

Positive-Congruent Training: Towards Regression-Free Model Updates

Sijie Yan* Yuanjun Xiong Kaustav Kundu Shuo Yang Siqi Deng
Meng Wang Wei Xia Stefano Soatto

AWS/Amazon AI

ysijie@gmail.com, {yuanjx, kaustavk, shuoy, siq Deng, mengw, wxia, soattos}@amazon.com

Abstract

Reducing inconsistencies in the behavior of different versions of an AI system can be as important in practice as reducing its overall error. In image classification, sample-wise inconsistencies appear as “negative flips”: A new model incorrectly predicts the output for a test sample that was correctly classified by the old (reference) model. Positive-congruent (PC) training aims at reducing error rate while at the same time reducing negative flips, thus maximizing congruency with the reference model only on positive predictions, unlike model distillation. We propose a simple approach for PC training, Focal Distillation, which enforces congruence with the reference model by giving more weights to samples that were correctly classified. We also found that, if the reference model itself can be chosen as an ensemble of multiple deep neural networks, negative flips can be further reduced without affecting the new model’s accuracy.

1. Introduction

Imagine a “new and improved” version of the software that manages your photo collection exhibiting mistakes absent in the old one. Even if the average number of errors has decreased, every new mistake on your old photos feels like a step backward (Fig. 1), leading to perceived *regression*.¹ We tackle regression in the update of image classification models, where an old (reference) model is replaced by a new (updated) model. We call test samples that are correctly labeled by both the new and the old models *positive-congruent*. On the other hand, *negative flips* are samples correctly classified by the old model but incorrectly by the new one. Their

*Currently at The Chinese University of Hong Kong. Work conducted while at AWS.

¹In the Software Industry, regression refers to the deterioration of performance after an update. Even if the updated model has on average better performance, any regression risks breaking post-processing pipelines, which is why trained models are updated sporadically despite the steady stream of improvements reported in the literature.

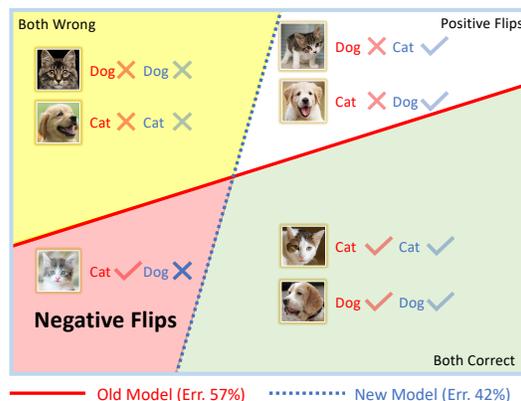


Figure 1: **Regression in model update:** When updating an old classifier (red) to a new one (dashed blue line), we correct mistakes (top-right, white), but we also introduce errors that the old classifier did not make (**negative flips**, bottom-left, red). While on average the errors decrease (from 57% to 42% in this toy example), regression can wreak havoc with downstream processing, nullifying the benefit of the update.

fraction of the total number is called negative flip rate (NFR), which measures regression. There are also *positive flips*, where the old model makes mistakes that the new one corrects, which are instead beneficial.

Reducing the NFR could be accomplished by reducing the overall error rate (ER). However, reducing the ER is neither necessary nor sufficient to reduce the NFR. In fact, models trained on the same data with different initial conditions, data augmentations, and hyperparameters tend to yield similar error rates, but with *errors occurring on different samples*. For instance, two ResNet152 models trained on ImageNet with just different initializations achieve the same accuracy of 78.3% but differ on 10.4% of the samples in the validation set (Fig. 2). Assuming an equal portion of positive and negative flips, updating from one model to another incurs a 5.2% NFR. This can be reduced simply by trading errors

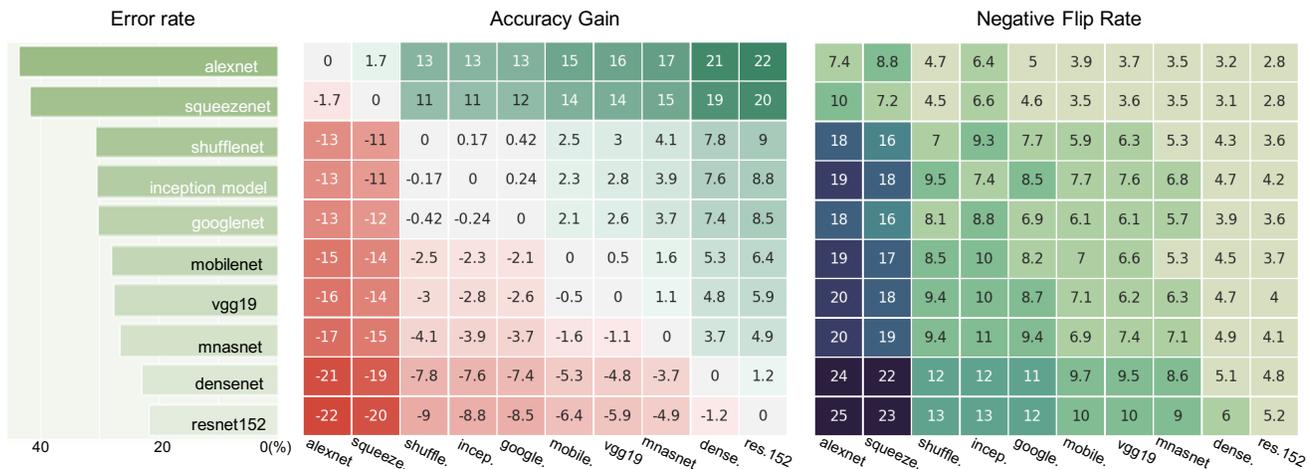


Figure 2: Differences in overall accuracy vs. negative flip rates. We measure the accuracy gains and negative flip rates on the ILSVRC12 [26] validation set when we update from various old models (y -axis) to new models (x -axis). Multiple CNN architectures [16, 11, 36] are used for comparison. First we observe training the same architecture twice can cause regression. If the new model has higher/lower capacity than the old one, the error rate decreases/increases (above/below the diagonal), but the NFR is always positive. Note that in some cases the NFR is of the same order of the error rate change.

