

QuEST: Quantized Embedding Space for Transferring Knowledge

Himalaya Jain¹, Spyros Gidaris¹, Nikos Komodakis^{2,3}, Patrick Pérez¹, and
Matthieu Cord^{1,4}

¹Valeo.ai ²University of Crete ³LIGM ⁴Sorbonne University

Abstract. Knowledge distillation refers to the process of training a student network to achieve better accuracy by learning from a pre-trained teacher network. Most of the existing knowledge distillation methods direct the student to follow the teacher by matching the teacher’s output, feature maps or their distribution. In this work, we propose a novel way to achieve this goal: by distilling the knowledge through a *quantized visual words* space. According to our method, the teacher’s feature maps are first quantized to represent the main visual concepts (i.e., visual words) encompassed in these maps and then the student is asked to predict those visual word representations. Despite its simplicity, we show that our approach is able to yield results that improve the state of the art on knowledge distillation for model compression and transfer learning scenarios. To that end, we provide an extensive evaluation across several network architectures and most commonly used benchmark datasets.

Keywords: Knowledge distillation, transfer learning, model compression.

1 Introduction

Knowledge distillation is an interesting learning problem with many practical applications. It was initially introduced by Hinton et al. [13] (KD method) as a means to achieve model compression. The main idea is to train a network using the output of another pre-trained network. Specifically, KD trains a low capacity “student network” (i.e., compressed model) to mimic the *softened classification predictions* of a higher capacity pre-trained “teacher network” (i.e., original uncompressed model). Adding such an auxiliary objective to the standard training loss of the student network, leads to learning a more accurate model. Apart from model compression, knowledge distillation has also been shown to be beneficial to semi-supervised learning [19,31] and transfer learning [20] problems.

Since KD, several other methods have been proposed for knowledge distillation [4,7,12,17,21,24,28,35,39]. A popular paradigm is that of transferring knowledge through intermediate features of the two networks [28,39]. For instance, in FitNet [28] the student is trained to regress the raw feature maps of the teacher. The main intuition is that these features tend to encode useful semantic knowledge for transfer learning. However, empirical evidence [1,11,12,36,39] suggests that



Fig. 1: **Teacher-word clusters from layer4 of ResNet34.** For each cluster we depict the 16 patch members with the smallest Euclidean distance to the cluster centroid. We see that they encode localized mid-to-high level image patterns

naively regressing the teacher features might be a difficult task which over-regularizes the student, thus leading to a suboptimal knowledge transfer solution. This can be because **(a)** the exact feature magnitudes (i.e., feature “details”) of the feature maps are not semantically important aspects of the teacher network in the context of knowledge distillation, and/or **(b)** features behaviour largely depends on the architecture design (e.g., ResNet versus MobileNet) and capacity (i.e., width and depth) of a network. For instance, AT [39] shows that transferring attention maps (feature maps aggregated into a single channel) might be a better option in many cases; FSP [36] opts to remove the spatial information by considering only the correlations among feature channels; AB [12] transfers binarized versions of feature maps; and VID [1] proposes a regression loss (inspired from information theory) that is meant to “ignore” difficult to regress channels.

In our work we propose to address the knowledge transfer task in the context of *a spatially dense quantization of feature maps into visual codes*. Specifically, first we learn with k-means a vocabulary of deep features from the teacher network, which we call visual teacher-words. Then, given an image, we use this vocabulary to quantize/assign each location of the teacher’s feature maps into the closest visual words. As distillation task, the student has to predict these visual teacher-word assignments maps using an auxiliary convolutional module. Hence, our method departs from prior work and essentially *transfers discretized (into words) versions of the feature maps*. Besides being remarkably simple, this strategy has the advantage of aligning the student’s behavior with that of the teacher by focusing only on the main local visual concepts (i.e., teacher-words) that the teacher has learnt to detect in an image (see Figure 1). As a result, it ignores unimportant feature “details” that are difficult to regress, thus avoiding over-regularizing the student. Compared to prior work (e.g., [12,36,39]), we argue that our strategy of transferring visual word representations is more effective as it does not withhold important aspects/cues of the teacher’s feature maps.

