

Low-shot Visual Recognition by Shrinking and Hallucinating Features

Bharath Hariharan and Ross Girshick

Facebook AI Research (FAIR)

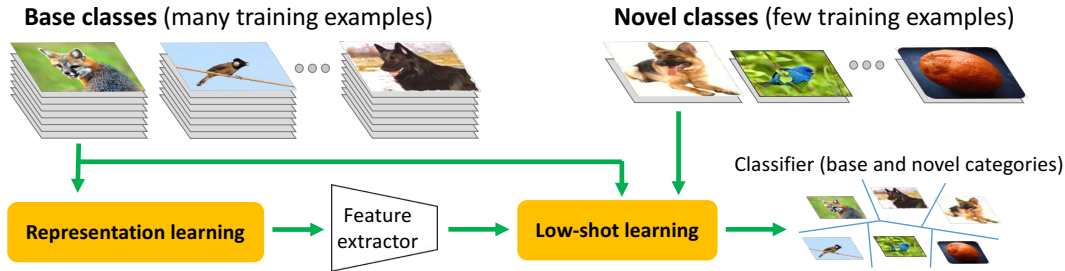


Figure 1: **Our low-shot learning benchmark in two phases: representation learning and low-shot learning.** Modern recognition models use large labeled datasets like ImageNet to build good visual representations and train strong classifiers (*representation learning*). However, these datasets only contain a fixed set of classes. In many realistic scenarios, once deployed, the model might encounter novel classes that it also needs to recognize, but with very few training examples available (*low-shot learning*). We present two ways of significantly improving performance in this scenario: (1) a novel loss function for representation learning that leads to better visual representations that generalize well, and (2) a method for hallucinating additional examples for the data-starved novel classes.

Abstract

Low-shot visual learning—the ability to recognize novel object categories from very few examples—is a hallmark of human visual intelligence. Existing machine learning approaches fail to generalize in the same way. To make progress on this foundational problem, we present a low-shot learning benchmark on complex images that mimics challenges faced by recognition systems in the wild. We then propose (1) representation regularization techniques, and (2) techniques to hallucinate additional training examples for data-starved classes. Together, our methods improve the effectiveness of convolutional networks in low-shot learning, improving the one-shot accuracy on novel classes by $2.3\times$ on the challenging ImageNet dataset.

1. Introduction

Recently, error rates on benchmarks like ImageNet [8] have been halved, and then halved again. These gains come from deep convolutional networks (ConvNets) that learn rich feature representations [24]. It is now clear that if an application has an a priori fixed set of visual concepts and thousands of examples per concept, an effective way to build an object recognition system is to train a deep ConvNet. But what if these assumptions are not satisfied and the network

must learn novel categories from very few examples?

The ability to perform *low-shot learning*—learning novel concepts from very few examples—is a hallmark of the human visual system. We are able to do this not only for natural object categories such as different kinds of animals, but also for synthetic objects that are unlike anything we’ve seen before [39]. In contrast, in spite of significant improvements in recognition performance, computational recognition approaches fail to generalize well from few examples [25]. Our goal in this paper is to make progress towards imparting this human ability to modern recognition systems.

Our first contribution is a low-shot learning benchmark based on the challenging ImageNet1k dataset. As shown in Figure 1, our benchmark is implemented in two phases. In the *representation learning phase*, the learner tunes its feature representation on a set of base classes that have many training instances. In the *low-shot learning phase*, the learner is exposed to a set of novel classes with only a few examples per class and must learn a classifier over the *joint* label space of base and novel classes. This benchmark simulates a scenario in which the learner is deployed in the wild and has to quickly learn novel concepts it encounters from very little training data. Unlike previous low-shot learning tests (*e.g.*, [13, 25]) we measure the learner’s accuracy on both the base and novel classes. This provides a sanity check that accuracy gains on novel classes do not come at the expense of a large

loss in performance on the base classes. This evaluation protocol follows the standard way that image classifiers are evaluated on popular benchmarks like ImageNet, thus easing the comparison of progress on low-shot learning to the typical data-rich scenario.

Next, we investigate how to improve the learner’s performance on the benchmark. We build on the intuition that certain modes of intra-class variation generalize across categories (*e.g.*, pose transformations). We present a way of “hallucinating” additional examples for novel classes by transferring modes of variation from the base classes. These additional examples improve the one-shot top-5 accuracy on novel classes by *15 points* (absolute) while also maintaining accuracy on the base classes.

Finally, we show that the feature representation learnt in the first phase has a large impact on low-shot generalization ability. Specifically, we formulate a loss function that penalizes the difference between classifiers learnt on large and small datasets, and then draw connections between this loss and regularization of *feature activations*. We show that simply regularizing feature activations can increase one-shot, top-5 accuracy on novel classes by *9 points* (absolute) without harming base class performance. Combining this better representation with the hallucination strategy pushes our improvement up to *18 points* above the baseline.

2. Related work

One-shot and low-shot learning. One class of approaches to one-shot learning uses generative models of appearance that tap into a global [13] or a supercategory-level [38] prior. Generative models based on strokes [26] or parts [47] have shown promise in restricted domains such as hand-written characters [28, 25]. They also work well in datasets without much intra-class variation or clutter, such as Caltech 101 [13]. Dixit *et al.* [9] leverage a corpus with attribute annotations to generate additional examples by varying attributes. We also propose a way to generate additional examples, but our model is non-parametric and directly generates feature vectors. Jia *et al.* [22] present a promising alternative to generation using Bayesian reasoning to infer an object category from a few examples; however, in [22] the full, large-scale training set is available during training.