while maintaining an equal error rate. Thus, reducing the ER is not necessary to reduce the NFR. Reducing the ER to a value other than zero is not sufficient either: From Fig. 2 we see that, when updating an old AlexNet [16] to a new Resnet-152 [11] ER reduces significantly, but we still suffer a 3% NFR.

On the other hand, NFR can be made zero trivially by copying the old model, albeit with no ER reduction. Model distillation tries to bias the update towards the old model while reducing the ER. However, we wish to mimic the old model *only when it is right*. Thus, regression-free updates are not achieved by ordinary distillation.

Reducing the error rate and the NFR are two separate and independent goals. We call any training procedure that aims to minimize both the error rate and the negative flip rate *Positive-Congruent (PC) training*, or PCT for short.

We first propose a simple method for PC training when the old model is given. In this case, PCT is achieved by minimizing an additional loss term along with the standard classification loss (empirical cross-entropy). We consider several variants for the additional loss, and find *Focal Distillation* (FD), a variant of the distillation loss that we introduce to bias the model towards positive congruence (Sect. 4.2), to be most effective (Sect. 5).

We also explore the problem of “future-proofing” the reference model to facilitate subsequent PC training. This *forward setting* pertains to selecting the reference model, rather than PC training new ones. Motivated by the observation above that different training instantiations can yield

the same error rate with different erroneous samples, we propose a simple approach using ensembles. We show that this results in lower regression. In practice, ensembles are not viable at inference time in large-scale applications due to the high cost. Nonetheless, this approach can be used as a paragon for PC training of a single model in future works.

Our **contributions** can be summarized as follows: (i) We formalize the problem of quality regression in pairs of classifiers, and introduce the first method for positive congruent training (Sect. 3); (ii) we propose a variant of model distillation (Sect. 4.2) to perform PC training of a deep neural network (DNN); (iii) we show that reference models can be adapted for future PC-Training by replacing a single model with an ensemble of DNNs (Sect. 4.3). We conduct experiments on large scale image classification benchmarks (Sect. 5), providing both a baseline (Focal Distillation) and a paragon (Ensemble) for future evaluation.

2. Related Work

PC training relates to the general areas of **continual learning** [6, 15, 29, 31], incremental learning [23, 18], open set recognition [1, 27], and sequential learning [10, 20]. The goal is to evolve a model to incorporate additional data or concepts, as reviewed in [8, 21]. Specifically, [10, 15, 19, 18, 23, 25, 35] aim at training with a growing number of samples, classes, and tasks while maintaining similar performance to previously learned classes/tasks [20, 24].

PC training does not restrict the new model to reuse the

old model weights or architecture. Instead, it allows changing them along with the training algorithm, loss functions, hyper-parameters, and training set. Unlike most work focused on reducing forgetting [5, 18, 19], PC training focuses on reducing NFR.

PC Training has strong connections to **knowledge distillation** [12] and weight consolidation [15], that aim to keep a model close to a reference one, regardless of whether they are correct or not. PC training focuses on enforcing similarity only on cohorts that belong to negative flips.

Along a similar vein, **backward-compatible training** [28, 30] aims to design new models that are inter-operable with old ones. PC training is a form of backward-compatibility, but prior literature still measures compatibility using error rate, rather than NFR. As shown in experiments, these methods do not directly help reduce the NFR. Implicit negative flips during training of a single model are studied in [32]. We also conduct experiments on the evolution of negative flips during training of the new model.

3. Negative Flips in Model Updates

Let $x \in X$ (e.g., an image) and $y \in Y = \{y_1, \dots, y_K\}$ (e.g., a label), identified with an integer in $\{1, \dots, K\}$. Let p denote an unknown distribution from which a dataset is drawn $\mathcal{D} = \{(x_i, y_i) \sim p(x, y)\}_{i=1}^N$. Let p^{old} be the pseudo-distribution associated with the old (*reference*) model, and p^{new} the same for the new one.² These are parametric functions that, ideally, approximate the true posterior which is the optimal (Bayesian) discriminant, in the sense of minimizing the expected probability of error. More specifically, when evaluated on the sample (x_i, y_i) , $p^{\text{new}}(\cdot|\cdot)$ takes the form

$$p^{\text{new}}(y = y_i | x = x_i) = \frac{\exp(\langle \vec{y}_i, \phi_w^{\text{new}}(x_i) \rangle)}{\sum_{j=1}^K \exp(\langle \vec{y}_j, \phi_w^{\text{new}}(x_i) \rangle)} \quad (1)$$

where $\vec{y}_i \in \mathbb{R}_{\geq 0}^K$ is the “one-hot” (indicator) vector corresponding to the categorical variable $y_i \in \{1, \dots, K\}$. p^{old} has a similar form but with a different *architecture* ϕ that has known parameters. We refer to ϕ_w as the *model* or *discriminant*. The final class prediction of the model, ϕ_w , is denoted by \hat{y} , where, $\hat{y}(x_i) = \arg \max_y p(y|x_i)$.

