T(T-1)} \sum_{i=2}^T \sum_{j=1}^{i-1} \mathbb{1}(C_{ij} > C_{jj}), \quad (7)$$

where $\mathbb{1}(\cdot)$ denotes the indicator function.

Finally, we define the *Average Multi-model Accuracy (AM)* as the average of the entries of the Compatibility Matrix:

$$AM = \frac{2}{T(T+1)} \sum_{i=1}^T \sum_{j=1}^i C_{ij} \quad (8)$$

to provide an aggregate value of the accuracy metric M under compatible training.

4 LEARNING COMPATIBLE REPRESENTATIONS

It is well known that for different initializations a neural network learns the same subspaces but with different basis vectors [34], [35]. Therefore, training the network from scratch with different randomly initialized weights *does not* provides similar representations in terms of subspace geometry. This result excludes compatibility between two independently trained representation models.

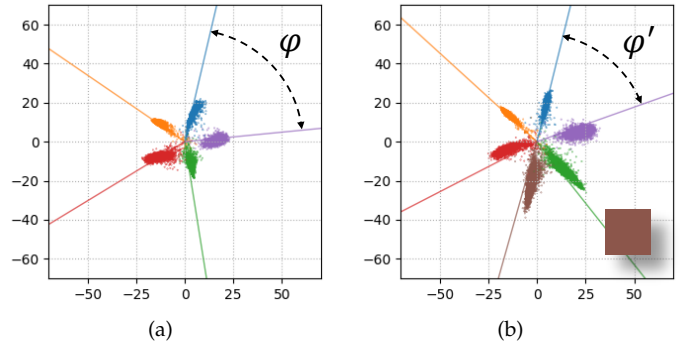


Fig. 3. Learning with incremental fine-tuning with MNIST dataset for 2D representation. Colored cloud points represent features from the test-set and gray lines represent classifier prototypes. (a) Initial configuration (5 classes); (b) Training by fine-tuning (adding the brown-class). The addition of the new class modifies the spatial configuration and angles between features.

The alternative of learning with incremental fine-tuning (i.e., weights are initialized from the previously learned model) appears to be a more favorable training procedure to compatibility. However, and perhaps counterintuitively, this does not help to keep the same subspace representation geometry regardless of the changes made.

We provide direct evidence of this aspect of feature learning in Fig. 3 with a toy problem. We trained the LeNet++ architecture [49] on a subset of the MNIST dataset setting the output size of the last hidden layer to two (so resulting in a two-dimensional representation space). The classifier weights were unit normalized and biases were set to zero to encourage learning cosine distance between features ([50], [51]) and the cross-entropy loss was used. The model was initially trained with five classes (red, orange, blue, purple, and green clouds in Fig. 3(a)); then a new class (brown cloud in Fig. 3(b)) was included in the training-set and the new model was trained by fine-tuning the old model on the new training-set of six classes. As the new class is included in the training-set and the representation is fine-tuned, the features of the old classes change their spatial configuration and the mutual angles between classifier prototypes change as well. This is due to the fact that linear classifiers maximize inter-class distance to better discriminate between classes [49]. As a consequence, the cosine distance comparison between old and new features cannot be guaranteed. The same effect holds for any number of classes and feature space dimension.

To limit such spatial configuration changes and therefore achieve feature compatibility, our approach learns stationary features exploiting the properties of fixed classifiers introduced in [27] that we briefly recall in the next subsection.

4.1 Learning Stationary Features with Fixed Classifiers

In [52], [53], a DCNN model with a fixed classification layer (i.e., not subject to learning) initialized by random weights was shown to be almost equally effective as a trainable classifier with substantial saving of computational and memory requirements. In fixed classifiers, the functional complexity of the classifier is fully demanded to the internal layers of the neural network. As the parameters of the

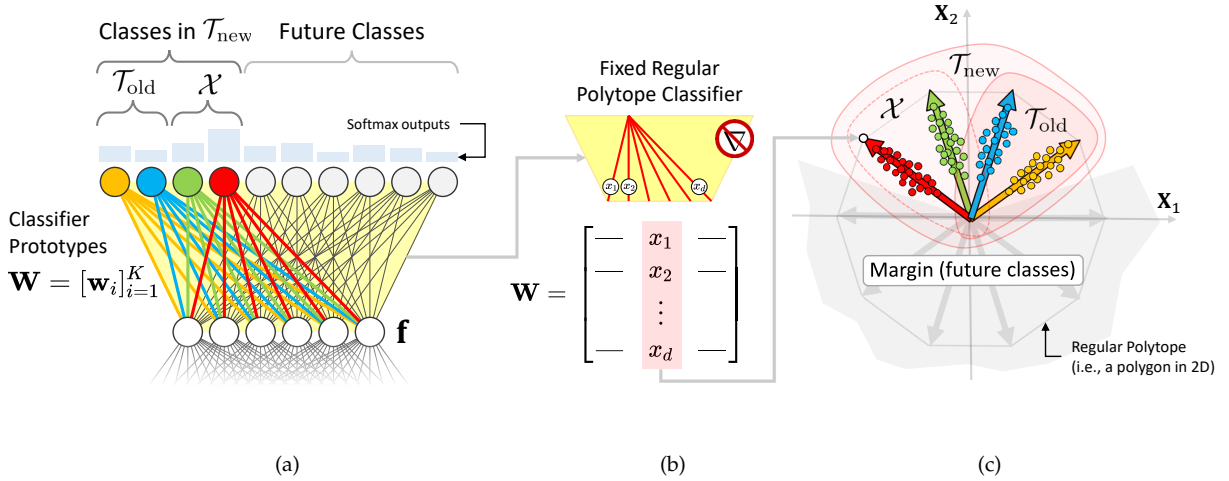


Fig. 4. Overview of the training procedure of CoReS based on feature stationarity showing: (a) fixed classifier with class prototypes $[\mathbf{w}_i]_{i=1}^K$ with evidence of those reserved for \mathcal{T}_{old} and the upgrade classes \mathcal{X} (both create \mathcal{T}_{new}), and those reserved for future classes; (b) class prototypes and their parameters x_1, x_2, \dots, x_d (parameters are the coordinate vertices of the regular polytope that defines the fixed classifier); (c) two-dimensional representation of the feature space generated by the fixed classifier. The colored point clouds represent the learned features. Prototypes of the future classes are represented with gray arrows. The gray region is the margin imposed by the future/unseen classes. As new features are learned they are pushed out from the margin.

classifier prototypes are not trainable, only the feature vector directions align toward the directions of the prototypes. In [27], we presented a special class of fixed classifiers where the weights of the classifier are fixed to values taken from the coordinates of the vertices of regular polytopes. Regular polytopes generalize regular polygons in any number of dimensions, and reflect the tendency of splitting the available space into approximately equiangular regions. There are only three possible polytopes in a multi-dimensional feature space with dimensions 5 and higher: the d -Simplex, the d -Cube and the d -Orthoplex. In these classifiers, the spatial configuration of the learned classes remains stationary as new classes are added and, at the same time, classes are maximally separated in the representation space [25].

4.2 Learning Compatibility via Stationarity

Compatibility has a close relationship with feature stationarity. In fact, stationarity requires that the representation that is learned in the future is statistically indistinguishable from the representation learned in the past, regardless of the number of model upgrades performed over time. We therefore argued that a compatibility training procedure should be defined by directly exploiting stationarity of the representation as provided by fixed classifiers.

According to this, in order to maintain the compatibility across model upgrades, at each upgrade we learn the features of the new classes in reserved regions of the representation space, while the features learned in the previous upgrades will not “move” thanks to the stationarity property. Region reservation is made at the beginning of training by setting a number of classifier outputs greater than the number of the initial classes and keeping not assigned regions for future upgrades. Fig. 4 provides an overview of the CoReS procedure for compatible representation learning. Fig. 4(a) shows the initial training-set \mathcal{T}_{old} (two classes for simplicity), the new data \mathcal{X} (two classes for simplicity), the regions left for the future classes, the class prototypes

\mathbf{w}_i , and the softmax outputs of the fixed classifier. As a result, the training-set \mathcal{T}_{new} that will be used to upgrade the initial representation has four classes. Fig. 4(c) shows the representation space of a 10-sided 2D regular polytope (a polygon) with the old and new classes, their prototypes and class samples. The regions reserved for the future classes are colored in gray. Fig. 4(b) highlights the coordinates of one of the fixed classifier prototypes.

As new data \mathcal{X}_t is available at time t to upgrade the representation, CoReS performs optimization of model ϕ_t according to the following loss:

$$\mathcal{L}_t = -\frac{1}{N_b} \sum_{i=1}^{N_b} \log \left(\frac{e^{\mathbf{w}_{y_i}^\top \phi_t(\mathbf{x}_i)}}{\sum_{j=1}^{K_t} e^{\mathbf{w}_j^\top \phi_t(\mathbf{x}_i)} + \sum_{j=K_t+1}^K e^{\mathbf{w}_j^\top \phi_t(\mathbf{x}_i)}} \right), \quad (9)$$

where: $\mathbf{x}_i \in \mathcal{T}_{t-1} \cup \mathcal{X}_t$ is a sample instance; $y_i \in \{1, 2, \dots, K_t\}$ its class; $\phi_t(\mathbf{x}_i) \in \mathbb{R}^d$ its feature vector as from model ϕ_t ; $\mathbf{W} \in \mathbb{R}^{d \times K}$ the fixed classifier weight matrix; K_t and K , respectively the number of classes in $\mathcal{T}_{t-1} \cup \mathcal{X}_t$ and the number of classifier outputs; and N_b the number of samples in the mini-batch.

