

Forward Compatible Training for Large-Scale Embedding Retrieval Systems

Vivek Ramanujan*
University of Washington†

Pavan Kumar Anasosalu Vasu
Apple

Ali Farhadi
Apple

Oncel Tuzel
Apple

Hadi Pouransari*
Apple

Abstract

In visual retrieval systems, updating the embedding model requires recomputing features for every piece of data. This expensive process is referred to as backfilling. Recently, the idea of backward compatible training (BCT) was proposed. To avoid the cost of backfilling, BCT modifies training of the new model to make its representations compatible with those of the old model. However, BCT can significantly hinder the performance of the new model. In this work, we propose a new learning paradigm for representation learning: forward compatible training (FCT). In FCT, when the old model is trained, we also prepare for a future unknown version of the model. We propose learning side-information, an auxiliary feature for each sample which facilitates future updates of the model. To develop a powerful and flexible framework for model compatibility, we combine side-information with a forward transformation from old to new embeddings. Training of the new model is not modified, hence, its accuracy is not degraded. We demonstrate significant retrieval accuracy improvement compared to BCT for various datasets: ImageNet-1k (+18.1%), Places-365 (+5.4%), and VGG-Face2 (+8.3%). FCT obtains model compatibility when the new and old models are trained across different datasets, losses, and architectures.¹

1. Introduction

Modern representation learning systems for vision use deep neural networks to embed very high-dimensional images into low-dimensional subspaces. These embeddings are multi-purpose representations that can be used for downstream tasks such as recognition, retrieval, and detection. In large systems in the wild, images are constantly being added in a highly distributed manner. Updat-

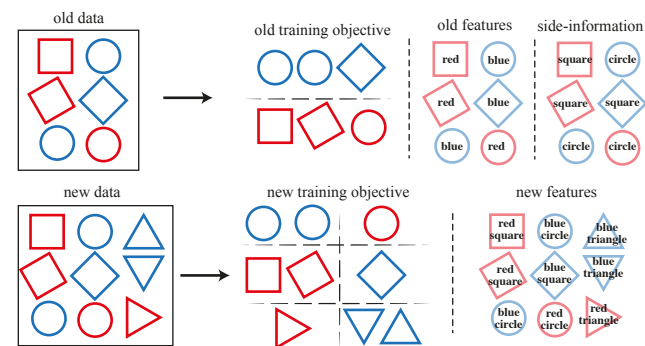


Figure 1. The old training objective (**top**) requires discrimination between *colors* of objects. The new training objective (**bottom**) further requires discrimination between *shapes* of objects. A model which learns the minimum description length feature (*color top middle*) of the old training objective will be unable to complete the new training objective. This demonstrates how side-information (**top right**) stores information which can be useful for future training objectives but not necessarily the present objective.

ing the models in these kinds of systems is challenging. The new model architecture and training objective may be completely disjoint from that of the old model, and therefore its embeddings are *incompatible* with downstream tasks reliant on the old embeddings.

The naive solution to feature incompatibility, recomputing embeddings (*backfilling*) for every image, is expensive. To avoid the backfilling cost, [44] proposed Backward Compatible Training (BCT), an algorithm for training of the new model to ensure its embeddings are compatible with those of the old model. BCT succeeds in maintaining retrieval accuracy in the no-backfilling case. However, we show that BCT struggles to achieve the same performance gain as independent training for the new model (Section 4.3). Further, it can potentially carry unwanted biases from the old to the new model [1].

Perfect direct compatibility between new (ϕ_{new}) and old (ϕ_{old}) embedding models runs counter to the goal of learning better features. Therefore, instead of trying to make

*Corr. ramanv@cs.washington.edu & mpouransari@apple.com.

†Work completed during internship at Apple.

¹Code available at <https://github.com/apple/ml-fct>.

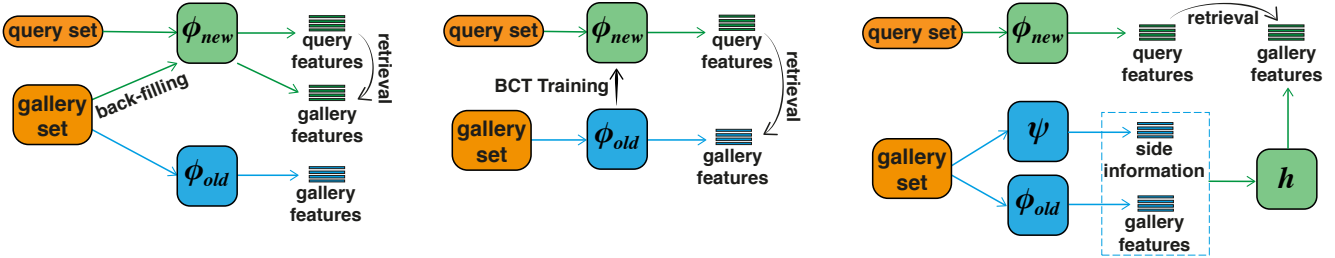


Figure 2. Independent training with back-filling (left), backward compatible training (middle), and forward compatible training (FCT) (right). In FCT, we store side-information along with gallery features to prepare for a future embedding model.

ϕ_{new} directly compatible with ϕ_{old} , as in [44], we learn a transformation to map from old to new embeddings similar to [30, 47]. If we are able to find a perfect transformation, the compatibility problem is solved because we can convert old to new embeddings without requiring access to the original images (no-backfilling). However, ϕ_{old} possibly discards information about the image that ϕ_{new} does not. This information mismatch between the embeddings makes it difficult to learn a perfect transformation.