3.1. Negative Flips

In Fig. 1 we illustrate the different errors occurring after a model upgrade, comparing the predictions \hat{y}^{old} and \hat{y}^{new} . The consistent areas are where ϕ_w^{old} and ϕ_w^{new} are both either correct ($\hat{y}^{\text{old}}(x_i) = \hat{y}^{\text{new}}(x_i) = y_i$) or incorrect ($\hat{y}^{\text{old}}(x_i) \neq y_i, \hat{y}^{\text{new}}(x_i) \neq y_i$). Positive flips occur when samples were

²In addition to having different architectures, the models can be trained with different optimization schemes, minimizing different loss functions, on different datasets. For simplicity, below we consider the case of models with same or different architectures, trained on the same dataset \mathcal{D} using the same loss function and optimization scheme.

incorrectly predicted by ϕ_w^{old} , but correctly predicted by ϕ_w^{new} . The most relevant to this paper are **negative flips**, where samples flipped from correct ϕ_w^{old} to incorrect predictions ϕ_w^{new} . The **negative flip rate (NFR)** measures the fraction of samples that are negatives

$$\text{NFR} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}(\hat{y}_i^{\text{new}} \neq y_i, \hat{y}_i^{\text{old}} = y_i) \quad (2)$$

where, $\mathbf{1}(\cdot)$ is the indicator function.

3.2. Persistence of Negative Flips

Deep neural network (DNN) classifiers are usually trained by minimizing the empirical cross-entropy (CE) loss:

$$\min_w \mathcal{L}_{\text{CE}}(\phi, w) = \min_w \frac{1}{N} \sum_{i=1}^N -\log p_w(y_i|x_i) \quad (3)$$

This objective is minimized by Stochastic Gradient Descent (SGD) [14] annealed to convergence to or near a local minimum.³ The final discriminant is determined by (a) the DNN architecture ϕ , (b) the training methodology, including optimization scheme and associated hyper-parameters such as learning rate and momentum, (η, μ) , (c) the dataset on which the model is trained \mathcal{D} and (d) the initialization $w^{(0)}$. Unless explicitly enforced, a new model is typically not “close” to an old one. Even if based on the same architecture and trained on the same dataset, different runs on SGD can converge to distant points in the loss landscape [4]. What is similar along limit cycles of solutions is the average error rate in the training set, which is minimal; what is different is the samples on which errors occur. These are negative flips, measured empirically in Fig. 2. There, we measure the difference of error rates and NFR between pairs of models trained on the ILSVRC12 [26] dataset⁴. We observe that, from earlier architecture design such as AlexNet [16], to recent ones such as DenseNet [13], although the overall error rates have dropped, the NFR remain non-negligible. In many cases, the NFR is of the same order of magnitude as the accuracy gain from the model update. This stubborn phenomenon calls for a dedicated solution that does not rely on reducing the error rate to zero in order to have regression-free model updates.

4. Positive Congruent Training

In most applications, the reference model is inherited and cannot be changed directly, so PC training is limited to the new model. However, in some cases we may get to design

³In reality, the landscape of DNNs trained with CE does not comprise isolated (Morse) critical points, but rather “wide valleys” [7, 3], and convergence is only forced by annealing the learning rate, without which the solution travels on limit cycles.

⁴All models in this experiment are publicly available in the PyTorch [22] online model zoo

Approach	\mathcal{F}	\mathcal{D}	q
Naive Baseline	$\mathbf{1}(\hat{y}^{\text{old}}(x_i) = y_i)$	Cross Entropy	\vec{y}_i
Focal Distillation - KL (FD-KL)	$\alpha + \beta \cdot \mathbf{1}(\hat{y}^{\text{old}}(x_i) = y_i)$	τ -scaled KL-Divergence	$p^{\text{old}}(y x_i)$
Focal Distillation - Logit Matching (FD-LM)	$\alpha + \beta \cdot \mathbf{1}(\hat{y}^{\text{old}}(x_i) = y_i)$	ℓ_2 distance	$\phi_w^{\text{old}}(x_i)$

Table 1: Different approaches for targeted positive congruent training and the corresponding choices of \mathcal{F} , \mathcal{D} , and q in the PC loss functions.

the reference model, aiming to make it easily updated with PCT. For the first ‘‘backward’’ case, we first present a simple approach to reducing NFR and achieve PCT. The idea is to simultaneously maximize the ‘‘both correct’’ region in Fig. 1, which measures *positive congruence*, along with the overall accuracy of the new model on the training set.

To achieve PC Training, we use the following objective function

$$\min_w \mathcal{L}_{\text{CE}}(\phi^{\text{new}}, w) + \lambda \mathcal{L}_{\text{PC}}(\phi^{\text{new}}, w; \phi^{\text{old}}) \quad (4)$$

where \mathcal{L}_{CE} is defined in (3) and $\mathcal{L}_{\text{PC}}(\phi^{\text{new}}, \phi^{\text{old}}, w)$ is the *positive congruence (PC) loss* with multiplier λ . We propose a generic form of the PC loss function as

$$\lambda \mathcal{L}_{\text{PC}}(\phi^{\text{new}}, w; \phi^{\text{old}}) = \mathcal{F}(x_i) \mathcal{D}(\phi_w^{\text{new}}(x_i), q(x_i)), \quad (5)$$

where \mathcal{F} is a filter function $\mathcal{F} \in f : X \rightarrow \mathbb{R}_{\geq 0}$ that applies a weight for each training sample based on the model outputs and \mathcal{D} is a distance function that measures the difference of the new model’s outputs to a certain target vector $q(x_i)$ conditioned on x_i . The PC loss serves to bias training towards maximal positive congruence. Favoring iso-error-rate solutions [4] with maximal positive congruence can lead to lower NFR relative to the reference model. In Table 1 we illustrate the different choices of \mathcal{F} and \mathcal{D} . We introduce different PC losses in the the next two sections.