We extensively evaluate our method on two knowledge distillation scenarios (i.e., model compression and transfer learning on small-sized datasets), with various pairs of teacher-student architectures, and across a variety of datasets. In all cases our simple but effective approach achieves state-of-the-art results.

2 Related work

The aim of knowledge distillation is to transfer knowledge from a trained teacher network to a student network. To carry out the knowledge transfer, the student is trained to imitate some facets of the teacher network. For instance, Buciluă et al. [3] proposed to train the student as an approximator of a large ensemble of teacher networks by predicting the output of this network ensemble and Hinton et al. [13] proposed to train the student to output the same softened predictions (i.e., classification probabilities) as the trained teacher network.

The teacher and student are deep networks, having a sequential structure that is known to learn hierarchical representations of increasing abstraction. Inspired by this, some methods propose to match the intermediate activation or feature maps of the teacher as the distillation task for the student. This encourages the student to follow the intermediate solutions produced by the teacher. FitNet [28] proposes to train the student with an additional layer to regress the feature maps of the teacher. While AT [39] builds “attention maps” by aggregating feature maps, and ℓ_2 error between normalized attention maps of student and teacher is used as distillation loss. In FSP [36], the difference between Gram matrices of feature maps of the two networks is minimized for distillation.

Another line of works focuses on matching distributions of feature maps rather than feature maps themselves. In NST [15], the maximum mean discrepancy between the distributions of feature maps of teacher and student is minimized. VID [1] and CRD [32] consider maximizing mutual information between the two networks as the knowledge distillation task. Mutual information is maximized by maximizing variational lower bound (VID) or contrastive loss (CRD).

SP [34] departs from matching feature maps or their distributions. This relieves the student network from the demanding task of copying the teacher network, which could be too ambitious given the difference in their capacity and architecture. SP proposes instead a similarity preserving constraint by using difference in pairwise similarity, computed on a mini-batch, as distillation loss. Similar methods are proposed in other recent works RKD [23] and CC [25].

In our work, similar to SP, we propose a different space for distillation than feature space. We propose to use a quantized space where the teacher’s feature maps are encoded by quantization with learned visual teacher-words. These visual words are learned by k-means clustering on the feature maps of the teacher thus, they represent useful semantic concepts. With the proposed quantized space we concentrate more on the important semantic concepts and their spatial correlation for knowledge distillation.

We note that quantizing features into visual words is key a ingredient of the bag-of-visual-word techniques that were extensively used in the past [5,16,26,30,33]. Although these visual-word techniques have now been used with deep learning [2,8,9], our work is the first to leverage a visual-word strategy with deep learning for knowledge distillation.

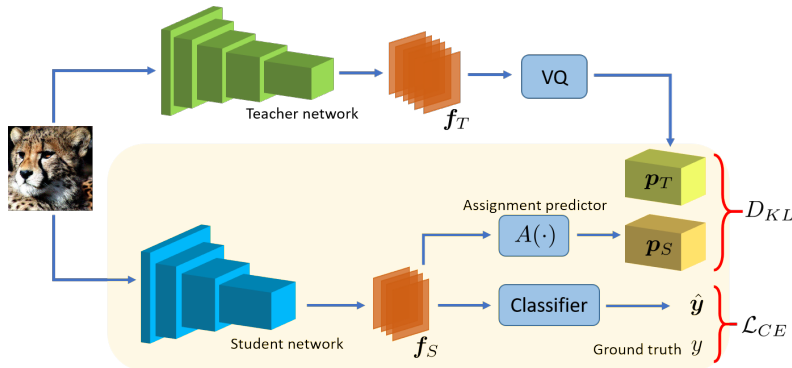


Fig. 2: Overview of the proposed method for knowledge transfer based on visual word assignments. The feature maps f_T of the teacher are converted into a visual words representation p_T via vector quantization (VQ) with soft-assignment. Then, as distillation task, the student has to predict the teacher’s words assignments p_T using its own feature maps f_S and an auxiliary assignment prediction module A . The distillation loss is the KL-divergence D_{KL} of the predicted word assignments p_S with p_T . The visual word vocabulary that is used in VQ is learned off-line by applying the k-means clustering algorithm to teacher feature maps extracted from the training dataset