In this context, we introduce the idea of *forward-compatible training* (FCT). This notion is borrowed from software-engineering, where software is developed with the assumption that it will be updated, and therefore is made easy to update. Here, we adopt a similar strategy for neural networks. When training the old model we know there will be an update in future, therefore we *prepare* to be compatible with some unknown successor model. Future model can vary in training dataset, objective, and architecture. We propose the concept of *side-information*: auxiliary information learned at the same time as the old model which facilitates updating embeddings in the future. Intuitively, *side-information* captures features of the data which are not necessarily important for the training objective of the old model but are potentially important for the future model.

In Figure 1 we show a toy example to demonstrate the concept. The old training objective is to classify red objects versus blue ones. The old model therefore learns *color* as the feature for each object, as this is the minimum discriminative attribute for this objective [46]. Later, additional data is added (triangles), and the new training objective is to classify objects based on both their shape and color. Using only what the old embeddings encode (*color*), we cannot distinguish between different shapes which have the same color, e.g. blue circles and blue squares. In FCT, we store *side-information* with the old embeddings to aid future updates. Here, *shape* is the perfect *side-information* to store. Note that the *shape side-information* in this scenario does not help with the old training objective, but is useful for future. Learning good *side-information* is a challenging and open-ended problem since we are not aware of the future model or new training objectives. We present our results on different possibilities for side information in Section 5.1.

To use *side-information* for model compatibility, we construct a transformation function h , which maps from pairs of old embeddings and *side-information* to new embeddings. h is trained to mix the auxiliary information provided in *side-information* with old embeddings to reproduce the new embedding. See Figure 2 for a schematic of our setup. Using FCT we show significant improvements on backward-compatibility metrics (Section 4). Since ϕ_{new} and ϕ_{old} are trained independently, ϕ_{new} performance remains unaffected by enforced compatibility. This is in contrast to prior works [44, 47], which have focused on making ϕ_{new} directly compatible with ϕ_{old} .

We note that FCT requires transforming all old embeddings, which BCT avoids. This increased computational cost could be seen as a drawback of the method. However, embeddings in general have much lower dimensionality than images. In Figure 3 we show the trade-off between computational cost (per example) and accuracy for different strategies in the ImageNet setup (with image size $3 \times 224 \times 224$) as described in Section 4. FCT does not affect new model performance, and obtains significantly higher backward accuracy for a small additional computation and storage. Note that backfilling cost (both computation and storage) scales with image resolution, but FCT transformation cost only depends on the embedding dimension. FCT is particularly effective in the paradigm where computations take place privately on-device, trivializing its cost. We elaborate this point further in Section K.

Our contributions are as follows:

1. We propose a new learning paradigm, forward compatible training (FCT), where we explicitly prepare for future model updates. Our goal is to be compatible with future models.
2. In the context of FCT for representation learning we propose *side-information*: an auxiliary features learned at the same time as ϕ_{old} which aids transfer to ϕ_{new} . Intuitively, *side-information* captures task-independent features of the data which are not necessarily important for the training objective ϕ_{old} is trained on but are likely to be important for ϕ_{new} .
3. We demonstrate substantial retrieval accuracy improvement compared to BCT for various datasets:

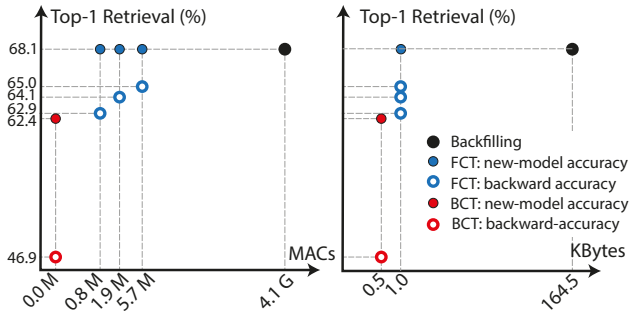


Figure 3. A comparison of compatibility methods’ accuracies and associated costs. We show **(left)** Number of multiply–accumulate operations (MACs) and **(right)** data storage requirement in KBytes. Transformation models of different capacity produce different FCT costs and accuracies. We added 0 to the logarithmic x -axis **(left)** for visual comparison.

ImageNet-1k (+18.1%), Places-365 (+5.4%), and VGG-Face2 (+8.3%). We show FCT outperforms the BCT paradigm by a large margin when new and old models are trained across different datasets, losses, and architectures. Unlike prior works [30, 44, 47], models in FCT are trained independently, and hence their accuracies are not compromised for the purpose of compatibility.

2. Related Works

Model Compatibility Our closest point of comparison is [44], which presents the problem of “backwards-compatibility”. They posit that future embeddings models should be compatible with old models when used in a retrieval setting. They present BCT, an algorithm that allows new embedding models to be compatible with old models through a joint training procedure involving a distillation loss. Other works [8, 30, 47] attempt to construct a unified representation space on which models are compatible. These procedures also modify training of individual models to ensure that they are easy to transform to this unified embedding space. This is in contrast to our work, which assumes the old and new models are trained independently.

[26] studies the relationship and degree of compatibility between two models trained on the same dataset. [49] discusses the transferability of features for the same model. These are both empirical works primarily studying deep learning phenomena.

Side-information The idea of side-information is commonly discussed in the zero-shot learning literature with a different context and purpose [9]. In zero-shot learning, side-information is acquired separate from pretraining datasets, and is used to provide information about unseen classes [13, 31, 40]. The recent CLIP [37] model and others [13] for example use a language model as side-information to initialize their classifiers for zero-shot inference of un-

seen classes. Our usage of *side-information* differs quite significantly, in that it is used per-example and our transfer to future tasks is not zero-shot. Further, our retrieval setup precludes the usage of side-information to define new categories.