4.1. Naive Baseline

We first consider a simple PC loss that measures CE on the samples where the old model made correct predictions, denoted as

$$\mathcal{L}_{\text{PC}}^{\text{naive}} = \frac{1}{N} \sum_{i=1}^N -\mathbf{1}(\hat{y}^{\text{old}}(x_i) = y_i) \log p_w^{\text{new}}(y_i|x_i). \quad (6)$$

This is just a re-weighting of samples that the old models classified correctly by a factor $1 + \lambda$. The resulting model is labeled ‘‘Naive’’ in the experiments in Sect. 5. This method does not help reduce the NFR: It is hard to find a suitable hyperparameter λ that gives sufficient weights to positive samples without inducing rote memorization. We now explore a more direct way to pass information from the old model to the new one.

4.2. Focal Distillation

Knowledge distillation [12] aims at biasing the new model to be ‘‘close’’ to the old one during training. In this sense it has the potential to reduce the NFR, but also to reduce positive flips, which are instead beneficial. Focal Distillation (FD) mitigates this risk by using the loss

$$\mathcal{L}_{\text{PC}}^{\text{Focal}} = - \sum_{i=1}^N [\alpha + \beta \cdot \mathbf{1}(\hat{y}^{\text{old}}(x_i) = y_i)] \mathcal{D}(\phi_w^{\text{new}}, \phi_w^{\text{old}}), \quad (7)$$

which penalizes a distance \mathcal{D} between the output of the two models, weighted by a filtering function $\mathcal{F} = \alpha + \beta \cdot \mathbf{1}(\hat{y}^{\text{old}}(x_i) = y_i)$, to the target $q = \phi_w^{\text{old}}(x_i)$. The filter function applies a basic weight α for all samples in the training set and an additional weight to the samples correctly predicted by the old model. This biases the new model towards the old one for positive samples in the old model, thus imposing a cost on NFR as well as the overall error rate. When $\alpha = 1$ and $\beta = 0$, focal distillation reduces to ordinary distillation [12]. When $\alpha = 0$ and $\beta > 0$, we are only applying the distillation objective to the training samples predicted correctly by the old model. We assess these choices empirically in Sec. 5.2.

A possible choice of \mathcal{D} is the temperature scaled KL-divergence of [12]:

$$\mathcal{D}^{\text{KL}}(\phi_w^{\text{new}}, \phi_w^{\text{old}}) = \text{KL} \left[\sigma \left(\frac{\phi_w^{\text{new}}(x_i)}{\tau} \right), \sigma \left(\frac{\phi_w^{\text{old}}(x_i)}{\tau} \right) \right]. \quad (8)$$

Here, σ is the ‘‘Softmax’’ function as shown in Eq. (1) and τ is the temperature scaling factor. In Sec. 4.2, we found that τ has to be set to a large number (e.g. $\tau = 100$) for this PC loss to reduce NFR. In this case, the above approaches the distance among ‘‘logits’’ [12, 2]

$$\mathcal{D}^{\text{LM}}(\phi_w^{\text{new}}, \phi_w^{\text{old}}) = \frac{1}{2} \|\phi_w^{\text{new}}(x_i) - \phi_w^{\text{old}}(x_i)\|_2^2. \quad (9)$$

In Sect. 5 we compare various settings of FD with the naive baseline and observe that training with the FD can indeed reduce the NFR at the cost of a slight increase in average accuracy for the new model.

4.3. PC Training by Ensembles

The fact that there are iso-error sets in the loss landscape of DNN (loci in weight space that correspond to models with

equal error rate³) [4] can be used to our advantage to reduce NFR as outlined in the introduction. In some cases, one may be able to choose both the new *and the old models* in a manner that reduces the NFR. Geometrically, one can think of each model as representing an iso-error equivalence class: A sphere in the space of models, centered at zero error, where all models achieve the same error rate but differ by which samples they mistake. Since cordal averaging reduces the distance to the origin [34] and therefore the average distance to other models (for instance, future PC trained ones), we hypothesize that replacing each model with an ensemble trained on the same data might reduce the NFR.

In particular, we assume that ϕ^{old} and ϕ^{new} are trained independently, each from a collection of models $\{\phi_j^{\text{old}}\}_{j=1}^N$ and $\{\phi_j^{\text{new}}\}_{j=1}^N$. We combine the results from all the models by averaging their discriminants

$$p^{\text{ensemble}}(y|x_i) = \sigma \left[\frac{1}{L} \sum_{j=1}^L \phi_{w_j}^{\text{old}}(y|x_i) \right]. \quad (10)$$

where L is the number of models in the ensemble. Similarly, we can define ϕ^{new} . Note that models in each ensemble are trained independently with different initialization.

We test our hypothesis empirically in Sect. 5 where we find that, indeed, the NFR between the two ensembles is lower than that between any two individual models.