Among discriminative approaches, early work attempted to use a single image of the novel class to adapt classifiers from similar base classes [3, 32] using simple hand-crafted features. Bertinetto *et al.* [4] regress from single examples to a classifiers, while Wang and Hebert [46] regress from classifiers trained on small datasets to classifiers trained on large datasets. Recent “meta-learning” techniques learn to directly map training sets and test examples to classification outputs [45, 15, 35]. We compare favorably with these approaches in our experiments.

Amongst representation learning approaches, metric

learning, such as the triplet loss [43, 40, 14] or siamese networks [23, 18], has been used to automatically learn feature representations where objects of the same class are closer together. Such approaches have shown benefits in face identification [43]. On benchmarks involving more general Internet imagery, such as ImageNet [8], these methods perform worse than simple classification baselines [36], and it is unclear if they can benefit low-shot learning.

Zero-shot learning. Zero-shot recognition uses textual or attribute-level descriptions of object classes to train classifiers. While this problem is different than ours, the motivation is the same: to reduce the amount of data required to learn classifiers. One line of work uses hand-designed attribute descriptions that are provided to the system for the novel categories [37, 27, 12]. Another class of approaches embeds images into word embedding spaces learnt using large text corpora, so that classifiers for novel concepts can be obtained simply from the word embedding of the concept [16, 42, 31, 48]. A final class of approaches attempts to directly regress to image classifiers from textual descriptions [11, 29] or from prototypical images of the category [21]. Similar to our benchmark, Chao *et al.* [5] propose that zero-shot learning evaluation should also include the training categories that do have examples. We believe this evaluation style is good for both zero and low-shot learning.

Transfer learning. The ability to learn novel classes quickly is one of the main motivations for multitask and transfer learning. Thrun’s classic paper convincingly argues that “learning the n -th task should be easier than learning the first,” with ease referring to sample complexity [44]. However, recent transfer learning research has mostly focussed on the scenario where large amounts of training data are available for novel classes. For that situation, the efficacy of pre-trained ConvNets for extracting features is well known [10, 33, 41]. There is also some analysis on what aspects of ImageNet training aid this transfer [1, 2]. For faces, Taigman *et al.* [43] find that low-dimensional feature representations transfer better on faces and Galanti *et al.* [17] provide some theoretical justification for this finding. This work hints at a link between the complexity of the feature representation and its generalizability, a link which we also observe in this paper. We find that stronger base classifiers generalize better than weaker classifiers (*e.g.* comparing ResNet-10 to ResNet-50 [19]). There have also been novel losses proposed explicitly to aid transfer, such as the multi-verse loss of Littwin and Wolf [30]. Our paper also proposes novel losses designed specifically for low-shot learning.

3. A low-shot learning benchmark

Our goal is to build a benchmark for low-shot learning that mimics situations that arise in practice. Current recognition systems require days or even weeks of training on expensive hardware to develop good feature representations.

The trained recognition systems may then be deployed as a service to be used by downstream applications. These downstream applications may need the ability to recognize novel categories, but they may have neither the training data required, nor the infrastructure needed to retrain the models. Thus, there are two natural phases: in the first phase, we have the data and resources to train sophisticated feature extractors on large labelled datasets, and in the second phase, we want to add additional categories to our repertoire at minimal computational and data cost.

Our low-shot learning benchmark implements a similar setup. It employs a learner, two training phases, and one testing phase. The learner is assumed to be composed of a feature extractor and a multi-class classifier. The benchmark is agnostic to the specific form of each component.

During *representation learning* (training phase one), the learner receives a fixed set of base categories C_{base} , and a dataset D containing a large number of examples for each category in C_{base} . The learner uses D to set the parameters of its feature extractor.

In the second phase, which we call *low-shot learning*, the learner is given a set of categories C^l that it must learn to distinguish. $C^l = C_{base} \cup C_{novel}$ is a mix of base categories C_{base} , and unseen novel categories C_{novel} . For each novel category, the learner has access to only n positive examples, where $n \in \{1, 2, 5, 10, 20\}$. For the base categories, the learner still has access to D . The learner may then use these examples and its feature extractor to set the parameters of its multi-class classifier while also optionally modifying the feature extractor.

In the testing phase, the learnt model predicts labels from the combined label space $C_{base} \cup C_{novel}$ on a set of previously unseen test images. To measure the variability in low-shot learning accuracy, we repeat the low-shot learning and testing phases for 5 trials, each time with a random draw of examples for the novel classes. We report the mean accuracy and the standard deviation over these trials.

The simplest, and commonly used, baseline approach is to train a ConvNet with label cross-entropy loss in the representation learning phase and then train a new linear classifier head in the low-shot learning phase. We now show significant improvements on this baseline, first by a novel strategy of hallucinating additional training examples (Section 4) and then by improving the representation itself (Section 5).

4. Better low-shot learning through generation

In the low-shot learning phase, our goal is to train good classifiers for novel categories from only a few examples. Intuitively, the challenge is that these examples capture very little of the category’s intra-class variation. For instance, if the category is a particular bird species, then we may only have examples of the bird perched on a branch, and none of it in flight. The classifier might then erroneously conclude

that this novel category only consists of perched birds.