For the $K - K_t$ outputs that have not yet been associated to classes at time t , the classifier responds with false positives of the training-set \mathcal{T}_t . To highlight the role of these outputs, in the denominator of the loss in Eq. 9, we distinguished the contributions corresponding to the K_t classes at time t (the first summation) from the $K - K_t$ future classes (the second summation). The contribution of the future classes enforces an angular margin penalty in the representation space that tends to push away the features already learned from the prototypes of future classes. Such a margin penalty has a key role for learning compatible features as, when upgrading the model with new data, it prevents the new data to affect the representation already learned. As is evident from Eq. 9, CoReS learns feature

TABLE 1

Compatibility comparative evaluation for open set verification on CIFAR 100/10 with one upgrade. Models are learned with 50% and upgraded with 100% of training data. Subscripts indicate the method used for compatibility learning. Columns indicate accuracy (M), whether the Empirical Compatibility Criterion is verified (ECC), the Update Gain (Γ) and the Absolute Gain.

COMPARISON PAIR	M	ECC	Γ (%)	ABSOLUTE GAIN
(ϕ_{old}, ϕ_{old}) (Lower Bound)	0.60	–	–	–
$(\phi_{new-ITM}, \phi_{old})$	0.12	×	–	–
$(\phi_{new-\ell^2}, \phi_{old})$	0.35	×	–	–
$(\phi_{new-IFT}, \phi_{old})$	0.57	×	–	–
$(\phi_{new-LwF}, \phi_{old})$	0.59	×	–	–
$(\phi_{new-BCT}, \phi_{old})$	0.60	✓	5.9	0.20
$(\phi_{new-CoReS}, \phi_{old})$	0.61	✓	21.3	1.05
$(\tilde{\phi}_{new}, \tilde{\phi}_{new})$ (Upper bound)	0.65	–	–	0.49

compatibility without using the previously learned model or classifier.

In our implementation, CoReS uses the d -Simplex fixed classifier as it exploits the only polytope that provides class prototypes equidistant from each other [27]. In this way, we avoid bias in the assignment of class labels to classifier outputs. The d -Simplex fixed classifier defined in a feature space of dimension d can accommodate a number of classes equal to its number of vertices, i.e., $K = d + 1$. As a result, the K classifier prototypes are computed as:

$$\mathbf{W} = \left[e_1, e_2, \dots, e_{K-1}, \frac{1 - \sqrt{K}}{K-1} \sum_{i=1}^{K-1} e_i \right] \quad (10)$$

where e_i denotes the standard basis in \mathbb{R}^d , with $i \in \{1, 2, \dots, K-1\}$. Increasing the dimensionality of the representation has no effect on both the training time and memory consumption, since fixed classifiers do not require back-propagation on the fully connected layers [27], [53].

4.3 CoReS Extension by Model Selection

At each upgrade, the training-set grows. This causes a shift in the distribution between the training-set and the gallery-set [54]. To cope with this shift, we extended CoReS with a simple model selection strategy, based on cross validation.

We assume that: (1) drifting occurs gradually such that the density ratio of the marginal distribution before and after the upgrade is close to uniform [55]; (2) there is a set of data $\hat{\mathcal{G}}$ available that is an i.i.d. sample of the gallery-set data \mathcal{G} (i.e., a split following the same distribution); (3) the previously learned model is available. Under these assumptions, at each upgrade step, for each epoch, we search the model with the highest self-test among the models that satisfy the compatibility criterion of Eq. 2 with respect to the previous model:

$$\begin{aligned} \phi_{t'} &= \arg \max_l M(\phi_{t',l}^Q, \phi_{t',l}^{\hat{\mathcal{G}}}) \\ \text{s. t.} \quad &M(\phi_{t',l}^Q, \phi_{t'}^{\hat{\mathcal{G}}}) > M(\phi_t^Q, \phi_t^{\hat{\mathcal{G}}}) \end{aligned} \quad (11)$$

where $t = t' - 1$ and l is the index of the training epoch.

TABLE 2

Compatibility evaluation for open set verification on CIFAR 100/10 with two upgrades with CoReS and BCT. Models are learned with with 33%, and upgraded with 66% and 100% of training data. Columns indicate accuracy(M), whether the Empirical Compatibility Criterion is verified (ECC) and the Update Gain (Γ).

COMPARISON PAIR	CORES			BCT		
	M	ECC	Γ (%)	M	ECC	Γ (%)
(ϕ_1, ϕ_1)	0.59	–	–	0.59	–	–
(ϕ_2, ϕ_1)	0.61	✓	54.7	0.59	✓	5.2
(ϕ_3, ϕ_1)	0.60	✓	18.5	0.58	×	–
(ϕ_2, ϕ_2)	0.63	–	–	0.61	–	–
(ϕ_3, ϕ_2)	0.61	×	–	0.59	×	–
(ϕ_3, ϕ_3)	0.65	–	–	0.64	–	–

5 EXPERIMENTAL RESULTS

In this section, we compare CoReS against baselines and the *Backward Compatible Training* (BCT) [18] state-of-the-art method for different visual search tasks on different benchmark datasets. In particular, in Sec. 5.1 we evaluate open-set verification with single and multi-model upgrading on the CIFAR-100/10 datasets [57]; in Sec. 5.2 and Sec. 5.3, we analyze single and multi-model upgrading in more challenging tasks, namely face verification (in Sec. 5.2) on the CASIA-WebFace/LFW datasets [58], [59] and person re-identification (in Sec. 5.3) on the Market1501 dataset [60]; in Sec. 5.4 and Sec. 5.5, we evaluate CoReS in challenging long-tail distribution datasets, namely Google Landmark Dataset v2 [43] and MET dataset [44]. The original datasets are split so that training-sets used for each upgrade have the same number of classes. Finally, in Sec. 5.6, we report a qualitative analysis that provides evidence of how stationarity contributes to compatibility. All the representation models were trained using 4 NVIDIA Tesla A-100 GPUs.

5.1 Compatibility Evaluation on CIFAR-100/10

We perform open-set verification with CIFAR-100/10 datasets. The CIFAR-100 and CIFAR-10 datasets consist of 60,000 32×32 RGB images (50,000 for training and 10,000 for test) in 100 and 10 classes, respectively. Their classes are mutually exclusive. We use the CIFAR-100 training-set to create the training-sets for upgrading the models and the CIFAR-10 test-set for generate the verification pairs used in the open-set verification protocol. We do not use the test-set of CIFAR-100 and the training-set of CIFAR-10.

The SENet-18 architecture [61] adapted to the CIFAR 32×32 image size¹ is used. The same random seed is used for initialization every time the model is upgraded. For CoReS, the output nodes of the d -Simplex fixed classifier are set to $K_p = 100$ so resulting in a feature space is 99-dimensional. Optimization is performed using SGD with 0.1 learning rate, 0.9 momentum, and $5 \cdot 10^{-4}$ weight decay. The batch size is set to 128. With every upgrade, training is terminated after 100 epochs. Learning rate is scheduled to decrease to 0.01 after 70 epochs.

1. SENet-18 code: <https://github.com/kuangliu/pytorch-cifar>

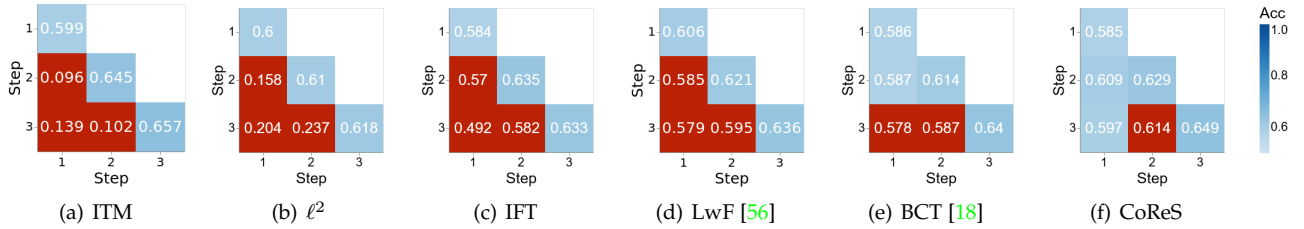


Fig. 5. Compatibility matrices for open set verification on CIFAR-100/10 with 2-step multi-model upgrading with different methods compared. Models are sequentially learned on CIFAR-100 with 33%, 66%, 100% of data. Diagonal and off-diagonal elements respectively report the self-test and the cross-test accuracy. Models that do not satisfy the compatibility criterion are highlighted in red.

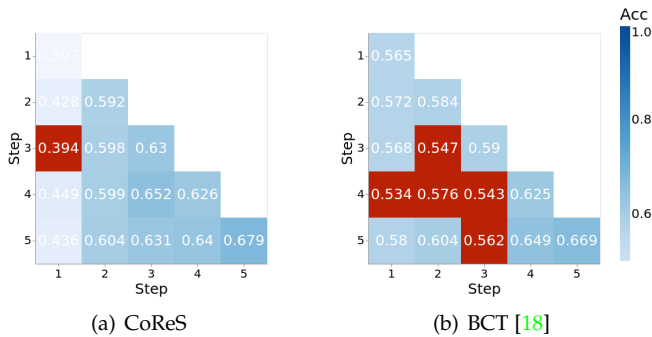


Fig. 6. Compatibility matrices of CoReS and BCT for open set verification on CIFAR-100/10 with 4-step multi-model upgrading. Diagonal and off-diagonal elements respectively report self-test and cross-test accuracy. Models are sequentially learned on CIFAR-100 with 20%, 40%, 60%, 80%, 100% of data. Models that do not satisfy the compatibility criterion are highlighted in red.