While this may not seem surprising at first, since ensembles reduce error rate, we note that, as anticipated in the introduction, NFR can be reduced independently of ER. In fact, empirically we observe that *the NFR decreases more rapidly than the average error* as the size of the ensemble grows (Fig. 3), indicating that ensembling improves PC training beyond simply lowering the average error rate. Note that this phenomenon holds for any pair of ensembles with either the same or different architectures. This makes it suitable as a paragon to explore the upper limit of PC training. While test-time ensembles are not practical at scale, this observation points to promising areas of future investigation by collapsing ensembles.

5. Experiments

We test the methods presented on image classification tasks using ImageNet [9] and iNaturalist [33]. We start with the simplest case: same architecture and training data, different training runs. We compare baseline methods for PCT and ensembles. We then extend the approaches to encompass 1) architecture changes; 2) changing number of training samples per class; 3) increase in the number of classes. Finally, we present empirical evidence on the source of negative flips and how PC training affects them.

Relative NFR. Since overall error rate (ER) is an upper bound to the NFR, comparison is challenging across datasets

PCT Approach	Error Rate (%)		NFR (%)	Rel. NFR (%)	#Params
	ϕ^{old}	ϕ^{new}			
No Treatment	30.24	30.29	6.44	30.48	12M
Naive	30.24	29.34	5.72	27.95	12M
BCT [28]	30.24	29.66	6.39	30.88	12M
FD-KL	30.24	30.63	2.50	11.70	12M
FD-LM	30.24	30.47	2.35	11.06	12M
Ensemble	26.07	25.98	1.70	8.85	187M

(a) ILSVRC12

PCT Approach	Error Rate (%)		NFR (%)	Rel. NFR (%)	#Params
	ϕ^{old}	ϕ^{new}			
No Treatment	40.69	41.06	7.77	31.91	14M
Naive	40.69	45.47	11.07	41.05	14M
FD-KL	40.69	41.81	2.83	11.41	14M
FD-LM	40.69	41.78	2.71	10.94	14M
Ensemble	35.68	35.42	2.01	8.82	221M

(b) iNaturalist

Table 2: Experiments on training the same model architecture multiple times on three datasets: (a) ILSVRC12 [26], (b) iNaturalist [33]. Applying PC training in the case can significantly reduce the NFR without hurting the new model’s accuracy. Among them, the focal distillation works the best for the backward compatibility setting. “Ensemble” refers to PC training with ensembles in the forward setting. “Rel. NFR” refers to the relative NFR value shown in Eq. 11. “#Params” refers to the number of parameters of the new model.

with different ERs. For this reason, we introduce the relative NFR:

$$\text{NFR}_{\text{rel}} = \frac{\text{NFR}}{(1 - \text{ER}_{\text{old}}) * \text{ER}_{\text{new}}}, \quad (11)$$

where ER_{new} and ER_{old} denote the error rate of ϕ^{new} and ϕ^{old} . The denominator is the expected error rate on the subset of samples predicted correctly by the old model. This is a naive estimation of NFR if the two models are independent of each other. The relative NFR is a measure of reduction in regression from a PCT method which factors out overall model accuracy.

5.1. Implementation Details

Unless otherwise noted, we minimize the CE loss (3) with batch-size 2048 for 90 epochs on each dataset. The learning rate starts at 0.1 and decreases by 1/10 every 30 epochs. We use $\lambda = 1$ for Eq. 4. Focal distillation and ensembles are implemented with PyTorch [22]. Models in Fig. 2 are from the PyTorch model zoo.

5.2. Positive Congruent Training

We start with the same architecture [11] trained twice on the same dataset, either ILSVRC12 [26] or iNaturalist [33],

and report results on their official validation sets in Table 2⁵. **Variants of PCT** include the **naive baseline** in Eq. 6, which unfortunately does not reduce the NFR markedly. Yet it is worth noting that, due to the fact that models are not independent, being trained on the same data, their relative NFR are always less than 100%. We additionally evaluate the recently proposed BCT method [28] for aligning representation between models. It does not reduce NFR.

Focal distillation in Eq. 7 (FD) has two variants, either using the KL-divergence after soft-max (*FD-KL*) or matching logits before soft-max (*FD-LM*). In FD, there are two parameters, α and β , controlling the filter function \mathcal{F} . We set $\alpha = 1$ and $\beta = 5$ in this set of experiments. The results on ILSVRC12 are summarized in Table 2a. We see that training with both variants of focal distillation, we can reduce the NFR from 6% to around 2.5%, a 60% relative reduction. We also note that the reduction of NFR comes as the cost of a slightly increased error rate of the new model.

Finally we evaluate the **ensemble** approach. We use two independently trained ensembles each composed of 16 ResNet-18 [11] for ϕ^{old} and ϕ^{old} . We see ensembles achieve the lowest absolute NFR, as low as 1.7%. Even when discounting the accuracy gain of ensembles, we also observe better relative NFR than other approaches in the backward setting. The price to pay is computational cost, a multiplier depending on the number of models in the ensemble.

Training from scratch vs fine-tuning. In practical scenarios, we often fine-tune from an existing model pretrained on a large-scale dataset. To compare the impact of this, on the ImageNet dataset we trained all models from scratch but on the iNaturalist dataset [33] we fine-tune from model weights pretrained on ImageNet. As shown in Table 2b, we observe starting from a common pretrained model does not guarantee regression-free. We observe a 7% NFR without treatment in finetuning on iNaturalist. But with PC training, we can reduce the NFR to 2.01%

Distances in focal distillation In focal distillation we presented two different choices of distance. In Table 2, we observe that the logit matching distance function (FD-LM) leads to marginally better NFR reduction. It is also worth to note that the KL-divergence only effectively reduces NFR when the temperature scaling parameter τ becomes large, say 100, at which point it resembles logit matching [12].