3 Approach

3.1 Preliminaries

Here we briefly explain the learning setting of knowledge distillation methods. Let S be a student network that we want to train using a dataset of (\mathbf{x}, y) examples, where \mathbf{x} is the image and y is the label. In the standard supervised learning setting the teacher S is trained to minimise for each example (\mathbf{x}, y) the classification loss $\mathcal{L}_{CLS}(S) = \mathcal{L}_{CE}(y, \sigma(\mathbf{z}_S))$, where \mathcal{L}_{CE} is the cross-entropy loss, $\mathbf{z}_S = S(\mathbf{x})$ are the classification logits predicted from S for the image \mathbf{x} , and $\sigma(\cdot)$ is the softmax function. Knowledge distillation methods assume that there is available a pre-trained teacher network T , which has higher capacity than that of the student S . The goal is to exploit the pre-trained knowledge of teacher T for training a better (i.e., more accurate) student S than the standard supervised learning setting. To that end, they define for each training example (\mathbf{x}, y) an additional loss term $\mathcal{L}_{DIST}(S; T)$ based on the pre-trained teacher T . For instance, in the seminal work of Hinton et al.[13], the additional distillation loss is the cross entropy of the two network output softmax distributions for the same training example (\mathbf{x}, y) :

$$\mathcal{L}_{DIST}(S; T) = \rho^2 \mathcal{L}_{CE}\left(\sigma\left(\frac{\mathbf{z}_T}{\rho}\right), \sigma\left(\frac{\mathbf{z}_S}{\rho}\right)\right), \quad (1)$$

where $\mathbf{z}_T = T(\mathbf{x})$ are the logits produced from T for the image \mathbf{x} , and $\rho > 0$ is a temperature that softens (i.e., lowers the peakiness) of the two softmax

distributions. Through this loss, the student learns from the teacher by mimicking the teacher’s output on the training data. Therefore, the final objective that the student S has to minimize per training example (\mathbf{x}, y) is:

$$\mathcal{L} = \alpha \mathcal{L}_{CLS}(S) + \beta \mathcal{L}_{DIST}(S; T), \quad (2)$$

where α and β are the weights of the two loss terms.

In the next subsection we explain the distillation loss $\mathcal{L}_{DIST}(S; T)$ that we propose in our work.

3.2 Distilling visual teacher-word assignments

An overview of our approach is provided in Figure 2. The proposed distillation task first converts the teacher feature maps into visual word assignments and then trains the student to predict those assignments from its own feature maps.

Converting the teacher feature maps into visual word assignments.

Given an image \mathbf{x} , let $\mathbf{f}_T \in \mathbb{R}^{C_T \times H_T \times W_T}$ be the feature map (with C_T channels and $H_T \times W_T$ spatial dimensions) that T generates at one of its hidden layers, and $\mathbf{f}_T^{(h,w)} \in \mathbb{R}^{C_T}$ be the feature vector at the location (h, w) of this feature map. We quantize \mathbf{f}_T in a spatially dense way using a predefined vocabulary $V = \{\mathbf{v}_k\}_{k=1\dots K}$ of K C_T -dimensional visual teacher-word embeddings. Specifically, we first compute for each location (w, h) the squared Euclidean distances between the feature vector $\mathbf{f}_T^{(h,w)}$ and the K teacher-words:

$$\mathbf{d}^{(h,w)} = \left[\left\| \mathbf{v}_k - \mathbf{f}_T^{(h,w)} \right\|_2^2 \right]_{k=1\dots K}. \quad (3)$$

Then, using $\mathbf{d}^{(h,w)}$ we compute the K -dimensional soft-assignment vector $\mathbf{p}_T^{(h,w)}$, which lies on a K -dimensional probability simplex, as:

$$\mathbf{p}_T^{(h,w)} = \sigma(-\mathbf{d}^{(h,w)}/\tau), \quad (4)$$

where $\tau > 0$ is a temperature value used for controlling the softness of the assignment. The visual teacher-word vocabulary V is learned off-line by k-means clustering feature vectors $\mathbf{f}_T^{(h,w)}$ extracted from the available training dataset. As we see in Figure 1, the vocabulary V includes the main local image patterns that are captured by the teacher’s feature maps (e.g., a dog’s leg on grass, the eye of a tiger, or a clock tower). Hence, $\mathbf{p}_T^{(h,w)}$ essentially encodes what semantic concepts (defined by V) the teacher network detects at location (h, w) of its feature map.