Transfer and cross-domain learning Transfer learning as a field is quite varied. Methods are variously classified under few-shot [25, 45], continual learning [35, 48], and life-long learning [28, 39]. In general, transfer learning methods seek to use knowledge learned from one domain in another to improve performance [34]. The assumption here is that transfer learning domains are related: one can exploit knowledge from one domain to aid in another domain. We have no such expectation when training old and new models. The goal of our feature transformation is not solve a new domain, but make two existing embedding models compatible, which is not a goal of transfer or cross-domain learning. In particular, even if a model can learn a new and old objective simultaneously, there is no guarantee that this model will be compatible with a model which has only learned the old objective (see Section 4.3 for details). Current practice for transfer learning largely centers around various fine-tuning schemes [24], which makes an assumption that the architecture used in the old domain and new domain are the same. We have no constraint on old and new architectures.

3. Method

3.1. Problem Setup

A gallery set \mathcal{G} is a collection of images, which are grouped into different clusters, $\{y_1 \dots y_n\}$. In visual retrieval, given a query image of a particular class, the goal is to retrieve images from \mathcal{G} with the same class. A set of such query images is called the query set \mathcal{Q} .

In embedding based retrieval, we have an embedding model $\phi : \mathbb{R}^D \rightarrow \mathbb{R}^d$ where $d \ll D$ and D is the dimensionality of the input image, trained offline on some dataset \mathcal{D} which is disjoint from \mathcal{Q} and \mathcal{G} . Then, using some distance function $\mathcal{D} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, we get the closest possible image to query image $x_q \in \mathcal{Q}$ as $\arg \min_{x \in \mathcal{G}} \mathcal{D}(\phi(x_q), \phi(x))$. We use L2-distance for \mathcal{D} in our case.

In our particular setup, we have two embedding models: the old model $\phi_{old} : \mathbb{R}^D \rightarrow \mathbb{R}^{d_{old}}$ and the new model $\phi_{new} : \mathbb{R}^D \rightarrow \mathbb{R}^{d_{new}}$, where $d_{old}, d_{new} \ll D$ are the embedding dimensions of old and new models, respectively. ϕ_{old} and ϕ_{new} are trained with datasets \mathcal{D}_{old} and \mathcal{D}_{new} , respectively, using some supervised loss function.

We assume ϕ_{old} is applied to every $x \in \mathcal{G}$, generating a collection of gallery embeddings, and ϕ_{new} is applied to every $x_q \in \mathcal{Q}$, generating a set of query embeddings. Our goal is to design a method for performing retrieval between these

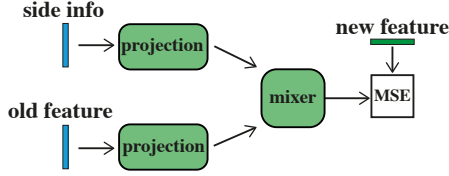


Figure 4. FCT transformation model architecture and training. We use an MLP with BatchNorm [20] for the projection layers and the mixer. The exact layer sizes are reported in Appendix D.

two sets of embeddings without directly using the images in \mathcal{G} . We quantify the model compatibility performance as the accuracy of this retrieval.

3.2. Forward Compatibility Setup

In FCT, we define a function $\psi : \mathbb{R}^D \rightarrow \mathbb{R}^{d_{side}}$, which takes in an input from the \mathcal{G} and produces our *side-information*. Along with every embedding $\phi_{old}(x)$ for $x \in \mathcal{G}$ we also store the corresponding *side-information* $\psi(x)$. Finally, we have a transformation model $h : \mathbb{R}^{d_{old}} \times \mathbb{R}^{d_{side}} \rightarrow \mathbb{R}^{d_{new}}$ which maps from ϕ_{old} and ψ to ϕ_{new} . In the FCT setup, we assume ψ and h are trained using \mathcal{D}_{old} and \mathcal{D}_{new} training sets, respectively. To perform retrieval for $x_q \in \mathcal{Q}$ using ψ and h , we take $\arg \min_{x \in \mathcal{G}} \mathfrak{D}(\phi_{new}(x_q), h(\phi_{old}(x), \psi(x)))$.

FCT has a few key properties and constraints:

1. We do not modify the training of the new model, so it gets the highest possible accuracy.
2. We make the old model representation compatible to that of the new model through a learned transformation.
3. As we discussed in the example of Figure 1, direct transformation from old embeddings to new may not be possible. When we train the old model, we *prepare* for a future update by storing *side-information* for each example in \mathcal{G} .

Efficiency FCT comes with additional costs compared to BCT in computing side-information and updating. However, with a decentralized system design, we can mitigate these costs. See Appendix K for more details on this decentralized design. Even without such a system, the computational cost of FCT is orders of magnitude lower than that of full backfilling (see Figure 3).

3.3. Training FCT Transformation

For each sample $x \in \mathcal{D}_{new}$, we compute the transformation $h(\phi_{old}(x), \psi(x))$. Our objective to minimize is $\ell(x) = \|h(\phi_{old}(x), \psi(x)) - \phi_{new}\|_2^2$. In our case, we write h as h_θ , a neural network parameterized by θ . Our final optimization problem can be written as $\theta^* = \arg \min_{\theta} \sum_{x \in \mathcal{D}_{new}} \ell(x)$. The transformation model consists of two projection layers corresponding to old embedding and side-information branches followed by a mixing layer to reconstruct the new embedding (Figure 4). We tested several alternative objectives, including KL divergence as in [18], but found L2 loss

to perform the best, in contrast to [44], which finds that using L2 loss does not allow for backwards compatibility. Recently, [4] also observed superior performance using L2 loss between features for knowledge distillation. We provide training details in Section D.