We compared CoReS against the ℓ^2 , the *Incremental Fine-Tuning* (IFT), the *Learning without Forgetting* (LwF) [56] baselines and the *Backward Compatible Training* (BCT) method [18]. All the baselines except IFT, are the public implementations of [18]². The ℓ^2 approach achieves compatibility using an auxiliary loss that takes into account the Euclidean distance ℓ^2 between features of images in \mathcal{T}_{old} as evaluated by $\phi_{\text{new-}\ell^2}$ and ϕ_{old} , that is:

$$\mathcal{L}_{\text{dist}} = \frac{1}{|\mathcal{T}_{\text{old}}|} \sum_{\mathbf{x} \in \mathcal{T}_{\text{old}}} \text{dist}(\phi_{\text{new-}\ell^2}(\mathbf{x}), \phi_{\text{old}}(\mathbf{x})).$$

According to this, the new representation model $\phi_{\text{new-}\ell^2}$ is trained as:

$$\phi_{\text{new-}\ell^2} = \arg \min_{\phi} \left(\mathcal{L}(\phi, \mathcal{T}_{\text{new}}) + \lambda \mathcal{L}_{\text{dist}}(\phi, \phi_{\text{old}}, \mathcal{T}_{\text{old}}) \right)$$

where \mathcal{L} is the standard cross-entropy loss and the scalar λ balances the two losses. The baseline ITM (Independently Trained Models) learns the representation of two versions of models trained independently with standard cross-entropy loss. When the gallery is re-indexed, ITM may be thought of as the standard approach to observe, if any, a reduction in performance of a compatibility-based method. The decrease is meant to accommodate the increased requirement that the learning model must fulfill to learn a compatible representation. ITM and CoReS are related since both methods use standard cross-entropy loss and linear classifiers. This relationship allows to quantify the representation capability

used for compatibility purposes. Indeed, since the self-test values in the compatibility matrices are calculated with both the query and the gallery images extracted from the same model (i.e., re-indexing the gallery), the difference of the main diagonal entries of ITM and CoReS quantifies the neural network model capacity used to learn compatibility.

To perform open-set verification, we use 3,000 positive and 3,000 negative pairs randomly generated from the CIFAR-10 test-set. Given a pair of images of CIFAR-10, verification assesses whether they are of the same class using the cosine distance. One of the images of each pair can be considered gallery and the other the query. Compatibility is evaluated for one, two, four and nine upgrade steps.

In the one-upgrade case, the model is trained with 50% of CIFAR-100 classes and upgraded using 100%. The comparative evaluation is shown in Tab. 1 in terms of verification accuracy (M), whether the Empirical Compatibility Criterion is verified (ECC) and Update Gain (Γ). It can be noticed that only BCT and CoReS satisfy Eq. 2. However, CoReS obtains higher Update Gain with respect to BCT.

In the two-upgrade case, the model is initially trained on 33% of CIFAR-100 classes and upgraded with 66% and 100%. The compatibility matrix is shown in Fig. 5 for all the methods compared. We can notice that none of the methods achieves full compatibility in this case, however BCT and CoReS achieve higher accuracy values with respect to the others. CoReS misses compatibility only once between ϕ_3 and ϕ_2 . In Tab. 2 the behavior of CoReS and BCT in the two-upgrade case is reported in more detail. We can observe that CoReS has a substantially higher compatibility and update gain than BCT. This may be due to the fact that BCT learns compatibility of ϕ_3 with no direct relationship with ϕ_1 , while CoReS learns ϕ_3 taking into account both ϕ_1 and ϕ_2 , because all models share the same class prototypes in the feature space.

The compatibility matrices of CoReS and BCT for the cases of four and nine upgrade steps are shown in Fig. 6 and 7, respectively. The superior behavior of CoReS with respect to BCT is clearly evident in both cases, scoring 0.9 versus 0.5 and 0.51 versus 0.11 of AC , respectively. This indicates that the more upgrades are made, the better CoReS performs with respect to BCT.

5.2 Compatibility Evaluation on CASIA-WebFace/LFW

In this evaluation, we perform open-set face verification using CASIA-WebFace dataset to create the training-sets and LFW as the test set. The CASIA-WebFace dataset includes 494,414 RGB face images of 10,575 subjects. The

2. <https://github.com/YantaoShen/openBCT>

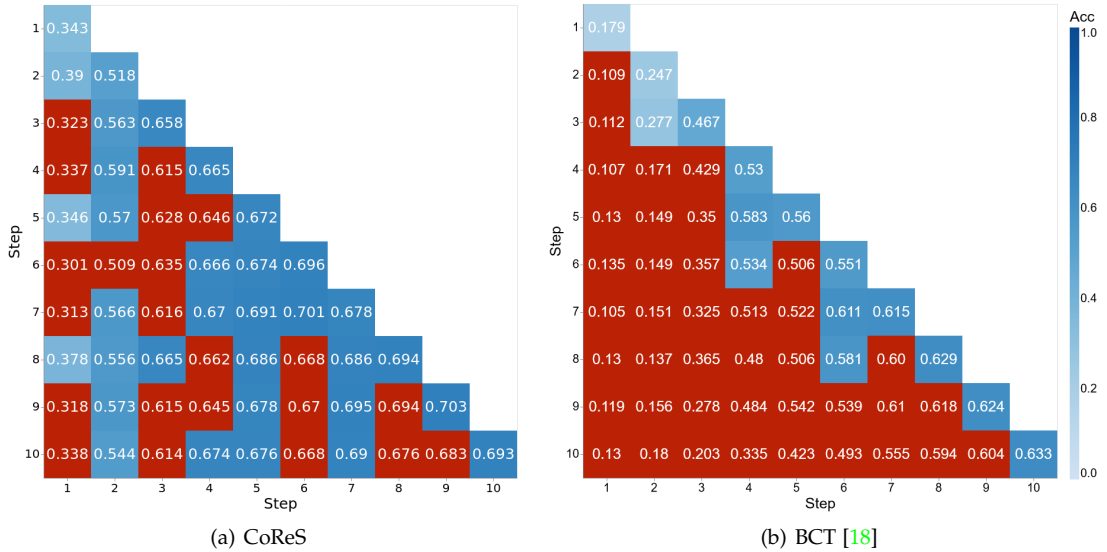


Fig. 7. Compatibility matrices of CoReS and BCT for open-set verification on CIFAR 100/10 with 9-step multi-model upgrading. Diagonal and off-diagonal elements report self-test and cross-test accuracy, respectively. Models are sequentially learned with 10%, 20%, ..., 100% of training data. Models that do not satisfy compatibility are highlighted in red.

TABLE 3

Compatibility of CoReS and BCT for open-set face verification on the CASIA-WebFace/LFW dataset with one upgrade. Models are learned on CASIA-WebFace with 50% and upgraded with 100% of data. Columns indicate accuracy (M), whether the Empirical Compatibility Criterion is verified (ECC) and the Update Gain (Γ).

COMPARISON PAIR	CORES			BCT		
	M	ECC	Γ (%)	M	ECC	Γ (%)
(ϕ_{old}, ϕ_{old})	0.90	–	–	0.91	–	–
(ϕ_{new}, ϕ_{old})	0.91	✓	0.32	0.91	✓	0.29
(ϕ_{new}, ϕ_{new})	0.92	–	–	0.91	–	–

LFW dataset contains 13,233 target face images of 5,749 subjects. Of these, 1,680 have two or more images, while the remaining 4,069 have only one single image. ResNet50 [62] with input size of 112×112 is used as backbone. Optimization is performed using SGD with 0.1 learning rate, 0.9 momentum, and $5 \cdot 10^{-4}$ weight decay. The batch size is 1,024. With every upgrade, training is terminated after 120 epochs. Learning rate is scheduled to decrease to 0.01, 0.001, and 0.0001 at epoch 30, 60, 90 respectively.

Compatibility is evaluated for one, two, three, four, five, and nine upgrade steps.

In the one-upgrade case, models are learned with 50% of the CASIA-WebFace dataset and upgraded with 100%. Tab. 3 shows that, with 50% of CASIA-WebFace dataset, CoReS and BCT achieve similar performance. This is because there is already sufficient data variability to learn compatible features.