Soft-assignment vs. hard-assignment. In our experiments we noticed that the optimal τ value produces very peaky soft-assignments, e.g., the softmax probability for the closest visual teacher-words is on average around 0.996, which is very close to hard assignment. However, very peaky soft-assignment still leads to slightly better students, which can be attributed to either **(a)** the fact that, as in KD [13], the probability mass for the remaining visual teacher-words carries useful teacher knowledge, or **(b)** soft-assignments regularize the visual teacher-words prediction task that the student has to perform [37].

Predicting the teacher’s visual word assignments. The distillation task that we propose is to train the student S to predict the teacher soft-assignment map \mathbf{p}_T (of visual words) based on its own feature map $\mathbf{f}_S \in \mathbb{R}^{C_S \times H_S \times W_S}$. As \mathbf{p}_T encodes semantic information with spatial structure, we posit that, to predict \mathbf{p}_T the student must build an understanding of semantics similar to the teacher. We assume for simplicity that the student’s feature map \mathbf{f}_S has the same spatial size as that of the teacher¹, i.e., $H_S = H_T$ and $W_S = W_T$. We also note that we impose *no constraint on the number of student feature channels* C_S . Therefore, to predict \mathbf{p}_T , we use an *assignment predictor* A , which consists of one cosine-similarity-based convolutional layer with 1×1 kernel size followed by a learnable scaling factor and a softmax function. Specifically, at each location (h, w) , the student predicts the K -dimensional soft-assignment vector $\mathbf{p}_S^{(h,w)}$ from its feature vector $\mathbf{f}_S^{(h,w)}$, as:

$$\mathbf{p}_S^{(h,w)} = A(\mathbf{f}_S^{(h,w)}) = \sigma \left(\left[\gamma \frac{\langle \mathbf{W}_i, \mathbf{f}_S^{(h,w)} \rangle}{\|\mathbf{W}_i\| \|\mathbf{f}_S^{(h,w)}\|} \right]_{i=1 \dots K} \right), \quad (5)$$

where γ is the learnable scaling factor, $\mathbf{W} \in \mathbb{R}^{C_S \times K}$ are the parameters of the convolutional layer, and \mathbf{W}_i is the i -th column of \mathbf{W} . Hence, \mathbf{W} plays essentially the role of a new student vocabulary, and soft assignment is done according to cosine similarity (instead of Euclidean distance).

Finally, we define our distillation loss $\mathcal{L}_{DIST}(S; T)$ (for a training example) as the summation over all the locations (h, w) of the KL-divergence D_{KL} of the predicted soft-assignment distribution $\mathbf{p}_S^{(h,w)}$ from the target distribution $\mathbf{p}_T^{(h,w)}$:

$$\mathcal{L}_{DIST}(S; T) = \sum_{h,w} D_{KL}(\mathbf{p}_T^{(h,w)} \parallel \mathbf{p}_S^{(h,w)}) , \quad (6)$$

where $D_{KL}(\mathbf{p}_T^{(h,w)} \parallel \mathbf{p}_S^{(h,w)})$ consists of the cross-entropy of $\mathbf{p}_T^{(h,w)}$ and $\mathbf{p}_S^{(h,w)}$, and the entropy of $\mathbf{p}_T^{(h,w)}$. The latter is independent to the student and thus does not actually affect the training. We minimize the total loss $\mathcal{L} = \alpha \mathcal{L}_{CLS}(S) + \beta \mathcal{L}_{DIST}(S; T)$ of the student network over W , γ , and the parameters of S . Note that it is possible to apply the proposed distillation loss to more than one feature levels. In this case the final distillation loss is the sum of all the per-level losses.