However, this mode of variation is common to many bird species, including those we have encountered in the base classes. From the many base class examples we have seen, we can understand the transformation that relates perched bird images to the image of the corresponding bird in flight, and then use this transformation to “hallucinate” additional examples for our novel bird category. If we were given the set of all such category-independent transformations, then we can hallucinate as many new examples for each novel category example as there are transformations.

However, we do not have a pre-defined set of transformations that we can apply. But we can take a non-parametric approach. Any two examples z_1 and z_2 belonging to the same category represent a plausible transformation. Then, given a novel category example x , we want to apply to x the transformation that sent z_1 to z_2 . That is, we want to complete the transformation “analogy” $z_1 : z_2 :: x : ?$.

We do this by training a function G that takes as input the concatenated feature vectors of the three examples $[\phi(x), \phi(z_1), \phi(z_2)]$. It produces as output a “hallucinated” feature vector (of the same dimensionality as ϕ), which corresponds to applying the $z_1 \rightarrow z_2$ transformation to x . We use an MLP with three fully connected layers for G .

We first describe how we train G , and then show how we use the generated examples in the low-shot learning phase.

4.1. Learning to generate new examples

To train G , we first collect a dataset of completed analogies from our base classes. To do this we first cluster the feature vectors of the examples in each base category into a fixed number of clusters (100). This is to keep computational complexity manageable. Next, for each pair of centroids c_1^a, c_2^a in one category a , we search for another pair of centroids c_1^b, c_2^b from another category b , such that the cosine distance between $c_1^a - c_2^a$ and $c_1^b - c_2^b$ is minimized. We collect all such quadruplets $(c_1^a, c_2^a, c_1^b, c_2^b)$ with cosine similarity greater than zero into a dataset D_G . See Figure 2 for example transformation analogies.

We now use the dataset D_G to train G . For each quadruplet $(c_1^a, c_2^a, c_1^b, c_2^b)$, we feed (c_1^a, c_1^b, c_2^b) to the generator. Let $\hat{c}_2^a = G([c_1^a, c_1^b, c_2^b])$ be the output of the generator. We then minimize $\lambda L_{mse}(c_2^a, \hat{c}_2^a) + L_{cls}(W, \hat{c}_2^a, a)$, where:

1. $L_{mse}(\hat{c}_2^a, c_2^a)$ is the mean squared error between the generator’s output and the true target of the analogy c_2^a .
2. $L_{cls}(W, \hat{c}_2^a, a)$ is the classification loss, where W is the fixed linear classifier on the base classes learnt during representation learning, and $L_{cls}(W, x, y)$ is the log loss of the classifier W on the example (x, y) .

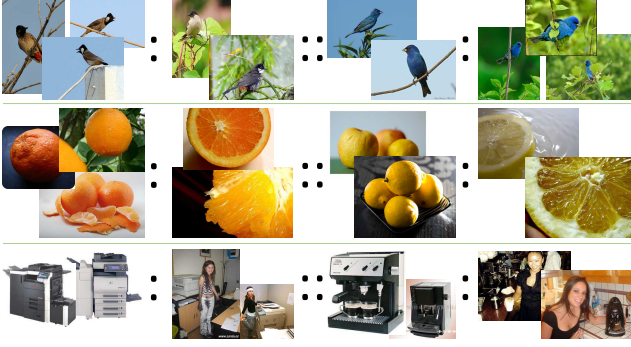


Figure 2: Example mined analogies. Each row shows the four image clusters that form the four elements in the analogy. **Row 1:** birds with a sky backdrop vs birds with greenery in the background. **Row 2:** whole fruits vs cut fruit. **Row 3:** machines (printer, coffee making) in isolation vs the same machine operated by a human.

4.2. Using generated examples for low-shot learning

Our generated examples are unlikely to be as good as real examples, but should provide a useful bias to the classifier when only a few real examples are present. Therefore we want to rely on generated examples only when the number of real examples is low.

Concretely, we have a hyperparameter k (set through cross-validation), which is the minimum number of examples per novel category that we want to have. If the actual number of real examples for a novel category, n , is less than k , then we additionally generate $k - n$ hallucinated examples. To generate a synthetic example for a novel category l , we sample the feature vector of a “seed” example $\phi(x)$ from one of the n real examples for this category, and a pair of cluster centroids c_1^a, c_2^a from a base category a chosen uniformly at random. We then pass this triplet through G , and add the hallucinated feature vector $G([\phi(x), c_1^a, c_2^a])$ to our training set with label l . We then train the logistic regression classifier on this mix of real and generated data in the usual manner.

5. Better representations for low-shot learning

We now turn to the question of improving representation learning so as to enable better low-shot learning. As described above, the learner consists of a *feature extractor* ϕ and a *classifier* W . The goal of representation learning is a *good feature extractor*: one that enables learning of effective classifiers from few examples. Intuitively, our goal is to reduce the difference between classifiers trained on large datasets and classifiers trained on small datasets so that those trained on small datasets generalize better.

We first describe a proposal that encodes this goal in a loss that can be minimized during representation learning. Then, we draw connections to several alternatives.