A big difference between CoReS and compared methods becomes evident when multiple upgrades are considered as shown in Fig. 8(a). In this figure, the values of the AC are reported for each method over a set of different experiment respectively with one, two, three, four, five, and nine upgrades. CoReS achieves full compatibility ($AC = 1$) for one, two, three, and four upgrades and starts decreasing

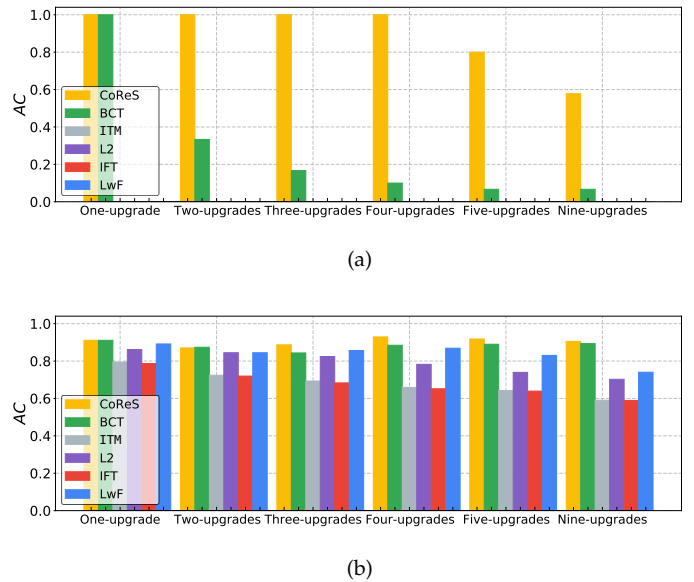


Fig. 8. Compatibility of CoReS and compared methods (shown color-coded) for open-set face verification on the CASIA-WebFace/LFW dataset with multi-model upgrading. Bins show: (a) AC scores for different number of upgrades; (b) AM scores for different number of upgrades.

ing AC from five upgrade steps up to $AC = 0.58$ with nine upgrades. In contrast BCT loses performance already with two upgrade steps finishing at $AC = 0.09$ with nine upgrade steps. Baseline methods report $AC = 0$ in all of the scenario. The compatibility matrices for the case of four upgrades are shown in detail in Fig. 9(f). In this case, we can appreciate full compatibility of CoReS with respect to the poor compatibility of BCT and other methods. In Fig. 8(b), we report the AM metric for these experiments. It can be observed that the CoReS and BCT score almost the same average verification accuracy in all the experiments, while

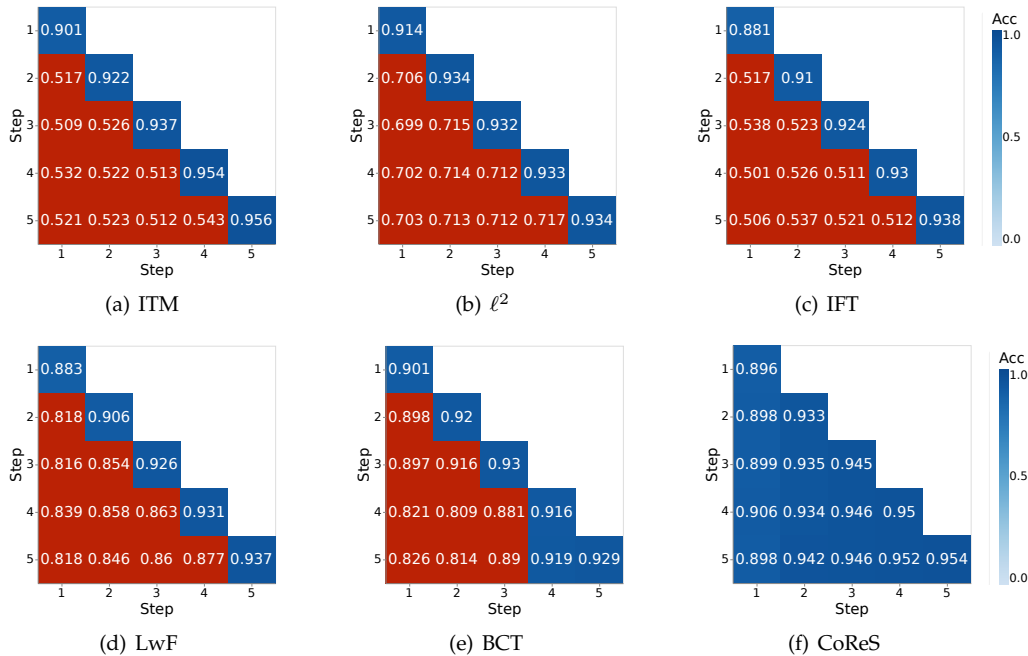


Fig. 9. Compatibility matrices of CoReS and compared methods for open-set verification on CASIA-WebFace-LFW dataset with 4-step multi-model upgrading. Models are sequentially learned on CASIA-WebFace with 20%, 40%, 60%, 80%, 100%. Models that do not satisfy compatibility are highlighted in red.

TABLE 4
Compatibility evaluation for person re-identification on the Market1501 dataset for one and two-step upgrading. AC and AM are used to compare CoReS with other methods. mAP is used as evaluation metric M .

METHOD	One-upgrade		Two-upgrades	
	AC	AM	AC	AM
STANDARD	0	0.30	0	0.22
ℓ^2	0	0.36	0	0.31
IFT	0	0.30	0	0.21
LwF	0	0.40	0	0.36
BCT	1	0.49	0.3	0.50
CoReS	1	0.57	1	0.51

in the others values are always lower. This is due to the fact that cross-test values are low since no compatibility is reported.

We conclude that for multi-model upgrading, CoReS, while having the same verification performance of BCT, largely improves compatibility across model upgrades with 544% relative improvement over BCT for the challenging scenario of nine-step upgrading. The lower AC of BCT appears to be related to the fact that in this method compatibility is obtained only through transitivity from the model previously learned.

5.3 Compatibility Evaluation on Market1501

In this experiment, we perform person re-identification (1:N search) using the Market1501 dataset. Differently from CASIA-WebFace/LFW, Market1501 includes images of identities with severe occlusions and pose variations. This makes learning compatible features largely more challenging. The

Market1501 dataset contains 1,501 identities, split in 751 for training and 750 for test. The two sets do not share identities. Images of each identity are captured by several cameras so that cross-camera search can be performed. For 1:N re-identification, a set of templates is used in the gallery-set. Then templates of the query-set are used to search against the gallery templates. Mean Average Precision (mAP) is used as quality metric. Following [63], we use a pre-trained ResNet101 [62] as backbone and Adam [64] as optimizer with an initial learning rate of $3 \cdot 10^{-4}$. The batch size is 256. With every upgrade, training is terminated after 25 epochs. Learning rate is scheduled to decrease to $3 \cdot 10^{-5}$ after 20 epochs.

Tab. 4 shows the results for one and two upgrade steps. We observe that CoReS always achieves full compatibility, while BCT loses compatibility already with two upgrade steps. The baseline methods never achieve compatibility in both one and two-step upgrade. While it is evident that the stationarity of the CoReS representation still plays a key role to ensure compatibility across the upgrades also in this challenging task, the results confirm the limits of transitive compatibility used by BCT.

5.4 Compatibility Evaluation on GLDv2

Google Landmark Dataset v2 (GLDv2), cleaner version, consists of 81,313 classes and 1,580,470 images [43]. The main issue is that with such a large number of classes, the d -Simplex classifier matrix will be large and will require additional GPU memory. Although we have already evaluated the fixed classifier by training with more than 10k classes using CASIA-WebFace dataset, with 81k classes it is necessary to also reduce the representation dimension. This is a result of the linear growth of the feature dimension of

TABLE 5

Compatibility of CoReS and BCT on the retrieval setup of GLDv2 with 1-step upgrade. Models are learned on GLDv2 with 50% of classes and upgraded with 100% of data. Columns indicate mAP@100 (M), whether the Empirical Compatibility Criterion is verified (ECC) and the Update Gain (Γ).

COMPARISON PAIR	CORES			BCT		
	M	ECC	Γ (%)	M	ECC	Γ (%)
(ϕ_{old}, ϕ_{old})	0.19	–	–	0.18	–	–
(ϕ_{new}, ϕ_{old})	0.20	✓	0.15	0.16	×	–
(ϕ_{new}, ϕ_{new})	0.21	–	–	0.19	–	–

the d -Simplex fixed classifier with the number of classes. However, training can still be achieved effectively because fixed weights do not require additional storage for the gradient, thus resulting in a decrease in the amount of memory needed. The challenging situation remains at test time in which the gallery-set is composed of more than 760k elements, each of which is 81,312 dimensional, requiring more than 200Gb of storage. With such a memory requirement, the Faiss GPU accelerated indexing [65] available in the GLDv2 repository³ may not be fully suitable for use. To overcome this issue, inspired by [66], we adopted top- k sparsification⁴ and perform nearest neighbor using the PySparNN library⁵. Tab. 5 shows the retrieval results on GLDv2 with one-step upgrading for both CoReS and BCT. In our experiments, we set the top- k sparsification to $k = 512$, use the ResNet50 backbone and image size 256×256 . As it can be seen from the table, CoReS achieves compatibility also under this learning settings, while BCT results not compatible as the cross-test is 0.16 which is lower than the previous self-test which results 0.18. CoReS performs better than BCT in learning compatible representation with a long-tail distribution dataset. One possible reason is that CoReS is based on a d -Simplex fixed classifier that has recently been shown to perform better than a learnable classifier in long-tail/imbalanced datasets [41], [42].

5.5 Compatibility Evaluation on MET

MET, mini version, is an instance-level artwork recognition dataset consisting of 38,307 images and 33,501 classes [44]. The main issue for this dataset is that standard supervised learning with cross-entropy is not converging. The official repository⁶ provides a self-supervised training baseline based on contrastive learning [1] which, starting from a pre-trained model on ImageNet, converges to a working representation. We have extended CoReS to support the contrastive learning setup. The extension consists in using the output logits of the d -Simplex fixed classifier as the contrasted representation. Since classes are not required in contrastive learning, under this condition the dimension d of the fixed classifier becomes a free parameter (we set $d = 2048$). Our training in miniMET starts by pre-training a d -Simplex fixed classifier based ResNet50 on ImageNet as in

3. <https://github.com/cvdfoundation/google-landmark>

4. All the entries of a feature representation vector are set to zero except for the top k .

5. <https://github.com/facebookresearch/pysparnn>

6. <https://github.com/nikosips/met>

TABLE 6

Compatibility of CoReS and BCT on the kNN classification of MET with 1-step upgrade. Models are learned on miniMET with 50% of classes and upgraded with 100% of data. Columns indicate accuracy (M) and whether the Empirical Compatibility Criterion is verified (ECC) and the Update Gain (Γ).