Weights and focus in focal distillation In Focal Distillation, we use a default focal weight $\beta = 5$ and a base weight $\alpha = 1$. We now experiment changing the focal weight and base weight and report results in Table 3. First note two special cases: 1) $\alpha = 0$ and $\beta = 1$, when we only apply distillation to samples classified correctly by the old model, which is ineffective: The new model must still learn the old model’s behavior from all samples in order to be able to

α	β	ϕ^{new} Error Rate (%)	NFR (%)	Rel. NFR(%)
0	0	30.29	6.44	30.48
0	1	31.52	5.25	23.88
1	0	31.12	3.90	17.96
1	1	30.59	2.75	12.89
1	2	30.79	2.73	12.71
1	5	30.47	2.35	11.06
1	10	30.44	2.39	11.26
1	20	33.55	6.94	29.65

Table 3: Effect of different α and β values in focal distillation. We train ResNet-18 for both ϕ^{old} and ϕ^{new} on ILSVRC12 dataset [26]. We use the logit matching distance function for focal distillation because it has better effect in reducing NFR.

overcome negative flips. 2) $\alpha = 1$ and $\beta = 0$, corresponding to ordinary distillation, which can help reduce NFR only to a limited extent. Focal Distillation is most effective when $\alpha = 1$ and β is between 5 to 10, reducing NFR to 2.35%.

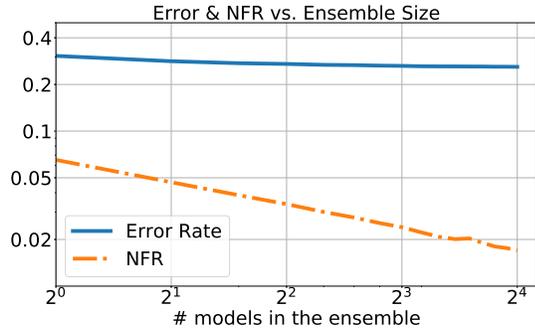
Impact of ensemble size In ensemble experiments on ILSVRC12, we used a default ensemble size of 16. In Figure 3 we explore the behavior of NFR as the number of components in an ensemble changes. Although the accuracy plateaus with the growing ensemble size, the NFR keeps decreasing in a logarithmic scale. This observation is also confirmed in the case of ensembles with different architectures. Although having a large ensemble is impractical for real-world applications, these results corroborate the geometric picture sketched in Sect. 4.3 and suggest that it may be possible to achieve zero NFR in a model update without having to drive the ER to zero, which may be infeasible, as suggested in the Introduction.

5.3. Other Changes in Model Updates

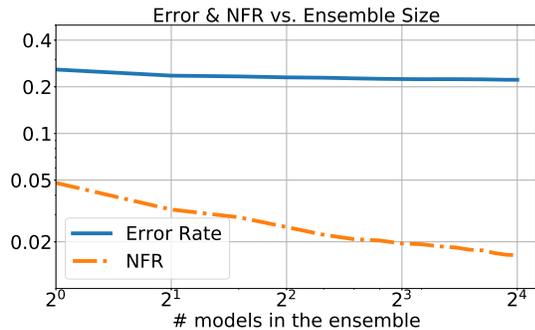
Having studied the simplest case of PC training, we now move to more practical cases of image classification model updates. We enumerate and study the following types of changes: 1) model architecture; 2) number of training samples per class; 3) number of classes.

Changing model architectures to a larger model, with either a similar architecture (Resnet-18 to Resnet-50) or a dissimilar one (ResNet18 to DenseNet-161), yields results summarized in Table 4. A larger model is expected to have a lower overall error rate. However, without any treatment, it still suffers from a significant number of negative flips. We then apply the focal distillation approach in this case, we see a reduction in NFR, although the new model incurs a slight increase in error rate. We also observe that ensembles reduce NFR, this time without increasing the new model’s error rate. This is noteworthy because the two ensembles in this case do not share any information other than the dataset on which they are trained. In Fig. 3 we also visualize the

⁵We follow the common practice to use the official validation sets on these datasets as test set and a small portion training set for validation.



(a) ResNet-18 → ResNet-18



(b) ResNet-18 → ResNet-50

Figure 3: The effect of ensemble size (number of individual models) on error rates and the NFRs when both the old and new models are both independently trained ensembles. Results are reported on ILSVRC12. (a) shows the results when the old model and new one are both composed of ResNet-18 [11] models. (b) shows the case when the old model is an ensemble of ResNet-18 [11] models and the new one is formed by ResNet-50 [11] models. In both case, the error rates plateau quickly but the NFRs keep decreasing as the ensemble size increases.

trend of NFR vs. the ensemble size. It suggests that, even in the case of architecture changes, it may be possible to achieve zero NFR.

Increasing training samples per visual category is shown in Table 5a for the ILSVRC12 dataset, with the old model trained on all 1000 classes using 50% of the samples per category, and the new model trained on 100% set. We observe that all new models decrease the ER, but PC trained ones achieve better reduction of NFR.