5.1. Squared gradient magnitude loss (SGM)

We assume that the classifier W is linear, e.g., the last layer of a ConvNet. Let D denote a large labeled dataset of base class images. Typically, training the feature extractor ϕ and the classifier W on D involves minimizing a classification objective with respect to ϕ and W :

$$\min_{W, \phi} L_D(\phi, W) = \min_{W, \phi} \frac{1}{|D|} \sum_{(x, y) \in D} L_{cls}(W, \phi(x), y) \quad (1)$$

where $L_{cls}(W, x, y)$ is the multiclass logistic loss on an example x with label y for a linear classifier W :

$$L_{cls}(W, x, y) = -\log p_y(W, x) \quad (2)$$

$$p_k(W, x) = \frac{\exp(w_k^T x)}{\sum_j \exp(w_j^T x)}. \quad (3)$$

We modify this training procedure as follows. We *simulate* low-shot learning experiments on the base classes by considering several tiny training sets $S \subset D$, $|S| \ll |D|$. We then want to reduce the difference between classifiers trained on the large dataset D (using the feature extractor ϕ) and classifiers trained on these small datasets S .

The classifier trained on D is just W . Training a classifier on S involves solving a minimization problem:

$$\min_V L_S(\phi, V) = \min_V \frac{1}{|S|} \sum_{(x, y) \in S} L_{cls}(V, \phi(x), y) \quad (4)$$

We want the minimizer of this objective to match W . In other words, we want W to minimize $L_S(\phi, V)$. $L_S(\phi, V)$ is convex in V (Fig. 3), so a necessary and sufficient condition for this is that the gradient of $L_S(\phi, V)$ at $V = W$, denoted by $\nabla_V L_S(\phi, V)|_{V=W}$, is 0. More generally, the closer W is to the global minimum of $L_S(\phi, V)$, the lower the magnitude of this gradient. Thus, we want to minimize:

$$\tilde{L}_S(\phi, W) = \|\nabla_V L_S(\phi, V)|_{V=W}\|^2 \quad (5)$$

The gradient $\nabla_V L_S(\phi, V)$ has a simple analytical form (see supplemental material for details):

$$\nabla_V L_S(\phi, V) = [g_1(S, V), \dots, g_K(S, V)] \quad (6)$$

$$g_k(S, V) = \frac{1}{|S|} \sum_{(x, y) \in S} (p_k(V, \phi(x)) - \delta_{yk}) \phi(x) \quad (7)$$

where K is the number of classes, δ_{yk} is 1 when $y = k$ and 0 otherwise, and p_k is as defined in equation (3).

This leads to an analytical form for the function $\tilde{L}_S(\phi, W) : \frac{1}{|S|^2} \sum_{k=1}^K \|\sum_{(x, y) \in S} (p_k(W, \phi(x)) - \delta_{yk}) \phi(x)\|^2$. We use this analytical function of W and ϕ as a loss.




Figure 3: Motivation for the SGM loss. We want to learn a representation ϕ such that the arg min of the small set training objective $L_S(\phi, V)$ matches W , the classifier trained on a large dataset D .

We consider an extreme version of this loss where S is a *single example* (x, y) . In this case,

$$\tilde{L}_S(\phi, W) = \sum_{k=1}^K (p_k(W, \phi(x)) - \delta_{yk})^2 \|\phi(x)\|^2 \quad (8)$$

$$= \alpha(W, \phi(x), y) \|\phi(x)\|^2. \quad (9)$$

where $\alpha(W, \phi(x), y) = \sum_k (p_k(W, \phi(x)) - \delta_{yk})^2$ is a per-example weight that is higher for data points that are misclassified. Thus the loss becomes a weighted L_2 regularization on the feature activations.

Our final loss, which we call SGM for Squared Gradient Magnitude, averages this over all examples in D .

$$L_D^{SGM}(\phi, W) = \frac{1}{|D|} \sum_{(x,y) \in D} \alpha(W, \phi(x), y) \|\phi(x)\|^2 \quad (10)$$

We train our feature representation by minimizing a straightforward linear combination of the SGM loss and the original classification objective.

$$\min_{W, \phi} L_D(\phi, W) + \lambda L_D^{SGM}(\phi, W) \quad (11)$$

λ is obtained through cross-validation.

Batch SGM. Above, we used singleton sets as our tiny training sets S . An alternative is to consider every *mini-batch* of examples B that we see during SGD as S . Hence, we penalize the squared gradient magnitude of the average loss over B , yielding the loss term: $\lambda \tilde{L}_B(\phi, W)$. In each SGD iteration, our total loss is thus the sum of this loss term and the standard classification loss. Note that because this loss is defined on mini-batches the number of examples per class in each mini-batch is a random variable. Thus this loss, which we call “batch SGM”, optimizes for an expected loss over a distribution of possible low-shot values n .

5.2. Feature regularization-based alternatives

In Eq. (9), it can be shown that $\alpha(W, \phi(x), y) \in [0, 2]$ (see supplementary). Thus, in practice, the SGM loss is

dominated by $\|\phi(x)\|^2$, which is much larger. This suggests a simple squared L_2 norm as a loss:

$$\min_{W, \phi} L_D(\phi, W) + \lambda \frac{1}{|D|} \sum_{(x,y) \in D} \|\phi(x)\|^2. \quad (12)$$

While L_2 regularization is a common technique, note that here we are regularizing the *feature representation*, as opposed to regularizing the *weight vector*. Regularizing the feature vector norm has been a staple of unsupervised learning approaches to prevent degenerate solutions [34], but to the best of our knowledge it hasn’t been considered in supervised classification.