COMPARISON PAIR	CORES			BCT		
	M	ECC	Γ (%)	M	ECC	Γ (%)
(ϕ_{old}, ϕ_{old})	0.59	–	–	0.63	–	–
(ϕ_{new}, ϕ_{old})	0.65	✓	0.6	0.54	×	–
(ϕ_{new}, ϕ_{new})	0.69	–	–	0.68	–	–

[27]. Contrastive learning with CoReS is then accomplished by using a distinct d -Simplex fixed final matrix which operates as a projector rather than a classifier. MiniMET is used to create the training-sets for the upgrade and the gallery-set, while the MET dataset of the queries is used as query-set. Tab. 6 shows the result with one-step upgrading of CoReS and BCT. As evidenced by the table, only CoReS achieves compatibility. This finding further supports the positive performance of the proposed method, and suggests that it may also be used for contrastive/self-supervised learning. The positive results may be motivated by the fact that the extension of CoReS to contrastive learning is related to DirectCLR [67] in which the SimCLR contrastive method [68] is trained with a final fixed projector based on random weights. We can consider the d -Simplex classifier in this “contrastive extended CoReS” as a special fixed projector. The positive results we achieved are consistent with the improved results on SimCLR reported by DirectCLR.

5.6 Qualitative Results

Fig. 10 shows the different behaviors of IFT, BCT, and CoReS for a simple two-step upgrading with the MNIST dataset. In this experiment, the training set is upgraded from seven to eight classes and then to nine classes using the same setting as in Sec. 4.

It can be observed that IFT changes the spatial configuration of the representation as every new class is learned. In fact, IFT has no mechanism to prevent the rearrangement of the features when the model is upgraded. By exploiting the class prototypes of the classifier at the previous upgrade step, BCT keeps the representation reasonably stationary, although small changes are introduced to accommodate the new classes. Differently from the others, the features learned by CoReS remain aligned with the class prototypes. The configuration of the feature space does not change, thus providing compatibility with previous representations.

6 ABLATION STUDIES

As CoReS is a single building block method, ablation study consists in tuning training hyperparameters. In the following, we present experiments to evaluate the factors that affect the performance of our training procedure. In particular: in Sec. 6.1, we discuss how much the number of training epochs affects learning compatible representations; in Sec. 6.2, we evaluate the influence of Model Selection over compatibility; in Sec. 6.3, we analyze the effects of starting

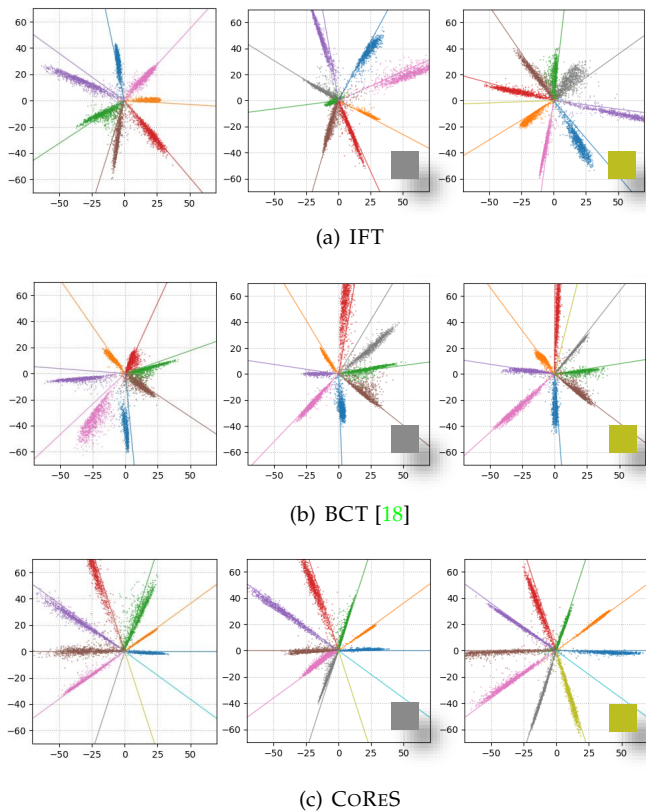


Fig. 10. Feature spatial configuration with 2-step upgrading (1 class per upgrade) and MNIST dataset with 2D features. Colored cloud points are the features from the test-set and lines represent the classifier prototypes. Compatibility methods compared: (a) IFT; (b) BCT; (c) CoReS.

from a pre-trained model; in Sec. 6.4, we examine model initialization, whether using same and different random initialization or fine-tuning from a previously learned model; in Sec. 6.5, we consider the effect of using different class sequence ordering; in Sec. 6.6 we study how much different network architectures impact compatibility; in Sec. 6.7, we examine whether the number of pre-allocated future classes in the d -Simplex classifier impacts the performance of CoReS. Finally, in Sec. 6.8, we explore whether different source of upgrade data influences the compatibility performance.

The ablation experiments are performed on the verification task of Sec. 5.1 using CIFAR-100/10 dataset with nine upgrades.

6.1 Number of Epochs

In this experiment, models were learned sequentially with 30, 70, 100, 200 and 350 epochs at each upgrade with 9-step upgrading. We modified the learning rate as follows: no change for the 30-epochs case; scheduled to decrease to 0.1 at epoch 50 for the 70-epochs case; scheduled to decrease to 0.01 at epoch 70 for the 100-epochs case (already discussed in Sec. 5.1); scheduled to decrease to 0.1 at epoch 70 and 0.01 at epoch 140 for the 200 epochs case; scheduled to decrease to 0.1 at epoch 150 and to 0.01 at epoch 250 for the 350 epochs case.

Fig. 11 shows plots of AC (Fig. 11(a)) and AM (Fig. 11(b)) for this experiment. We can notice that the

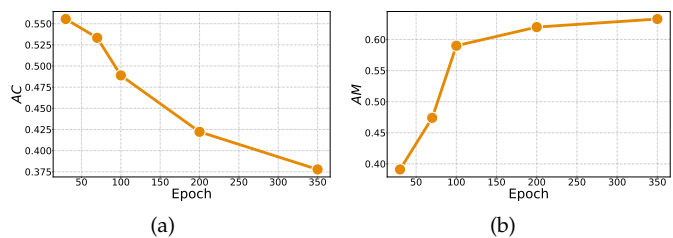


Fig. 11. Compatibility of CoReS as a function of the number of training epochs (30, 70, 100, 200, 350) at each upgrade, with 9-step upgrading. AC (a) and AM (b).

TABLE 7
Compatibility of CoReS for training from scratch without Model Selection (CoReS w/O MS) and with Model Selection (CoReS w MS) at each upgrade, with 9-step upgrading over 20 runs. Mean and standard deviation of AC and AM .

	AC	AM
CoReS w/o MS	0.51 ± 0.10	0.59 ± 0.08
CoReS w MS	0.53 ± 0.11	0.59 ± 0.09

trade-off between AC and AM occurs near 100 epochs. As AC decreases, the AM loses its meaning as a metric for compatibility evaluation.

6.2 Model Selection

In Sec. 4.3, we considered the case in which at each upgrade the training-set is increased in size and extended the basic CoReS formulation introducing compatibility learning with Model Selection. Tab. 7 compares AC and AM of the basic CoReS implementation against CoReS with Model Selection on CIFAR-100 and nine upgrade steps in 20 runs. It can be observed that Model Selection provides some improvement of AC while keeping the same AM . A larger improvement is observed when Model Selection combined with pre-training as discussed in the next subsection.

6.3 Pre-training

In this section we analyze the effect of pre-training over compatibility. We reserved 50 classes of CIFAR-100 to pre-train the model and 50 classes to create 10 distinct training-sets to upgrade the model. The 50-classes pre-trained model was used for initialization. Five classes were added at each upgrade. Since the size of pre-training changes over the upgrades, we expected to observe the correlation between pre-training and compatibility.

The values of AM and AC in relationship with the number of classes leveraged for pre-training are reported in Tab. 8 for CoReS without and with Model Selection over 20 runs. From the table, it appears that the more data are used for pre-training the more the model improves the AC and the AM . The Model Selection strategy provides an additional increase of AC of about 8% at any level of pre-training with nearly the same improvement of AM .

For the sake of comparison, we implemented pre-training with Model Selection also on BCT in the same test scenario. The compatibility matrices of CoReS and BCT are shown in Fig. 12(a) and Fig. 12(b), respectively. We see that

TABLE 8

Compatibility of CoReS for training from a pre-trained model without (w/o MS) and with (w MS) Model Selection at each upgrade, with 9-step upgrading over 20 runs. Mean and standard deviation of AC and AM as a function of the number of classes in the pre-trained model.