3.4. Training FCT Side-Information

The ideal side-information encodes compressed features for each example with which we can reconstruct new features given the old ones. This is a challenging task since we do not know about the future embedding model when learning the *side-information*. The old model mainly learns features that are useful for its own training objective, and is (ideally) invariant to extraneous information. *Side-information*, on the other hand, should be task independent, and capture features complementary to the old embeddings, that are useful for possible future updates.

Learning general features has been extensively studied in the context of unsupervised and contrastive learning [3, 9, 12, 15, 32, 51] when labels are not available. Here, our objective is to learn task-independent features that are not necessarily relevant for the old task, and hence the labels. This makes unsupervised learning method a particularly interesting choice to learn *side-information* for FCT. We use SimCLR contrastive learning [9] to train ψ . As a constraint of our setup, we can only train ψ with data available at old training. We study other choices of side-information for FCT in Section 5.1. We provide details of the improved SimCLR training in the Appendix.

4. Experiments

In our experimentation, we consider different model update setups by changing training dataset, gallery/query sets, architecture, and loss. We show FCT consistently results in high model compatibility while not affecting accuracy of the new model.

4.1. Evaluation Metrics

Cumulative Matching Characteristics (CMC) corresponds to top- k accuracy. A similarity ordering is produced using the query embedding and every embedding in the gallery set, sorted by lowest L2 distance. If an image with the same class or identity (for face retrieval) appears in the top- k retrievals, this is recorded as correct. We report CMC top-1 and 5 percentages for all models.

Mean Average Precision (mAP) is a standard metric which summarizes precision and recall metrics by taking the area under the precision-recall curve. We compute mAP@1.0, which is the average precision over recall values in [0.0, 1.0].

Notation Following [44], we use the notation Query Embedding / Gallery Embedding to denote which model we are

using for query and gallery embeddings, respectively. For example ϕ_{new}/ϕ_{old} refers to using the new model for the query embedding and the old model for the gallery embedding. The evaluation task (retrieval from the gallery set) is fixed, while the old and new embedding models are trained differently.

All FCT results are *backwards-compatible* with the definition provided in [44] on CMC and mAP metrics.

4.2. Datasets

ImageNet-1k [41] is a large-scale image recognition dataset used in the ILSVRC 2012 challenge. It has 1000 image classes. Its train set is approximately class balanced with $\sim 1.2k$ images per class, and its validation set is exactly class balanced with 500 images per class. We call a subset of the ImageNet-1k training set consisting of images of the first 500 classes, ImageNet-500. This subset is used for training of the old model. This is a biased split: the second 500 classes are generally more difficult (e.g. they contain all the dog breeds). For retrieval, we utilize full ImageNet-1k validation set for both the query set and gallery set. When we query with an image, we remove it from the validation set to fairly compute retrieval accuracy.

VGGFace2 [5] is a large-scale face recognition dataset of 3.31 million images of 9131 subjects. It is split into a train set of 8631 subjects and a validation set of 500 (disjoint) subjects. On average, there are 362.6 images per subject in the train set and exactly 500 images per subject in the validation set. The old train set is constructed from the first 10% of subjects. We call this train set VGGFace2-863 and the full train set as VGGFace2-8631. For the validation, we generate a (fixed) random subsample of 50 images per subject in the validation set to perform retrieval on, as in [5].

Places-365 [52] is a large-scale scene recognition dataset with ~ 1.8 million images of 365 scene categories with between 3068 and 5000 images per category. We use the first 182 classes as the old train set, which we refer to as Places-182. The validation set for Places-365 contains 36500 images, 100 images per scene category. We use the full validation set for both the query and gallery set, as with ImageNet.

4.3. Old to New Update Experiments

By convention, we use the notation (Architecture)-(Embedding Dimension). We choose the embedding dimension with the highest accuracy on the new model. Our ablation on this is provided in the Appendix. Note that (ϕ_{new}/ϕ_{new}) is the upper bound and (ϕ_{old}/ϕ_{old}) is the lower bound for model compatibility. We further clarify that in each set of experiments, the evaluation criterion is the same regardless of training objective: retrieval on a fixed gallery set.

ImageNet-500 to ImageNet-1k In this experiment, we

train ϕ_{old} (ResNet50-128) on ImageNet-500 and ϕ_{new} (ResNet50-128) on ImageNet-1k using the Softmax cross-entropy loss. The side-information model ψ for FCT is also a ResNet50-128 trained on the ImageNet-500 dataset with SimCLR [9]. For retrieval evaluation, we use the ImageNet-1k validation set as both the query and gallery sets. Results are provided in Table 1.

FCT obtains substantial improvement over BCT. First, FCT, by construction, does not affect accuracy of the new model (68.1% top-1, which is the same as independent training). This is opposed to 5.7% top-1 accuracy drop for BCT. When comparing the CMC top-1 in the compatibility setup FCT outperforms BCT even more significantly: 46.9% to 65.0%, +18.1% improvement in CMC top-1. In fact, FCTs’ top-1 compatibility accuracy (65.0%) is more than the upper bound accuracy in BCT (62.4%) by 2.6%. For reference, in the case where we use no side-info, where $\psi(x) = 0$, we get a drop of top-1 compatibility accuracy to 61.8%, showing the importance of side-information. See Section 5.1 for a more detailed ablation on the affects of side-information.

We also report accuracy of the transformed old model, $h(\phi_{old}, \psi)/h(\phi_{old}, \psi)$. We see significant improvement over the independently trained old model: 46.5% to 59.3%. This shows that a model update with FCT, not only results in model compatibility, but also significantly improves quality of the old features through the transformation. This has significant consequences for practical use-cases such as embedding based clustering. We report comparisons to [30] and [47], modified for our setting, in the Appendix.