Increasing the number of classes is tested on ILSVRC12 with the old model trained on a random subset of 500 classes and the new model on the entire ILSVRC12 training set. We use ResNet-50 for this study. The evaluation is done on the validation subset which comprises all samples of the original 500 classes on which the older model is trained on. The results are shown in Table 5b. Since the number of

Approach	Error Rate (%)		NFR (%)	Rel. NFR (%)	#params
	ϕ^{old}	ϕ^{new}			
No Treatment	30.24	25.85	4.89	27.12	25M
Naive	30.24	24.41	3.78	22.20	25M
FD-KD	30.24	26.32	2.90	15.79	25M
FD-LM	30.24	26.53	2.92	15.78	25M
Ensemble	26.07	22.20	1.64	9.99	409M

(a) ResNet-18 → ResNet-50

Approach	Error Rate (%)		NFR (%)	Rel. NFR (%)	#params
	ϕ^{old}	ϕ^{new}			
No Treatment	30.24	22.86	4.03	25.24	25M
Naive	30.24	21.60	3.28	21.76	25M
FD-KD	30.24	23.52	2.50	15.24	25M
FD-LM	30.24	23.83	2.56	15.40	25M

(b) ResNet-18 → DenseNet161

Table 4: Experiments of PC training methods in changes of model architectures on ILSVRC12 [26]. Here the old model architecture is ResNet-18 [11]. We experiment with the new models with both Resnet-50 [11] and Denset161 [13] architectures. Results suggest that PC training method are effective in reducing regression in the face of model architecture changes.

Approach	Error Rate (%)		NFR (%)	Rel. NFR (%)	#params
	ϕ^{old}	ϕ^{new}			
No Treatment	28.58	23.65	3.98	19.47	25M
Naive	28.45	24.46	3.27	18.68	25M
FD-KD	28.45	25.20	2.89	15.84	25M
FD-LM	28.45	24.77	2.85	16.09	25M
Ensemble	26.39	22.09	2.75	16.88	100M

(a) Increase in # samples

Approach	Error Rate (%)		NFR (%)	Rel. NFR (%)	#params
	ϕ^{old}	ϕ^{new}			
No Treatment	19.30	23.65	8.05	41.90	25M
Naive	19.28	23.74	7.70	40.20	25M
FD-KD	19.28	24.14	7.07	36.29	25M
FD-LM	19.28	25.11	7.37	36.35	25M
Ensemble	17.53	21.98	6.72	37.06	100M

(b) Increase in # classes

Table 5: Experiments of training data change on ILSVRC12 [26] dataset. We use Resnet-18 architecture for both ϕ^{old} and ϕ^{old} .

categories increase for the newer model, so does the total error on the validation set. Yet, PC training is able to reduce NFR.

5.4. Prediction of negative flips

First we look at the potential relationship between a test sample being a negative flip and its prediction uncertainty. In

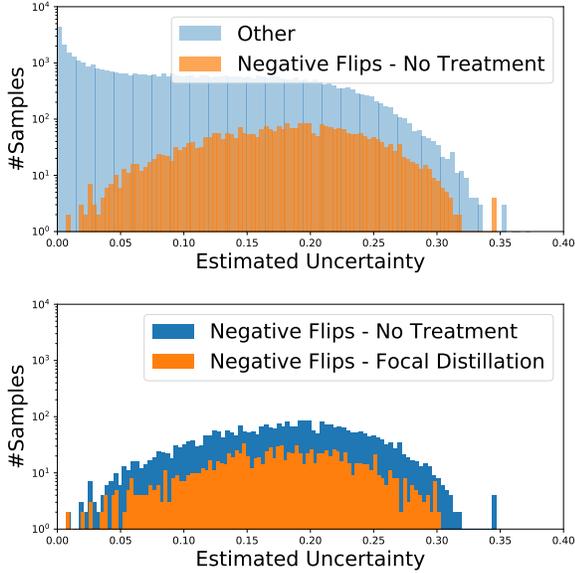


Figure 4: Distribution of sample uncertainty estimated by [17], measured on test validation set of ILSVRC12 [26]. Here we use ResNet-18 [11] as the model architecture in experiment’s. **Top:** Uncertainty histogram of negative flipped samples and other test samples between two ResNet-18 models trained without PC training. **Bottom:** Uncertainty histogram of negative flipped sample in two model pairs, the first without PC training. The new model in the second pair is trained with focal distillation based PC training. The y -axes are in log-scale.

this experiment we use the uncertainty estimated by a deep ensemble of ResNet-18 [11] as described in [17]. We visualize the results in Fig. 4. We observe that test samples that become negative flips have relatively high estimated uncertainty but our measure does not clearly separate them from other samples. We also visualize the histogram of uncertainty estimates before and after we apply focal distillation based PCT, as shown in Fig. 4. We observe that PCT with focal distillation reduces the NFR with an almost uniform chance for samples with different level of uncertainty. This observation may suggest alternate strategies to designing PC training methods.

The second study is on the evolution of negative flip rates during the training of the new model. We visualize the relative NFRs and error rates at every epoch when training a new model. We present results for both training without PCT and with focal distillation based PCT (FD-LM) in Fig. 5. Our first observation is that the NFRs change in a similar trend as the new models’ error rates and reduce as the models are trained longer. By applying PCT during training, we observe that the relative NFR drops faster in early epochs compared with no treatment. As training goes on, the new