PRE-TRAINING #CLASSES	w/o MS		w MS	
	AC	AM	AC	AM
20	0.56 ± 0.16	0.65 ± 0.13	0.64 ± 0.11	0.64 ± 0.09
30	0.60 ± 0.14	0.67 ± 0.14	0.68 ± 0.09	0.65 ± 0.12
40	0.62 ± 0.19	0.69 ± 0.09	0.71 ± 0.12	0.67 ± 0.08
50	0.67 ± 0.13	0.70 ± 0.12	0.74 ± 0.14	0.69 ± 0.11

TABLE 9

Compatibility of CoReS for different model initializations at each upgrade, with 9-step upgrading over 20 runs: with same random initialization (SAME); with a different random initialization (RANDOM); with fine tuning from the previously learned model (FINE-TUNED). Mean and standard deviation of AC and AM .

INITIALIZATION	AC	AM
SAME	0.51 ± 0.10	0.59 ± 0.08
RANDOM	0.45 ± 0.09	0.60 ± 0.11
FINE-TUNED	0.38 ± 0.12	0.57 ± 0.13

CoReS achieves compatibility in almost all the upgraded models, achieving an AC of 0.78, with an evident large improvement over BCT.

6.4 Model Initialization

In our previous experiments, the network was trained from scratch at each upgrade, using the same, randomly selected, initial parameters. According to this, all the upgraded models started optimization from the same configuration of parameters. Alternatively, parameters can be randomly selected at each upgrade. As a third option, learning at each upgrade can also be performed by starting from the previous upgrade and then performing fine-tuning incrementally.

Tab. 9 reports mean and standard deviation of AC and AM for 9-step upgrading and 20 runs for these three cases (SAME, RANDOM, and FINE-TUNED, respectively). It can be noticed that random initialization at each upgrade negatively affects the final performance, resulting into a sensible reduction of AC with respect to initialization with the same parameters at each upgrade. This result can be supported by the observations in [69], [70] and is likely related to the concept of “flat minima” [71], [72] and to the fact that with the same initialization, the SGD optimization explores the same basin of the loss landscape near the minimum. We argue that the basins of the loss landscape near the minimum explored by CoReS across upgrades has close relationship with the stationarity/compatibility of the learned representation. On its turn, FINE-TUNED results into a strong reduction in AC . We argue that this is due to the fact that in this case, differently from learning from scratch, optimization should go back to earlier weight configurations to find more complex error landscapes.

6.5 Different class order

Tab. 10 reports mean and standard deviation of AC and AM of 20 runs for different sequential class orderings, namely:

TABLE 10

Compatibility of CoReS as a function of sequential class ordering with 9-step upgrading over 20 runs: with alphabetical order (ALPHABETICAL); with different random permutation of the classes at each run (RANDOM). Mean and standard deviation of AC and AM .

CLASS ORDERING	AC	AM
ALPHABETICAL	0.51 ± 0.10	0.59 ± 0.08
RANDOM	0.52 ± 0.13	0.57 ± 0.05

(1) classes are alphabetically ordered, (ALPHABETICAL); (2) each run has a different random permutation of the classes (RANDOM), which we adopted as default choice in our experiments. It can be noticed, AC and AM metrics are almost the same in the two cases showing that class ordering has essentially no effect on learning compatible features.

6.6 Different Model Architectures

It has been shown that the expressive power of the network architecture has impact on the classification accuracy of fixed classifiers since the complexity of learning compatible representations is fully demanded to the internal layers of the neural network [27], [53].

According to this, it is relevant to evaluate the impact of the expressive power of network architectures over CoReS compatible learning. We compared different architectures with increasing expressive power, namely: ResNet20 (0.27M parameters), ResNet32 (0.46M parameters), SENet-18 (1.23M parameters) and RegNetX_200MF (2.36M parameters)⁷. We set the SGD initial learning rate to 0.1 and the weight decay factor to $5 \cdot 10^{-4}$. ResNet20 was trained for 40 epochs with constant learning rate. ResNet32 was trained for 70 epochs, with learning rate reduced to 0.01 at epoch 50, and to 0.001 at epoch 65. RegNetX_200MF was trained for 150 epochs, with learning rate reduced to 0.01 at epoch 70 and to 0.001 at epoch 120.

Fig. 13(a), Fig. 13(b), Fig. 13(c) and Fig. 13(d) show the CoReS compatibility matrices with ResNet20, ResNet32, SENet-18, and RegNetX_200MF respectively. It can be noticed that both compatibility and verification accuracy follow the expressive power of the network model. ResNet32 in Fig. 13(b) has compatibility similar to Fig. 13(a), but the verification accuracy is higher. RegNetX_200MF improves both compatibility and verification accuracy Fig. 13(d). This experiment also confirms the general applicability of our compatibility learning approach.

Finally, we analyzed the practical case in which a deployed system is upgraded not only with fresh data but also with recent and more powerful network architectures. To this end, we evaluate the case in which the ResNet20 is first upgraded with SENet-18, and subsequently with RegNetX_200MF. Fig. 14 reports the compatibility matrix for this experiment with the CIFAR-100/10 datasets. As can be observed, the learned representations remain compatible as the network architecture is upgraded. The increasing expressive power of the architectures also positively reflects on compatibility.

⁷. Code source for ResNet20 and RegNetX_200MF can be found at <https://github.com/kuangliu/pytorch-cifar> and for ResNet32 at https://github.com/arthurdouillard/incremental_learning_pytorch

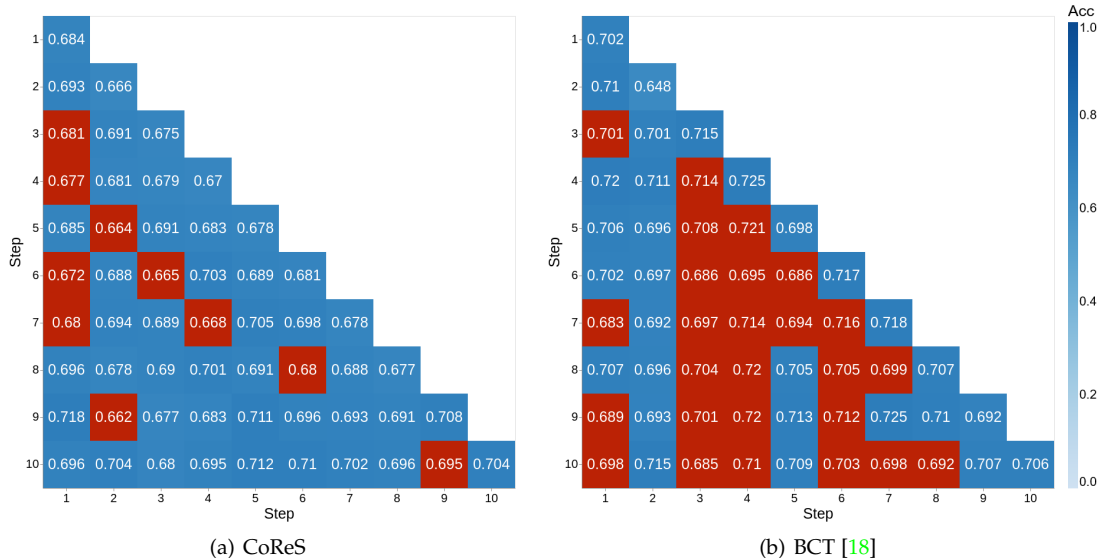


Fig. 12. Compatibility matrices of (a) CoReS and (b) BCT with 9-step upgrading. Models are pre-trained with 50 CIFAR-100 classes and the Model Selection strategy is used.

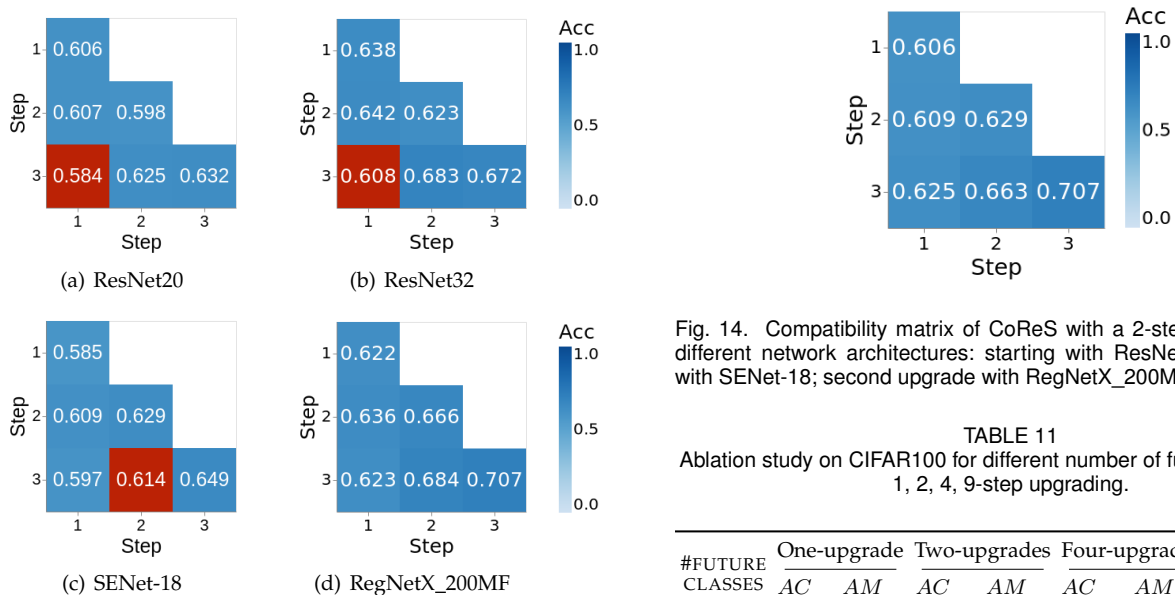


Fig. 13. Compatibility matrices of CoReS for different network architectures with 2-step upgrading: (a) ResNet20; (b) ResNet32; (c) SENet-18; (d) RegNetX_200MF.