Places-182 to Places-365 This experiment uses the same setup as ImageNet-500 to ImageNet-1k, except with training the old model (ResNet50-512) on Places-182 and the new model (ResNet50-512) on Places-365. The SimCLR side-information used for these experiments is trained on ImageNet-1k, showing that side-information can be transferred between domains. For retrieval evaluation, we use the Places-365 validation set as both the query and gallery sets. These results are shown in Table 2. Our compatibility performance $(\phi_{new}/h(\phi_{old}, \psi))$ yields a 5.4% improvement in CMC top-1 over BCT, even outdoing its upper bound by 0.9% top-1.

VGGFace2-863 to VGGFace2-8631 In this experiment, we train the old model (ResNet18-128) on VGGFace2-863 and the new model (ResNet50-128) on VGGFace2-8631 using the ArcFace objective [11]. The exact hyperparameters for these experiments is provided in the Appendix. For retrieval evaluation, we use the VGGFace2 validation set. Following [5], we randomly sample 50 images for each subject and use this as both the query and gallery sets. In this experiment we use the *alternate old model* side-information, another training run of the old model, differing only in SGD randomness and initialization (Section 5.1).

Method	Case	CMC		mAP@1.0
		top-1	top-5 %	
Independent	ϕ_{old}/ϕ_{old}	46.5	64.6	29.9
	ϕ_{new}/ϕ_{old}	0.1	0.5	0.003
	ϕ_{new}/ϕ_{new}	68.1	84.4	45.0
BCT [44]	$\phi_{new}^{BCT}/\phi_{old}$	46.9	65.4	30.1
	$\phi_{new}^{BCT}/\phi_{new}^{BCT}$	62.4	81.9	41.1
FCT (ours)	$h(\phi_{old}, \psi)/h(\phi_{old}, \psi)$	59.3	76.4	41.3
	$\phi_{new}/h(\phi_{old}, \psi)$	65.0	82.3	43.6
	ϕ_{new}/ϕ_{new}	68.1	84.4	45.0

Table 1. Comparison of different compatible training methods. The old model is trained on the ImageNet-500, and new model on the ImageNet-1k. The ImageNet validation set is used as gallery and query sets. Both models are ResNet50-128.

Method	Case	CMC		mAP@1.0
		top-1	top-5 %	
Independent	ϕ_{old}/ϕ_{old}	29.6	58.2	11.6
	ϕ_{new}/ϕ_{old}	0.3	1.5	0.12
	ϕ_{new}/ϕ_{new}	37.0	65.1	17.0
BCT [44]	$\phi_{new}^{BCT}/\phi_{old}$	30.4	58.7	12.6
	$\phi_{new}^{BCT}/\phi_{new}^{BCT}$	34.9	64.2	16.0
FCT (ours)	$h(\phi_{old}, \psi)/h(\phi_{old}, \psi)$	34.0	62.3	17.3
	$\phi_{new}/h(\phi_{old}, \psi)$	35.8	64.5	18.1
	ϕ_{new}/ϕ_{new}	37.0	65.1	17.0

Table 2. Comparison of different compatible training methods. The old model is trained on Places-182, and new model on Places-365. The Places-365 validation set is used as gallery and query sets. Both models are ResNet50-512.

ImageNet-1k SimCLR side-information provided marginal improvement over no side information (91.6% to 92.0% top-1). This makes sense because ImageNet as a domain is very distant from cropped faces. These results are shown in Table 3 and are consistent with the results on other datasets. Our compatibility performance yields a 8.3% improvement in CMC top-1 over BCT. This shows our algorithm generalizes to models trained with objectives other than softmax cross-entropy.

ImageNet-500 to ImageNet-1k w/ changing architecture

This experiment has the same setup as ImageNet-500 to ImageNet-1k, except we consider model architecture is also changed during the update. Changing architecture when updating a model is frequent practice. We show FCT results in Table 4 on architectures of varying depths, structure, and embedding dimension compared to the new model. We see that as old model performance drops, compatibility performance also drops. However, in all cases compatibility performance remains quite high, still outperforming the BCT upper bound on ResNet50.

ImageNet-500 to ImageNet-1k evaluated on Places-365

In this experiment we evaluate the same embedding models as in the ImageNet-500 to ImageNet-1k experiment using Places-365 validation set as gallery and query sets. This is a

Method	Case	CMC		mAP@1.0
		top-1	top-5 %	
Independent	ϕ_{old}/ϕ_{old}	84.0	93.5	44.6
	ϕ_{new}/ϕ_{old}	0.2	1.0	0.2
	ϕ_{new}/ϕ_{new}	96.6	98.4	77.8
BCT [44]	$\phi_{new}^{BCT}/\phi_{old}$	84.2	93.9	45.0
	$\phi_{new}^{BCT}/\phi_{new}^{BCT}$	95.1	98.3	74.3
FCT (ours)	$h(\phi_{old}, \psi)/h(\phi_{old}, \psi)$	87.2	94.9	53.4
	$\phi_{new}/h(\phi_{old}, \psi)$	92.5	97.5	62.6
	ϕ_{new}/ϕ_{new}	96.6	98.4	77.8

Table 3. Comparison of different compatible training methods on VGGFace2. The old model is trained on VGGFace2-863, and the new model is trained on VGGFace2-8631 with the ArcFace training loss. The architecture is ResNet18-128 for the old model and ResNet50-128 for the new model.