We can also consider other ways of regularizing the representation, such as an L_1 regularization:

$$\min_{W, \phi} L_D(\phi, W) + \lambda \frac{1}{|D|} \sum_{(x,y) \in D} \|\phi(x)\|_1. \quad (13)$$

We also evaluate other forms of feature regularization that have been proposed in the literature. The first of these is dropout [20], which was used in earlier ConvNet architectures [24], but has been eschewed by recent architectures such as ResNets [19]. Another form of feature regularization involves minimizing the correlation between the features [6, 7]. We also compare to the multiverse loss [30] which was shown to improve transfer learning performance.

Why should feature regularization help? When learning the classifier and feature extractor jointly, the feature extractor can choose to encode less discriminative information in the feature vector because the classifier can learn to ignore this information. However, when learning new classifiers in the low-shot phase, the learner will not have enough data to identify discriminative features for the unseen classes from its representation. Minimizing the norm of the feature activations might limit what the learner can encode into the features, and thus force it to only encode useful information.

5.3. Metric-learning based approaches

A common approach to one-shot learning is to learn a good distance metric that generalizes to unseen classes. We train a ConvNet with the triplet loss as a representative baseline method. The triplet loss takes as input a triplet of examples (x, x_+, x_-) , where x and x_+ belong to the same category while x_- does not:

$$L_{\text{triplet}}(\phi(x), \phi(x_+), \phi(x_-)) = \max(\|\phi(x_+) - \phi(x)\| - \|\phi(x_-) - \phi(x)\| + \gamma, 0). \quad (14)$$

The loss encourages x_- to be at least γ farther away from x than x_+ is.

6. Experiments and discussion

6.1. Low-shot learning setup

We use the ImageNet1k challenge dataset for experiments because it has a wide array of classes with significant intra-class variation. We divided the 1000 ImageNet categories randomly into 389 base categories and 611 novel categories (listed in the supplementary material).

Many of the methods we evaluate have hyperparameters that need to be cross-validated. Since we are interested in generalization to novel classes, we did not want to cross-validate on the same set of classes that we test on. We therefore constructed two disjoint sets of classes by dividing the base categories into two subsets C_{base}^1 (193 classes) and C_{base}^2 (196 classes) and the novel categories into C_{novel}^1 (300 classes) and C_{novel}^2 (311 classes). Then, for cross-validating hyperparameters, we provided the learner with $C^{cv} = C_{base}^1 \cup C_{novel}^1$ in the low-shot learning and testing phase, and evaluated its top-5 accuracy on the combined label set C^{cv} . The hyperparameter setting that gave the highest top-5 accuracy was then frozen. We then conducted our final experiments using these hyperparameter settings by providing the learner with $C^{fin} = C_{base}^2 \cup C_{novel}^2$. All reported numbers in this paper are on C^{fin} .

Our test images are a subset of the ImageNet1k validation set: we simply restricted it to only include examples from the classes of interest (C^{cv} or C^{fin}). Performance is measured by top-1 and top-5 accuracy on the test images for each value of n (number of novel examples per category). We report the mean and standard deviation from 5 runs each using a different random sample of novel examples during the low-shot training phase.

To break down the final performance metrics, we report separately the average accuracy on the test samples from the novel classes and on all test samples. While our focus is on the novel classes, we nevertheless need to ensure that good performance on novel classes doesn't come at the cost of lower accuracy on the base classes.

6.2. Network architecture and training details

For most of our experiments, we use a small ten-layer ResNet architecture [19] as our feature extractor ϕ (details in supplementary material). When trained on all 1000 categories of ImageNet, it gives a validation top-5 error rate of 16.7% (center crop), making it similar to AlexNet [24]. We use this architecture because it's relatively fast to train (2 days on 4 GPUs) and resembles state-of-the-art architectures. Note that ResNet architectures, as described in [19], do not use dropout. Later, we show some experiments using the larger and deeper ResNet-50 architecture.

For all experiments on representation learning, except the triplet embedding, the networks are trained from scratch for 90 epochs on the base classes. The learning rate starts

at 0.1 and is divided by 10 every 30 epochs. The weight decay is fixed at 0.0001. For the triplet embedding, we first pretrain the network using a softmax classifier and log loss for 90 epochs, and then train the network further using the triplet loss and starting with a learning rate of 0.001. We stop training when the loss stops decreasing (55 epochs). This schedule is used because, as described in [36], triplet networks train slowly from scratch.

For methods that introduce a new loss, there is a hyperparameter that controls how much we weigh the new loss. Dropout also has a similar hyperparameter that governs what fraction of activations are dropped. We set these hyperparameters by cross-validation.

For our generator G , we use a three layer MLP with ReLU as the activation function. We also add a ReLU at the end, since ϕ is known to be non-negative. All hidden layers have a dimensionality of 512.

In the low-shot learning phase, we train the linear classifier using SGD for 10000 iterations with a mini-batch size of 1000. We cross-validate for the learning rate.

6.3. Training with class imbalance

The low-shot benchmark creates a heavily imbalanced classification problem. During low-shot learning the base classes may have thousands of examples, while each novel class has only a few examples. We use two simple strategies to mitigate this issue. One, we oversample the novel classes when training the classifier by sampling uniformly over classes and then uniformly within each chosen class. Two, we L_2 regularize the multi-class logistic classifier's weights by adding weight decay during low-shot learning. We find that the weight of the classifier's L_2 regularization term has a large impact and needs to be cross-validated.