6.7 Different Number of Future Classes

It is relevant to study how the number of future classes of Eq. 9 influences the compatibility of a learned representation model. In Tab. 11, we evaluate CoReS with a different number of future classes with one, two, four and nine-step upgrades. As can be noticed from the table, performance does not vary as the number of future classes increases. The *AC* stabilizes as the number of upgrades increases. The *AC* may have sudden variations when the number of upgrades is small as it basically averages the number of times compatibility is achieved.

A related ablation for the classification task in Class-incremental Learning has already been studied in [25]. CiL

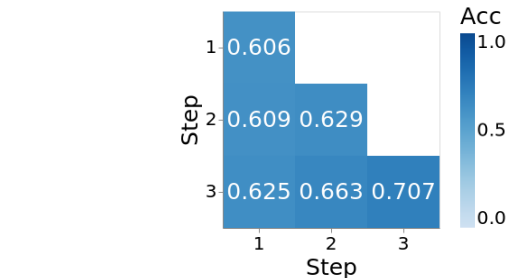


Fig. 14. Compatibility matrix of CoReS with a 2-step upgrading, with different network architectures: starting with ResNet20; first upgrade with SENet-18; second upgrade with RegNetX_200MF.

TABLE 11
Ablation study on CIFAR100 for different number of future classes with 1, 2, 4, 9-step upgrading.

#FUTURE CLASSES	One-upgrade		Two-upgrades		Four-upgrades		Nine-upgrades	
	<i>AC</i>	<i>AM</i>	<i>AC</i>	<i>AM</i>	<i>AC</i>	<i>AM</i>	<i>AC</i>	<i>AM</i>
100	1	0.62	0.67	0.62	0.9	0.59	0.51	0.59
200	1	0.68	1	0.66	0.8	0.62	0.53	0.57
500	1	0.67	1	0.64	0.7	0.64	0.51	0.60
1000	1	0.66	0.67	0.63	0.9	0.62	0.56	0.55
10000	1	0.63	1	0.62	0.7	0.59	0.51	0.56

addresses incremental classification under catastrophic forgetting. Given the substantial difference between the classification and the open-set verification that is investigated in this work, the evaluation of Tab. 11 provides stronger evidence that pre-allocating a large number of future classes in the final classification layer does not reduce or impair the general applicability of the *d*-Simplex fixed classifier [27] used to learn compatible features.

6.8 Different Source of Upgrade Data

In this section we study how performance changes if new data used to upgrade the models belongs to some of the

TABLE 12

Compatibility of CoReS for training with different source of upgrade data with 1-step upgrading: from new classes (NEW CLASSES); from the old classes (OLD CLASSES); from both new and old classes (OLD & NEW CLASSES).

COMPARISON PAIR	SOURCE OF UPGRADE DATA								
	NEW CLASSES			OLD CLASSES			OLD & NEW CLASSES		
	M	ECC	Γ (%)	M	ECC	Γ (%)	M	ECC	Γ (%)
(ϕ_{old}, ϕ_{old})	0.60	-	-	0.64	-	-	0.59	-	-
(ϕ_{new}, ϕ_{old})	0.61	✓	21.3	0.65	✓	33.3	0.64	✓	55.6
(ϕ_{new}, ϕ_{new})	0.65	-	-	0.66	-	-	0.66	-	-

already seen classes. In Tab. 12 “OLD CLASSES” represents this new case in which ϕ_{old} is trained on the 100% of classes with the 50% of samples per class and upgraded using the 100% of training classes with 100% of samples per class; the “NEW CLASSES” column refers to the standard case reported in Sec. 4.1 of the submitted manuscript in which the old model ϕ_{old} is first trained using 50% of classes with 100% of samples per class, while the upgraded model ϕ_{new} using the 100% of classes with 100% of samples; and finally “OLD & NEW CLASSES” indicates the performance of training ϕ_{old} using 50% classes with 50% of samples per class while the upgraded model ϕ_{new} is trained using the 100% of classes with 100% of samples. As can be seen from the table, CoReS achieves a compatible representation between ϕ_{old} and ϕ_{new} in all of the three scenarios, showing that it is effective also in the case in which the upgrade data come from already seen classes. This is substantially due to the fact that the d -Simplex fixed classifier introduced in [27] is less susceptible to the reduced (i.e., imbalanced) number of data samples as recently shown in [41] and in [42].

7 CONCLUSIONS

In this paper, we presented CoReS, a new training procedure that provides compatible representations grounding on the stationarity of the representation as provided by fixed classifiers. In particular, we exploited a special class of fixed classifiers where the classifier weights are fixed to values taken from the coordinate vertices of regular polytopes. They define classes that are maximally separated in the representation space and maintain their spatial configuration stationary as new classes are added. With our solution, there is no need to learn any mappings between representations nor to impose pairwise training with the previously learned model. The functional complexity is fully demanded to the internal layers of the network. We analyzed and discussed the factors that affect the performance of our method and performed extensive comparative experiments on visual search applications of different complexity. We demonstrated that our solution for compatibility learning based on feature stationarity largely outperforms the current state-of-the-art and is particularly effective in the case of multiple upgrades of the training-set, that is the typical case in real applications. CoReS also provides large compatibility when the network architecture is upgraded to newer and more powerful models.

ACKNOWLEDGMENTS

This work was partially supported by the European Commission under European Horizon 2020 Programme, grant number 951911 - AI4Media.

The authors also acknowledge the CINECA award under the IS CRA initiative (IS CRA-C - “ILCoRe”, ID: HP10CRM187), for the availability of high-performance computing resources and thank Giuseppe Fiameni (Nvidia) for his support.

REFERENCES

- [1] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 539–546. IEEE, 2005.
- [2] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- [3] Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 806–813, 2014.
- [4] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [5] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [6] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 507–516, 2016.
- [7] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1701–1708, 2014.
- [8] Yi Sun, Yuheng Chen, Xiaogang Wang, and Xiaoou Tang. Deep learning face representation by joint identification-verification. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [9] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019.
- [10] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Deep metric learning for person re-identification. In *2014 22nd International Conference on Pattern Recognition*, pages 34–39. IEEE, 2014.
- [11] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 152–159, 2014.
- [12] Liang Zheng, Hengheng Zhang, Shaoyan Sun, Manmohan Chandraker, Yi Yang, and Qi Tian. Person re-identification in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1367–1376, 2017.
- [13] Tianlong Chen, Shaojin Ding, Jingyi Xie, Ye Yuan, Wuyang Chen, Yang Yang, Zhou Ren, and Zhangyang Wang. Abd-net: Attentive but diverse person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8351–8361, 2019.
- [14] Artem Babenko, Anton Slesarev, Alexandr Chigorin, and Victor Lempitsky. Neural codes for image retrieval. In *European conference on computer vision*, pages 584–599. Springer, 2014.
- [15] Albert Gordo, Jon Almazán, Jerome Revaud, and Diane Larlus. Deep image retrieval: Learning global representations for image search. In *European conference on computer vision*, pages 241–257. Springer, 2016.