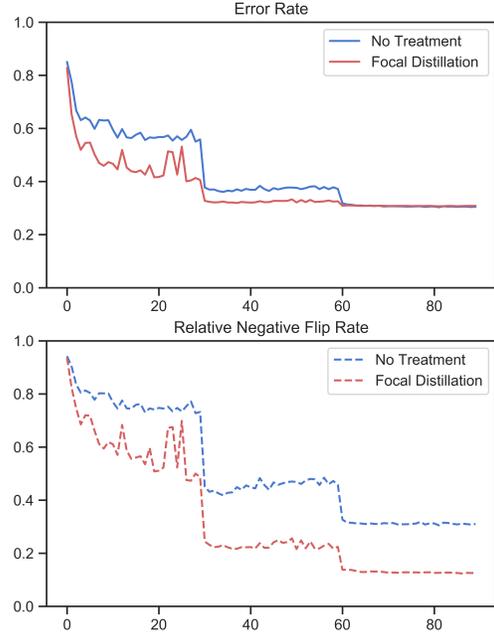


Figure 5: Evolution of error rates and relative NFR during training of the new models on ILSVRC12 [26]. We compare focal distillation-based PC training with no treatment. Generally NFR follows the trend of the error rate during training. Focal distillation lead to a gap in relative NFR in early epochs and keeps the gap as training evolves.

model with PCT maintains the same gap in NFR until the training terminates.

6. Discussion

Large-scale DNN-based classifiers are typically only a part of more complex systems that involve additional post-processing. In the simplest cases, the output of the DNN is used to map the data to a metric space, where it is clustered and searched. In more complex cases, classifiers are part of an elaborate system that includes high-level reasoning. In all cases, changing the classifier can break the system, which is why DNN models are seldom updated despite the steady improvements reported in the literature. PC training could enable seamless adoption of improved models, and ensure steady progress and increased accessibility to the state of the art in image classification. In this paper, we have merely scratched the surface of PC training, as the methods proposed have obvious limitations: Ensembling is not viable in large-scale systems, even though it achieves the highest NFR reduction. Focal distillation reduces the NFR, but at the price of a slight increase in error rate. Further exploration is needed to identify methods that can achieve paragon performance at the baseline cost of single model distillation.

References

- [1] Abhijit Bendale and Terrance E Boult. Towards open set deep networks. In *CVPR*, 2016. [2](#)
- [2] Cristian Bucila, R. Caruana, and Alexandru Niculescu-Mizil. Model compression. In *KDD*, 2006. [4](#)
- [3] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: Biasing gradient descent into wide valleys. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124018, 2019. [3](#)
- [4] Pratik Chaudhari and Stefano Soatto. Stochastic gradient descent performs variational inference, converges to limit cycles for deep networks. In *ITA*, 2018. [3](#), [4](#), [5](#)
- [5] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *ECCV*, 2018. [3](#)
- [6] Zhiyuan Chen and Bing Liu. Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 2018. [2](#)
- [7] Anna Choromanska, Mikael Henaff, Michael Mathieu, Gérard Ben Arous, and Yann LeCun. The loss surfaces of multilayer networks. In *Artificial intelligence and statistics*, 2015. [3](#)
- [8] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. Continual learning: A comparative study on how to defy forgetting in classification tasks. *arXiv:1909.08383*, 2019. [2](#)
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. [5](#)
- [10] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv:1312.6211*, 2013. [2](#)
- [11] Kaiying He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. [2](#), [5](#), [6](#), [7](#), [8](#)
- [12] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015. [3](#), [4](#), [6](#)
- [13] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, 2017. [3](#), [7](#)
- [14] Jack Kiefer, Jacob Wolfowitz, et al. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 23(3):462–466, 1952. [3](#)
- [15] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. [2](#), [3](#)
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. [2](#), [3](#)
- [17] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *NIPS*, 2017. [8](#)
- [18] Zhizhong Li and Derek Hoiem. Learning without forgetting. *TPAMI*, 40(12):2935–2947, 2017. [2](#), [3](#)
- [19] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *NIPS*, 2017. [2](#), [3](#)
- [20] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*. 1989. [2](#)
- [21] 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. [2](#)
- [22] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *NIPS*. 2019. [3](#), [5](#)
- [23] Ameeya Prabhu, Philip Torr, and Puneet Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *ECCV*, 2020. [2](#)
- [24] Roger Ratcliff. Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review*, 1990. [2](#)
- [25] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *CVPR*, 2017. [2](#)
- [26] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015. [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [27] Walter J Scheirer, Anderson de Rezende Rocha, Archana Sapkota, and Terrance E Boult. Toward open set recognition. *TPAMI*, 35(7):1757–1772, 2012. [2](#)
- [28] Yantao Shen, Yuanjun Xiong, Wei Xia, and Stefano Soatto. Towards backward-compatible representation learning. In *CVPR*, 2020. [3](#), [5](#), [6](#)
- [29] Daniel L Silver and Robert E Mercer. The task rehearsal method of life-long learning: Overcoming impoverished data. In *Conference of the Canadian Society for Computational Studies of Intelligence*, 2002. [2](#)
- [30] Megha Srivastava, Besmira Nushi, Ece Kamar, Shital Shah, and Eric Horvitz. An empirical analysis of backward compatibility in machine learning systems. In *KDD*, 2020. [3](#)
- [31] Sebastian Thrun. Lifelong learning algorithms. In *Learning to learn*, pages 181–209. Springer, 1998. [2](#)
- [32] Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J

- Gordon. An empirical study of example forgetting during deep neural network learning. *arXiv:1812.05159*, 2018. 3
- [33] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *CVPR*, 2018. 5, 6
- [34] Geoffrey S Watson. *Statistics on spheres*. Wiley-Interscience, 1983. 5
- [35] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. *ICML*, 2017. 2
- [36] Hang Zhang, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Zhi Zhang, Haibin Lin, Yue Sun, Tong He, Jonas Mueller, R Manmatha, et al. Resnest: Split-attention networks. *arXiv:2004.08955*, 2020. 2