Old Architecture	Case	CMC		mAP@1.0
		top-1	top-5 %	
N/A	ϕ_{new}/ϕ_{new}	68.1	84.4	45.0
ResNet18-128 [17]	ϕ_{old}/ϕ_{old}	40.1	60.1	19.9
	$h(\phi_{old}, \psi)/h(\phi_{old}, \psi)$	57.0	75.0	39.4
	$\phi_{new}/h(\phi_{old}, \psi)$	64.2	82.1	42.6
MobileNet-128 [19]	ϕ_{old}/ϕ_{old}	42.3	62.7	20.3
	$h(\phi_{old}, \psi)/h(\phi_{old}, \psi)$	57.6	75.7	39.8
	$\phi_{new}/h(\phi_{old}, \psi)$	64.3	82.1	42.9
ResNet50-64 [17]	ϕ_{old}/ϕ_{old}	44.7	62.7	30.6
	$h(\phi_{old}, \psi)/h(\phi_{old}, \psi)$	56.6	74.1	39.0
	$\phi_{new}/h(\phi_{old}, \psi)$	63.6	81.9	42.0

Table 4. FCT compatibility results when old and new architectures are different. For all cases, the new model is ResNet50-128. The old model is trained on the ImageNet-500, and new model on the ImageNet-1k. The ImageNet validation set is used as gallery and query sets.

challenging case since Places365 is a different domain from ImageNet. This is clear from the drop in old and new model retrieval performance from ImageNet to Places365 in this setting: 46.3% to 14.5% and 68.1% to 21.9% top-1, respectively. With FCT we obtain 20.3% top-1 compatibility performance, $\phi_{new}/h(\phi_{old}, \psi)$, only 1.6% below the accuracy of the new model. This shows generalization of the side-information and transformation models to out-of-domain data. In the Appendix, we report more results with this setup. We show that when the training objectives of ϕ_{old} and ϕ_{new} are disjoint (trained on disjoint training sets), a particularly difficult scenario, we can maintain model compatibility with performance very close to that of ϕ_{new} .

4.4. Sequence of Model Updates

Consider a sequence of model updates: $v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_n$. The first model, v_1 , directly computes features and side-information from the gallery set. The v_i embedding and side-information models, denoted by ϕ_i and ψ_i , respectively, are trained independently on the v_i training set. When updating from v_i to v_{i+1} , we apply FCT transformation on both features and side-information of v_i to make

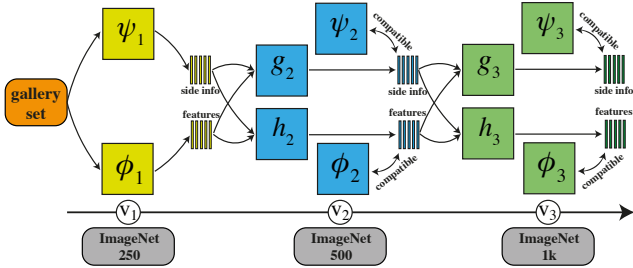


Figure 5. FCT for a sequence of model updates. Here, we transform both side-information and embedding to be compatible with their next versions.

them compatible with those of v_{i+1} :

$$\phi_{i+1} \xleftrightarrow{\text{compatible}} h_{i+1}(\phi_i, \psi_i), \psi_{i+1} \xleftrightarrow{\text{compatible}} g_{i+1}(\phi_i, \psi_i)$$

We use the same architectures for h_{i+1} and g_{i+1} transformations, as shown in Figure 4. Both transformations are trained using the MSE loss over the v_{i+1} training set, the same process as in Section 3.3.

Here, we experiment with a sequence of three model updates: the v_1 model is trained on the ImageNet-250 training set (a quarter of the full ImageNet training set). v_2 and v_3 models have the same setup as v_1 but with larger training sets: ImageNet-500 and ImageNet-1k, respectively. The v_i side-information model is a ResNet50-128 backbone trained on v_i training set using SimCLR. We compare two FCT scenarios: (i) v_1 is first updated to v_2 and then v_2 is updated to v_3 as shown in Figure 5, and (ii) v_1 is directly updated to v_3 using $h_{1 \rightarrow 3}$ transformation.

We show compatibility results in Table 5. Last two rows correspond to scenario (i) and (ii), respectively. FCT demonstrates a great performance (57.4% top-1) even in the two-hop update, Scenario (i). Scenario (ii) is slightly more accurate (+1.5%), but in practice requires extra storage of v_1 's feature and side-information (even after updating to v_i with $i > 1$). In Scenario (i) all v_i features and side-information are replaced when updating to v_{i+1} as shown in Figure 5.

If we do not store side-information, we observe significant drop in compatibility performance for both scenarios: 57.4% to 44.9% in Scenario (i) and 59.9% to 53.9% in Scenario (ii). This shows storing side-information is crucial for compatibility performance in a sequence of model updates. In the Appendix, we further demonstrate the importance of side-information by scaling to a longer sequence of small updates using subsets of the CIFAR-100 dataset [23].

5. Analysis

5.1. Types of Side-Information

No side-information We use no side-information as a simple baseline. In terms of implementation, we pass the zero-

Case	CMC top-1 — top-5 %	mAP@1.0
ϕ_1/ϕ_1	29.6 — 44.1	15.5
ϕ_2/ϕ_2	46.5 — 65.1	28.7
ϕ_3/ϕ_3	68.1 — 84.4	45.0
$\phi_2/h_2(\phi_1, \psi_1)$	40.5 — 58.8	26.4
$\phi_3/h_3(\phi_2, \psi_2)$	65.0 — 82.3	43.6
$\phi_3/h_3(h_2(\phi_1, \psi_1), g_2(\phi_1, \psi_1))$	57.4 — 78.0	33.5
$\phi_3/h_{1 \rightarrow 3}(\phi_1, \psi_1)$	59.9 — 79.3	37.0

Table 5. Model compatibility in sequence of updates. We see only a small degradation in performance in performing a sequence of updates $v_1 \rightarrow v_2 \rightarrow v_3$ versus going directly from $v_1 \rightarrow v_3$, corresponding to the last two rows.