- [16] Giorgos Tolias, Ronan Sifre, and Hervé Jégou. Particular Object Retrieval With Integral Max-Pooling of CNN Activations. In *ICLR 2016 - International Conference on Learning Representations*, International Conference on Learning Representations, pages 1–12, San Juan, Puerto Rico, May 2016.
- [17] Sultan Daud Khan and Habib Ullah. A survey of advances in vision-based vehicle re-identification. *Computer Vision and Image Understanding*, 182:50–63, 2019.
- [18] Yantao Shen, Yuanjun Xiong, Wei Xia, and Stefano Soatto. Towards backward-compatible representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6368–6377, 2020.
- [19] Richard Van Noorden. The ethical questions that haunt facial-recognition research. *Nature*, 587(7834):354–358, 2020.
- [20] Rahul Duggal, Hao Zhou, Shuo Yang, Yuanjun Xiong, Wei Xia, Zhuowen Tu, and Stefano Soatto. Compatibility-aware heterogeneous visual search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10723–10732, 2021.
- [21] Ken Chen, Yichao Wu, Haoyu Qin, Ding Liang, Xuebo Liu, and Junjie Yan. R3 adversarial network for cross model face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9868–9876, 2019.
- [22] Jie Hu, Rongrong Ji, Hong Liu, Shengchuan Zhang, Cheng Deng, and Qi Tian. Towards visual feature translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3004–3013, 2019.
- [23] Chien-Yi Wang, Ya-Liang Chang, Shang-Ta Yang, Dong Chen, and Shang-Hong Lai. Unified representation learning for cross model compatibility. In *31st British Machine Vision Conference 2020, BMVC 2020, Virtual Event, UK, September 7-10, 2020*. BMVA Press, 2020.
- [24] Qiang Meng, Chixiang Zhang, Xiaoqiang Xu, and Feng Zhou. Learning compatible embeddings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9939–9948, October 2021.
- [25] Federico Pernici, Matteo Bruni, Claudio Baccchi, Francesco Turchini, and Alberto Del Bimbo. Class-incremental learning with pre-allocated fixed classifiers. In *25th International Conference on Pattern Recognition, ICPR 2020, Milan, Italy, January 10-15, 2021*. IEEE Computer Society, 2020.
- [26] Federico Pernici, Matteo Bruni, Claudio Baccchi, and Alberto Del Bimbo. Maximally compact and separated features with regular polytope networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 46–53, June 2019.
- [27] F. Pernici, M. Bruni, C. Baccchi, and A. D. Bimbo. Regular polytope networks. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–15, 2021.
- [28] Gagan Bansal, Besmira Nushi, Ece Kamar, Daniel S Weld, Walter S Lasecki, and Eric Horvitz. Updates in human-ai teams: Understanding and addressing the performance/compatibility tradeoff. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 2429–2437, 2019.
- [29] Megha Srivastava, Besmira Nushi, Ece Kamar, Shital Shah, and Eric Horvitz. An empirical analysis of backward compatibility in machine learning systems. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 3272–3280, 2020.
- [30] Sijie Yan, Yuanjun Xiong, Kaustav Kundu, Shuo Yang, Siqi Deng, Meng Wang, Wei Xia, and Stefano Soatto. Positive-congruent training: Towards regression-free model updates. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14299–14308, 2021.
- [31] F. Träuble, J. von Kügelgen, M. Kleindessner, F. Locatello, B. Schölkopf, and P. Gehler. Backward-compatible prediction updates: A probabilistic approach. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*, December 2021.
- [32] Michael Gygli, Jasper Uijlings, and Vittorio Ferrari. Towards reusable network components by learning compatible representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7620–7629, 2021.
- [33] Mateusz Budnik and Yannis Avrithis. Asymmetric metric learning for knowledge transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8228–8238, 2021.
- [34] Yixuan Li, Jason Yosinski, Jeff Clune, Hod Lipson, and John Hopcroft. Convergent learning: Do different neural networks learn the same representations? In *Feature Extraction: Modern Questions and Challenges*, pages 196–212. PMLR, 2015.
- [35] Liwei Wang, Lunjia Hu, Jiayuan Gu, Yue Wu, Zhiqiang Hu, Kun He, and John Hopcroft. Towards understanding learning representations: To what extent do different neural networks learn the same representation. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, page 9607–9616, Red Hook, NY, USA, 2018. Curran Associates Inc.
- [36] Vardan Pappayan, XY Han, and David L Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020.
- [37] Vignesh Kothapalli, Ebrahim Rasromani, and Vasudev Awatramani. Neural collapse: A review on modelling principles and generalization. *arXiv preprint arXiv:2206.04041*, 2022.
- [38] Florian Graf, Christoph Hofer, Marc Niethammer, and Roland Kwitt. Dissecting supervised contrastive learning. In *International Conference on Machine Learning*, pages 3821–3830. PMLR, 2021.
- [39] Zhihui Zhu, Tianyu Ding, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features. *Advances in Neural Information Processing Systems*, 34:29820–29834, 2021.
- [40] Cong Fang, Hangfeng He, Qi Long, and Weijie J Su. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *Proceedings of the National Academy of Sciences*, 118(43):e2103091118, 2021.
- [41] Yibo Yang, Shixiang Chen, Xiangtai Li, Liang Xie, Zhouchen Lin, and Dacheng Tao. Inducing neural collapse in imbalanced learning: Do we really need a learnable classifier at the end of deep neural network? In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 37991–38002. Curran Associates, Inc., 2022.
- [42] Tejaswi Kasarla, Gertjan Burghouts, Max van Spengler, Elise van der Pol, Rita Cucchiara, and Pascal Mettes. Maximum class separation as inductive bias in one matrix. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 19553–19566. Curran Associates, Inc., 2022.
- [43] Tobias Weyand, Andre Araujo, Bingyi Cao, and Jack Sim. Google landmarks dataset v2-a large-scale benchmark for instance-level recognition and retrieval. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2575–2584, 2020.
- [44] Nikolaos-Antonios Ypsilantis, Noa Garcia, Guangxing Han, Sarah Ibrahim, Nanne Van Noord, and Giorgos Tolias. The met dataset: Instance-level recognition for artworks. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [45] Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira. Re-evaluating continual learning scenarios: A categorization and case for strong baselines. In *NeurIPS Continual Learning Workshop*, 2018.
- [46] Guido M Van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019.
- [47] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost van de Weijer. Class-incremental learning: survey and performance evaluation. *arXiv preprint arXiv:2010.15277*, 2020.
- [48] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989.
- [49] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *European Conference on Computer Vision*, pages 499–515. Springer, 2016.
- [50] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks. In Maria-Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 507–516. JMLR.org, 2016.
- [51] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphereface: Deep hypersphere embedding for face recognition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 6738–6746. IEEE Computer Society, 2017.
- [52] Moritz Hardt and Tengyu Ma. Identity matters in deep learning. In *5th International Conference on Learning Representations, ICLR*

2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings, 2017.

- [53] Elad Hoffer, Itay Hubara, and Daniel Soudry. Fix your classifier: the marginal value of training the last weight layer. In *International Conference on Learning Representations*, 2018.
- [54] Joaquin Quiñero-Candela, Masashi Sugiyama, Neil D Lawrence, and Anton Schwaighofer. *Dataset shift in machine learning*. Mit Press, 2009.
- [55] Masashi Sugiyama, Matthias Krauledat, and Klaus-Robert Müller. Covariate shift adaptation by importance weighted cross validation. *Journal of Machine Learning Research*, 8(5), 2007.
- [56] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.
- [57] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [58] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014.
- [59] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007.
- [60] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE international conference on computer vision*, pages 1116–1124, 2015.
- [61] Jie Hu, Li Shen, Samuel Albanie, Gang Sun, and Enhua Wu. Squeeze-and-excitation networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(8):2011–2023, 2020.
- [62] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [63] Kaiyang Zhou and Tao Xiang. Torchreid: A library for deep learning person re-identification in pytorch. *arXiv preprint arXiv:1910.10093*, 2019.
- [64] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [65] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547, 2019.
- [66] Yujun Lin, Song Han, Huizi Mao, Yu Wang, and William J Dally. Deep Gradient Compression: Reducing the communication bandwidth for distributed training. In *The International Conference on Learning Representations*, 2018.
- [67] Li Jing, Pascal Vincent, Yann LeCun, and Yuandong Tian. Understanding dimensional collapse in contrastive self-supervised learning. In *International Conference on Learning Representations*, 2021.
- [68] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [69] Vaishnavh Nagarajan and J. Zico Kolter. Uniform convergence may be unable to explain generalization in deep learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [70] Behnam Neyshabur, Hanie Sedghi, and Chiyuan Zhang. What is being transferred in transfer learning? In *NeurIPS*, 2020.
- [71] Sepp Hochreiter and Jürgen Schmidhuber. Simplifying neural nets by discovering flat minima. In *Advances in neural information processing systems*, pages 529–536, 1995.
- [72] Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural computation*, 9(1):1–42, 1997.



Niccolò Biondi Niccolò Biondi received the M.S. degree (cum laude) in Computer Engineering from the University of Firenze, Italy, in 2021. Presently, he is a Ph.D. student at the University of Firenze at MICC, Media Integration and Communication Center, University of Firenze. His research interests include machine learning and computer vision with special focus on compatible learning, representation learning, and incremental learning.



Federico Pernici Federico Pernici is Assistant Professor at the University of Firenze, Italy. He received the laurea degree in Information Engineering in 2002, the post-laurea degree in Internet Engineering in 2003 and the Ph.D. in Information and Telecommunication Engineering in 2005 from the University of Firenze, Italy. He has been working at the MICC Media Integration and Communication Center as a research assistant since 2002, and has also served as adjunct professor at the University of Firenze.

His scientific interests are computer vision and machine learning with a focus on different aspects of visual tracking, incremental learning and representation learning. He was Associate Editor of Machine Vision and Applications journal.



Matteo Bruni Matteo Bruni received the M.S. degree (cum laude) in Computer Engineering from the University of Firenze, Italy, in 2016, the Ph.D. in Information Engineering at the University of Firenze in 2020. His research interests include pattern recognition and computer vision with specific focus on feature embedding, face recognition, incremental learning and compatible representation.



Alberto Del Bimbo Prof. Del Bimbo is Full Professor at the University of Firenze, Italy. He is the author of over 350 scientific publications in computer vision and multimedia and principal investigator of technology transfer projects with industry and governments. He was the Program Chair of ICPR 2012, ICPR 2016 and ACM Multimedia 2008, and the General Chair of IEEE ICMCS 1999, ACM Multimedia 2010, ICMR 2011, ECCV 2012 and ICPR2020. He is the General Chair of ACM Multimedia 2021. He is the Editor in Chief of ACM TOMM Transactions on Multimedia Computing Communications and Applications and Associate Editor of Multimedia Tools and Applications and Pattern Analysis and Applications journals. He was Associate Editor of IEEE Transactions on Pattern Analysis and Machine Intelligence, IEEE Transactions on Multimedia and Pattern Recognition and also served as the Guest Editor of many Special Issues in highly ranked journals. Prof. Del Bimbo is the Chair of ACM SIGMM the Special Interest Group on Multimedia. He is IAPR Fellow and ACM Distinguished Scientist and is the recipient of the 2016 ACM SIGMM Award for *Outstanding Technical Contributions to Multimedia Computing Communications and Applications*.