6.4. Results

Impact of representation learning. We plot a subset of the methods¹ in Figure 4, and show the full set of numbers in Tables 1 and 2. The plots show the mean top-5 accuracy, averaged over 5 low-shot learning trials, for the novel classes, and over the combined set of novel and base classes. The standard deviations are low (generally less than 0.5%, see supplementary material) and are too small to display clearly as error bars. Top-1 accuracy and numerical values are in the supplementary material. We observe that:

- When tested just on base classes, many methods perform similarly (not shown), but their performance differs drastically in the low-shot scenarios, especially for small n . Thus, *accuracy on base classes does not generalize to novel classes, especially when novel classes have very few training examples.*

¹The subset reduces clutter, making the plots more readable. We omit results for Batch SGM, Dropout and L_1 because Batch SGM performs similarly to SGM and L_2 , while L_1 and Dropout perform worse.

Representation	Lowshot phase	n=1	2	5	10	20
<i>ResNet-10</i>						
Baseline	Classifier	14.1	33.3	56.2	66.2	71.5
Baseline	Generation* + Classifier	29.7	42.2	56.1	64.5	70.0
SGM*	Classifier	23.1	42.4	<i>61.7</i>	69.6	<i>73.8</i>
SGM*	Generation* + Classifier	<i>32.8</i>	46.4	<i>61.7</i>	<i>69.7</i>	<i>73.8</i>
Batch SGM*	Classifier	23.0	42.4	61.9	69.9	74.5
L1*	Classifier	20.8	40.8	59.8	67.5	71.6
L2*	Classifier	29.1	<i>47.4</i>	62.3	68.0	70.6
Triplets	Classifier	24.5	41.8	56.0	61.3	64.2
Dropout [20]	Classifier	26.8	43.9	59.6	66.2	69.5
Decov [7]	Classifier	13.0	33.9	59.3	68.9	73.4
Multiverse [30]	Classifier	13.7	30.6	52.5	63.8	71.1
Baseline	Data augmentation	16.0	31.4	52.7	64.4	71.8
Baseline	Model Regression [46]	20.7	39.4	59.6	68.5	73.5
Baseline	Matching Network [45]	41.3	51.3	<i>62.1</i>	67.8	71.8
Baseline-ft	Classifier	12.5	29.5	53.1	64.6	70.4
<i>ResNet-50</i>						
Baseline	Classifier	28.2	51.0	71.0	<i>78.4</i>	<i>82.3</i>
Baseline	Generation* + Classifier	<i>44.8</i>	59.0	<i>71.4</i>	77.7	<i>82.3</i>
SGM*	Classifier	37.8	57.1	72.8	79.1	82.6
SGM*	Generation* + Classifier	45.1	<i>58.8</i>	<i>72.7</i>	79.1	82.6

Table 1: Top-5 accuracy on only novel classes. Best are bolded and blue; the second best are italicized and red. *Our methods.

- Batch SGM, SGM, and L_2 are top performers overall with L_2 being better for small n . They improve novel class accuracy by more than 10 points for small n (1 or 2) and more than 3 points for $n > 10$. L_1 also improves low-shot performance, but the gains are much smaller.
- Dropout is on par with SGM for small n , but ends up being similar or worse than the baseline for $n \geq 5$ in terms of all class accuracy. Empirically, dropout also reduces feature norm, suggesting that implicit L_2 feature regularization might explain some of these gains.
- Triplet loss improves accuracy for small n but is 5 points worse than the baseline for $n = 20$ in terms of all class accuracy. While more sophisticated variants of the triplet loss may improve performance [36], feature regularization is both effective and much simpler.
- The decov loss [7] provides marginal gains for higher values of n but is outperformed by the feature regularization alternatives.

As an additional experiment, we also attempted to fine-tune the baseline representation on all the base class examples and the small set of novel class examples. We found that this did not improve performance over the frozen representation (see Baseline-ft in Tables 1 and 2). This indicates that finetuning the representation is not only expensive, but also does not help in the low-shot learning scenario.

Impact of generation. Figure 5 shows the top-5 accuracies on novel classes and on base+novel classes for our generation method applied on top of the baseline representation and the SGM feature representation. The numbers

Representation	Lowshot phase	n=1	2	5	10	20
<i>ResNet-10</i>						
Baseline	Classifier	43.0	54.3	67.2	72.8	75.9
Baseline	Generation* + Classifier	52.4	59.4	67.5	72.6	76.9
SGM*	Classifier	49.4	60.5	<i>71.3</i>	75.8	<i>78.1</i>
SGM*	Generation* + Classifier	<i>54.3</i>	<i>62.1</i>	<i>71.3</i>	75.8	<i>78.1</i>
Batch SGM*	Classifier	49.3	60.5	71.4	<i>75.8</i>	78.5
L1*	Classifier	47.1	58.5	69.2	73.7	76.1
L2*	Classifier	52.7	63.0	71.5	74.8	76.4
Triplets	Classifier	47.6	57.1	65.2	68.4	70.2
Dropout [20]	Classifier	50.1	59.7	68.8	72.7	74.7
Decov [7]	Classifier	43.3	55.7	70.1	75.4	77.9
Multiverse [30]	Classifier	44.1	54.2	67.0	73.2	76.9
Baseline	Data Augmentation	44.9	54.0	66.4	73.0	77.2
Baseline	Model Regression [46]	46.4	56.7	66.8	70.4	72.0
Baseline	Matching Network [45]	55.0	61.5	69.3	73.4	76.2
Baseline-ft	Classifier	41.7	51.7	65.0	71.2	74.5
<i>ResNet-50</i>						
Baseline	Classifier	54.1	67.7	<i>79.1</i>	<i>83.2</i>	85.4
Baseline	Generation* + Classifier	<i>63.1</i>	71.5	78.8	82.6	85.4
SGM*	Classifier	60.0	<i>71.3</i>	80.0	83.3	<i>85.2</i>
SGM*	Generation* + Classifier	63.6	71.5	80.0	83.3	<i>85.2</i>