Side-Information (ψ)	$\phi_{new}/h(\phi_{old}, \psi)$ CMC top-1 — top-5 %	mAP@1.0	CKA
None	61.8 — 80.5	39.9	0.0
Autoencoder	62.0 — 80.9	40.5	0.04
ϕ_{old}^{alt}	63.5 — 81.6	42.1	0.70
ϕ_{old} w/ Mixup	63.8 — 81.9	42.4	0.66
SimCLR-128-ImageNet-500	65.0 — 82.3	43.6	0.27
SimCLR-128-ImageNet-1k	66.5 — 83.3	46.6	0.25

Table 6. Comparison of different kinds of side-info. Both old and new architectures are ResNet-50 with feature size 128. The old model is trained on ImageNet-500 and the new model is trained on ImageNet-1k

vector into the projection layer in Figure 4.

Autoencoder We train a simple autoencoder with L2 reconstruction loss. The encoder and decoder architectures are convolutional, and based on MobileNetv2 [42]. The exact architectures are in the Appendix.

Alternate Old Model With this method, we train another version of ϕ_{old} , with the only difference being randomness from data order and model initialization. This is very similar to ensembling [14]. The intuition is that each model captures a different facet of the data. We denote this model as ϕ_{old}^{alt} .

Alternate Model + Mixup We perform the same training as for ϕ_{old} but with Mixup augmentation [50] applied only on images (for labels we use one-hot vectors). This encourages learning features of the data different from the old model, which will aid in transformation as they capture different invariances of the data.

Contrastive Model Here we train a SimCLR [9] self-supervised contrastive learning model to use as side-information. Taking the previous ideas to their natural extreme, a self-supervised contrastive approach directly captures invariances in the data, which will be useful for transfer even if it is not as useful for retrieval.

5.2. Side-information Ablation Results

Results are shown in Table 6. First note that no side-information results in 61.8% top-1 retrieval accuracy, which is 14.9% greater than BCT in this setting. This demonstrates the strength of our transformation setup even without side-

information. The autoencoder provides a slight improvement over no side-information. This makes sense considering much of the information of the autoencoder is encoding is contextually dependent on its decoder, and is therefore hard to extract *a priori*.

The ϕ_{old}^{alt} and the ϕ_{old}^{alt} with Mixup both provide substantial improvements over no side-information. It is well-known in deep learning literature that different runs of SGD for the same model produce diverse predictions [14, 29], which is particularly useful for ensembling. Here we show that we can leverage this fact for side-information. It’s important to note that simply applying a transformation from ResNet50-256 with no side-info trained on ImageNet-500 to the new model only yields 62.0% CMC top-1, which is less than the improvement we get from training a second old model. Mixup regularization gets a further 0.3% improvement over ϕ_{old}^{alt} , even though it does not provide a better model for retrieval, with 46.4% (ϕ_{old}/ϕ_{old}) CMC top-1.

The SimCLR model with embedding dimension of 128 trained on half of the ImageNet training set, denoted by SimCLR-128-ImageNet-500, outperforms other side-information choices considered (+1.2% over Mixup). This is despite the fact that CMC top-1 with just the SimCLR model is only 35.41%, far worse than even ϕ_{old} at 46.5%. What’s even more interesting is that when we train our transfer model with only SimCLR side-information and no ϕ_{old} we get 63.2% top-1 retrieval accuracy, better than no side-info. This demonstrates the efficacy of contrastive features for forward transformation. Even given the large gap in retrieval performance between ϕ_{old} and SimCLR, transfer performance for SimCLR is still superior, validating our intuition that it is important for side-information to capture all the invariances of the target data source.

We also consider a case that for the old model training we have access to the next version training set (ImageNet-1k), yet without labels. We use the unlabeled dataset to train the side-information model, shown by SimCLR-128-ImageNet-1k. This side-information results in 66.5% backward compatibility accuracy, +1.5% boost compared to the case when SimCLR-128-ImageNet-500 is used as the side-information model.

5.3. Centered Kernel Alignment (CKA) Analysis

CKA [22] is a similarity index identifying correspondences between representations. We show CKA with linear kernel between the old model features and side-information over the gallery set (the ImageNet validation set) in Table 6. We use CKA as a measure of complementarity of the side-information to the old model. Evidently, using mixing to train an alternate old model results in lower CKA (0.7 to 0.66), and hence higher compatibility accuracy when used as a side-information. More interestingly, the SimCLR side-information model has significantly less CKA with the old

model (0.27). This demonstrates features learned by SimCLR are *complementary* to those of the old model. Hence, combining them in FCT transformation results in excellent accuracy (65.0). Note that no side-information (the zero vector) and the Autoencoder have very low CKA but poor performance, demonstrating that this is not the only factor important to selecting a good model for side-information.

6. Conclusion

In this paper we study the compatibility problem in representation learning. We presented a new learning paradigm: forward compatible training, where we prepare for the next version of the model when training the current model. FCT enables efficient embedding updates to ensure compatibility with future models in large retrieval systems. We have shown that it is important to learn side-information to transform old features to new embedding models. Side-information captures features of the data which might not be useful for the old task, but are potentially important in the future. We demonstrate that contrastive learning methods work well to train side-information models. Through experimentation across multiple datasets, architectures, and losses, we share insights into necessary components important to the design of model-to-model transformation using side-information. Finally, by ensuring that training of the old and new embedding models are completely independent, we do not degrade model performance or retain biases from old model training, a problem in prior compatibility literature.

7. Limitations

FCT provides a mechanism to make old embeddings compatible with their future version. As a result, after transforming gallery embeddings, only downstream models which are *non-parametric* will not need to be updated (e.g. nearest neighbor retrieval in our setup). This is not a requirement with BCT since new embeddings are directly comparable with old embeddings. However, the convenience of not updating the downstream task model comes with limited accuracy improvement when updating the embedding model. We have not evaluated on cross-domain compatibility: when the old and new model are trained on widely different domains. In practice, it is unclear how likely this scenario is, but we suspect it will make learning a transformation and useful side-information more difficult.

Acknowledgements

We would like to thank Floris Chabert and Vinay Sharma for their advice and discussions, and Jason Ramapuram and Dan Busbridge for their help with self-supervised trainings.

References

- [1] Samira Abnar, Mostafa Dehghani, and Willem Zuidema. Transferring inductive biases through knowledge distillation. *arXiv preprint arXiv:2006.00555*, 2020. [1](#)
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. [13](#)
- [3] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. *arXiv preprint arXiv:1906.00910*, 2019. [4](#)
- [4] 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. [4](#), [14](#)
- [5] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, pages 67–74. IEEE, 2018. [5](#), [16](#)
- [6] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Thirty-fourth Conference on Neural Information Processing Systems (NeurIPS)*, 2020. [13](#)
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *arXiv preprint arXiv:2104.14294*, 2021. [13](#)
- [8] 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. [3](#)
- [9] 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. [3](#), [4](#), [5](#), [7](#), [13](#)
- [10] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. *Advances in Neural Information Processing Systems*, 33:22243–22255, 2020. [13](#)
- [11] 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. [5](#), [13](#)
- [12] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015. [4](#)
- [13] Mohamed Elhoseiny, Babak Saleh, and Ahmed Elgammal. Write a classifier: Zero-shot learning using purely textual descriptions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2584–2591, 2013. [3](#)
- [14] Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep ensembles: A loss landscape perspective. *arXiv preprint arXiv:1912.02757*, 2019. [7](#), [8](#)
- [15] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018. [4](#)
- [16] Akhilesh Gotmare, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. A closer look at deep learning heuristics: Learning rate restarts, warmup and distillation. In *International Conference on Learning Representations*, 2018. [13](#)
- [17] 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. [6](#), [12](#), [13](#)
- [18] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. [4](#)
- [19] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. [6](#), [12](#)
- [20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. [4](#)
- [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [13](#)
- [22] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International Conference on Machine Learning*, pages 3519–3529. PMLR, 2019. [8](#)
- [23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [7](#)
- [24] Hao Li, Pratik Chaudhari, Hao Yang, Michael Lam, Avinash Ravichandran, Rahul Bhotika, and Stefano Soatto. Rethinking the hyperparameters for fine-tuning. *arXiv preprint arXiv:2002.11770*, 2020. [3](#)
- [25] Wei-Hong Li, Xialei Liu, and Hakan Bilen. Universal representation learning from multiple domains for few-shot classification. *arXiv preprint arXiv:2103.13841*, 2021. [3](#)
- [26] Yixuan Li, Jason Yosinski, Jeff Clune, Hod Lipson, John E Hopcroft, et al. Convergent learning: Do different neural networks learn the same representations? In *FE@ NIPS*, pages 196–212, 2015. [3](#)
- [27] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts, 2016. [13](#)
- [28] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018. [3](#)
- [29] Horia Mania, John Miller, Ludwig Schmidt, Moritz Hardt, and Benjamin Recht. Model similarity mitigates test set overuse. *arXiv preprint arXiv:1905.12580*, 2019. [8](#)

- [30] Qiang Meng, Chixiang Zhang, Xiaoqiang Xu, and Feng Zhou. Learning compatible embeddings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9939–9948, 2021. [2](#), [3](#), [5](#), [11](#), [15](#)
- [31] Thomas Mensink, Efstratios Gavves, and Cees GM Snoek. Costa: Co-occurrence statistics for zero-shot classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2441–2448, 2014. [3](#)
- [32] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, pages 69–84. Springer, 2016. [4](#)
- [33] NVIDIA. Deeplearningexamples/pytorch/classification. <https://github.com/NVIDIA/DeepLearningExamples/tree/master/PyTorch/Classification/RN50v1>. 5, 2019. [12](#), [13](#)
- [34] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009. [3](#)
- [35] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019. [3](#)
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019. [16](#)
- [37] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021. [3](#)
- [38] Jason Ramapuram, Dan BusBridge, Xavier Suau, and Russ Webb. Stochastic contrastive learning. *arXiv preprint arXiv:2110.00552*, 2021. [14](#), [15](#)
- [39] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. [3](#)
- [40] Marcus Rohrbach, Michael Stark, György Szarvas, Iryna Gurevych, and Bernt Schiele. What helps where—and why? semantic relatedness for knowledge transfer. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 910–917. IEEE, 2010. [3](#)
- [41] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. [5](#), [16](#)
- [42] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. [7](#)
- [43] Yantao Shen. openbct. <https://github.com/YantaoShen/openBCT>, 2020. [13](#)
- [44] 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. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [12](#), [14](#), [15](#)
- [45] Jake Snell, Kevin Swersky, and Richard S Zemel. Prototypical networks for few-shot learning. *arXiv preprint arXiv:1703.05175*, 2017. [3](#)
- [46] Elena Voita and Ivan Titov. Information-theoretic probing with minimum description length. *arXiv preprint arXiv:2003.12298*, 2020. [2](#)
- [47] Chien-Yi Wang, Ya-Liang Chang, Shang-Ta Yang, Dong Chen, and Shang-Hong Lai. Unified representation learning for cross model compatibility. *arXiv preprint arXiv:2008.04821*, 2020. [2](#), [3](#), [5](#), [11](#), [15](#)
- [48] Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. Supermasks in superposition. *arXiv preprint arXiv:2006.14769*, 2020. [3](#)
- [49] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? *arXiv preprint arXiv:1411.1792*, 2014. [3](#)
- [50] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. [7](#)
- [51] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016. [4](#)
- [52] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. 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](#), [16](#)