Table 2: Top-5 accuracy on base and novel classes. Best are bolded and blue; the second best are italicized and red. *Our methods.

are in Tables 1 and 2. Note that we only generate examples when $n < k$, with $k = 20$ for baseline representations and 5 for SGM (see Section 4.2). We observe that the generated examples provide a large gain of over 9 points for $n = 1, 2$ on the novel classes for the baseline representation. When using the SGM representation, the gains are smaller, but significant.

We also compared our generation strategy to common forms of data augmentation (aspect ratio and scale jitter, horizontal flips, and brightness, contrast and saturation changes). Data augmentation only provides small improvements (about 1 percentage point). This confirms that our generation strategy produces more diverse and useful training examples than simple data augmentation.

Comparison to other low-shot methods. We also compared to two recently proposed low-shot learning methods: matching networks [45] and model regression [46]. **Model regression** trains a small MLP to regress from the classifier trained on a small dataset to the classifier trained on the full dataset. It then uses the output from this regressor to regularize the classifier learnt in the low-shot learning phase. **Matching networks** proposes a nearest-neighbor approach that trains embeddings end-to-end for the task of low-shot learning. We apply both these techniques on our baseline representation.

For both these methods, the respective papers evaluated on the novel classes only. In contrast, real-world recognition systems will need to discriminate between data-starved novel concepts, and base classes with lots of data. We adapt these methods to work with both base and novel classes as follows.




Figure 4: **Representation learning comparison.** Top-5 accuracy on ImageNet1k val. Top-performing feature regularization methods reduce the training samples needed to match the baseline accuracy by 2x. Note the different Y-axis scales.




Figure 5: **Comparisons with and without example generation.** Top-5 accuracy on ImageNet1k val. Note the different Y-axis scales.




Figure 6: **Comparison to recently proposed methods.** Top-5 accuracy on ImageNet1k val. Note the different Y-axis scales.

For model regression, we only use the model regressor-based regularization on the novel classes, with the other classifiers regularized using standard weight decay. We use one-vs-all classifiers to match the original work.

Matching networks require the training dataset to be kept in memory during test time. To make this tractable, we use 100 examples per class, with the novel classes correspondingly oversampled.

Comparisons between these methods and our approach are shown in Figure 6. We find that model regression improves significantly over the baseline, but our generation strategy works better for low n . Model regression also hurts overall accuracy for high n .

Matching networks work very well on novel classes. In terms of overall performance, they perform better than our generation approach on top of the baseline representation, but worse than generation combined with the SGM representation, especially for $n > 2$. Further, matching networks are based on nearest neighbors and keep the entire training set in memory, making them much more expensive at test time than our simple linear classifiers.

Deeper networks. We also evaluated our approach on the ResNet-50 network architecture to test if our conclusions extend to deeper convnets that are now in use (Tables 1 and 2). First, even with the baseline representation and without any generation we find that the deeper architecture also leads to improved performance in all low-shot scenarios. However, our SGM loss and our generation strategy further improve this performance. Our final top-5 accuracy on novel classes is still more than 8 points higher for $n = 1, 2$, and our overall accuracy is about 3 points higher, indicating that our contributions generalize to deeper and better models.

7. Conclusion

This paper proposes (1) a low-shot recognition benchmark of realistic complexity, (2) the squared gradient magnitude (SGM) loss that encodes the end-goal of low-shot learning, and (3) a novel way of transferring modes of variation from base classes to data-starved ones. Source code and models are available at: <https://github.com/facebookresearch/low-shot-shrink-hallucinate>.

References

- [1] P. Agrawal, R. Girshick, and J. Malik. Analyzing the performance of multilayer neural networks for object recognition. In *ECCV*, 2014. 2
- [2] H. Azizpour, A. Sharif Razavian, J. Sullivan, A. Maki, and S. Carlsson. From generic to specific deep representations for visual recognition. In *CVPR*, 2015. 2
- [3] E. Bart and S. Ullman. Cross-generalization: Learning novel classes from a single example by feature replacement. In *CVPR*, 2005. 2
- [4] L. Bertinetto, J. Henriques, J. Valmadre, P. Torr, and A. Vedaldi. Learning feed-forward one-shot learners. In *NIPS*, 2016. 2
- [5] W.-L. Chao, S. Changpinyo, B. Gong, and F. Sha. An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In *ECCV*, 2016. 2
- [6] B. Cheung, J. A. Livezey, A. K. Bansal, and B. A. Olshausen. Discovering hidden factors of variation in deep networks. In *ICLR*, 2015. 5
- [7] M. Cogswell, F. Ahmed, R. Girshick, L. Zitnick, and D. Batra. Reducing overfitting in deep networks by decorrelating representations. In *ICLR*, 2016. 5, 7
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 1, 2
- [9] M. Dixit, R. Kwitt, M. Niethammer, and N. Vasconcelos. AGA: Attribute-Guided Augmentation. In *CVPR*, 2017. 2
- [10] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. DeCAF: A deep convolutional activation feature for generic visual recognition. In *ICML*, 2014. 2
- [11] M. Elhoseiny, B. Saleh, and A. Elgammal. Write a classifier: Zero-shot learning using purely textual descriptions. In *ICCV*, 2013. 2
- [12] A. Farhadi, I. Endres, and D. Hoiem. Attribute-centric recognition for cross-category generalization. In *CVPR*, 2010. 2
- [13] L. Fei-Fei, R. Fergus, and P. Perona. One-shot learning of object categories. *TPAMI*, 2006. 1, 2
- [14] M. Fink. Object classification from a single example utilizing class relevance metrics. *NIPS*, 2005. 2
- [15] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017. 2
- [16] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, M. Ranzato, and T. Mikolov. DeViSE: A deep visual-semantic embedding model. In *NIPS*, 2013. 2
- [17] T. Galanti, L. Wolf, and T. Hazan. A theoretical framework for deep transfer learning. *Information and Inference*, 2016. 2
- [18] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, 2006. 2
- [19] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2, 5, 6
- [20] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012. 5, 7
- [21] S. Jetley, B. Romera-Paredes, S. Jayasumana, and P. Torr. Prototypical priors: From improving classification to zero-shot learning. In *BMVC*, 2015. 2
- [22] Y. Jia and T. Darrell. Latent task adaptation with large-scale hierarchies. In *ICCV*, 2013. 2
- [23] G. Koch, R. Zemel, and R. Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML Deep Learning Workshop*, 2015. 2
- [24] A. Krizhevsky, I. Sutskever, and G. Hinton. ImageNet classification with deep convolutional neural networks. In *NIPS*, 2012. 1, 5, 6
- [25] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 2015. 1, 2
- [26] B. M. Lake, R. R. Salakhutdinov, and J. Tenenbaum. One-shot learning by inverting a compositional causal process. In *NIPS*, 2013. 2
- [27] C. H. Lampert, H. Nickisch, and S. Harmeling. Attribute-based classification for zero-shot visual object categorization. *TPAMI*, 2014. 2
- [28] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proc. of the IEEE*, 1998. 2
- [29] J. Lei Ba, K. Swersky, S. Fidler, and R. Salakhutdinov. Predicting deep zero-shot convolutional neural networks using textual descriptions. In *ICCV*, 2015. 2
- [30] E. Littwin and L. Wolf. The multiverse loss for robust transfer learning. *arXiv preprint arXiv:1511.09033*, 2015. 2, 5, 7
- [31] M. Norouzi, T. Mikolov, S. Bengio, Y. Singer, J. Shlens, A. Frome, G. S. Corrado, and J. Dean. Zero-shot learning by convex combination of semantic embeddings. In *ICLR*, 2014. 2
- [32] A. Opelt, A. Pinz, and A. Zisserman. Incremental learning of object detectors using a visual shape alphabet. In *CVPR*, 2006. 2
- [33] M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *CVPR*, 2014. 2
- [34] M. Ranzato. *Unsupervised learning of feature hierarchies*. PhD thesis, New York University, 2009. 5
- [35] S. Ravi and H. Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017. 2
- [36] O. Rippel, M. Paluri, P. Dollar, and L. Bourdev. Metric learning with adaptive density discrimination. In *ICLR*, 2016. 2, 6, 7
- [37] B. Romera-Paredes and P. Torr. An embarrassingly simple approach to zero-shot learning. In *ICML*, 2015. 2
- [38] R. Salakhutdinov, J. Tenenbaum, and A. Torralba. One-shot learning with a hierarchical nonparametric bayesian model. *Unsupervised and Transfer Learning Challenges in Machine Learning*, 2012. 2
- [39] L. A. Schmidt. *Meaning and compositionality as statistical induction of categories and constraints*. PhD thesis, Massachusetts Institute of Technology, 2009. 1

- [40] F. Schroff, D. Kalenichenko, and J. Philbin. FaceNet: A unified embedding for face recognition and clustering. In *CVPR*, 2015. 2
- [41] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. CNN features off-the-shelf: An astounding baseline for recognition. In *CVPR Workshops*, 2014. 2
- [42] R. Socher, M. Ganjoo, C. D. Manning, and A. Ng. Zero-shot learning through cross-modal transfer. In *NIPS*, 2013. 2
- [43] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Web-scale training for face identification. In *CVPR*, 2015. 2
- [44] S. Thrun. Is learning the n-th thing any easier than learning the first? *NIPS*, 1996. 2
- [45] O. Vinyals, C. Blundell, T. P. Lillicrap, K. Kavukcuoglu, and D. Wierstra. Matching networks for one shot learning. *CoRR*, abs/1606.04080, 2016. 2, 7
- [46] Y.-X. Wang and M. Hebert. Learning to learn: Model regression networks for easy small sample learning. In *ECCV*, 2016. 2, 7
- [47] A. Wong and A. L. Yuille. One shot learning via compositions of meaningful patches. In *ICCV*, 2015. 2
- [48] Z. Zhang and V. Saligrama. Zero-shot learning via semantic similarity embedding. In *ICCV*, 2015. 2