3.3 Discussion

Transferring “discretized” representations. In our method, the teacher’s feature maps are essentially discretized into visual words (i.e., as already described, the softmax probability for the closest visual teacher-words is on average around 0.996). Therefore, by minimizing the KL-divergence between the predicted visual word assignments \mathbf{p}_S and the teacher assignments \mathbf{p}_T , we essentially “discretize”

¹ Otherwise, we simply down-sample the biggest feature map to the size of the smaller one with an (adaptive) average pooling layer.

the student feature maps according to the “discretization” at the teacher side. The difference is that in the student case, the visual vocabulary is implicitly defined by the parameters W of the assignment predictor A , and the assignment is done via cosine similarity instead of Euclidean distance.² Hence, the student feature maps must learn to discriminate (i.e., to cluster together) the same local image patterns (represented by the visual words) as the teacher feature maps. This way, the student learns to align its feature maps with those of the teacher network both spatially and semantically, but without being forced to regress the exact feature “details” of its teacher (which can over-regularize the student).

Posing feature transfer as a classification problem. In a sense, our method converts the direct feature regression problem into a classification one where the teacher-words are the classification prototypes. As has been shown in prior works (e.g., see methods based on contrastive learning [22,32]), this leads to a training task that is easier to solve/optimize and, as a result, constitutes a far more effective learning strategy.

Relation to KD method [13]. Our method relates to KD in the sense that it also tries to “mimic” softened classification predictions. However, in our case those classification predictions are spatially dense and over the visual teacher-words, which represent visual concepts that are not only more localized but also much more numerous than the available image classes. Due to this fact, our distillation loss leads to a more efficient (i.e., richer) knowledge transfer. We should note at this point that many prior methods combine their proposed distillation loss with the KD loss. In our case, however, we empirically observed that adding the KD loss to our method is not required as it does not offer any significant performance improvement, which further verifies the effectiveness of our approach.

4 Experiments

In this section, we experimentally evaluate our proposed approach and compare it with several state-of-the art knowledge distillation methods. In the remainder of this section, we first compare our approach against prior art with extensive experiments on knowledge distillation in Section 4.1 followed by results on transfer learning in Section 4.2. Finally, in Section 4.3, we analyse and discuss the impact of various hyper-parameters and design choices of our approach.

4.1 Comparison with prior work

We extensively evaluate our approach on three different datasets: ImageNet [6], CIFAR-100 [18] and CIFAR-10. We start with ImageNet, as this is the most challenging and interesting dataset among the three. Then, on CIFAR-100 we conduct an extensive evaluation on many network architectures and compare against several state-of-the-art methods. Finally, we conclude with CIFAR-10.

² In Sec. A of supplementary, we discuss the use of cosine similarity on student side.

Table 1: **Evaluation on ImageNet.** Top-1 and Top-5 error rate of student network on ImageNet validation set, for two teacher-student combinations, ResNet34 (21.8M) to ResNet18 (11.7M) and ResNet50 (25.6M) to MobileNet (4.2M). The results for other methods are taken from [32] and [11]. FT refers to [17]

| | ResNet34 | ResNet18 | KD | AT+KD | SP | CC | CRD | CRD+KD | Ours |
|-------|----------|----------|-------|-------|-------|-------|-------|--------|--------------|
| Top-1 | 26.69 | 30.25 | 29.34 | 29.30 | 29.38 | 30.04 | 28.83 | 28.62 | 28.33 |
| Top-5 | 8.58 | 10.93 | 10.12 | 10.00 | 10.20 | 10.83 | 9.87 | 9.51 | 9.33 |

| | ResNet50 | MobileNet | KD | AT+KD | FT | AB+KD | OFD | Ours |
|-------|----------|-----------|-------|-------|-------|-------|-------|--------------|
| Top-1 | 23.87 | 31.13 | 31.42 | 30.44 | 30.12 | 31.11 | 28.75 | 27.46 |
| Top-5 | 7.14 | 11.24 | 11.02 | 10.67 | 10.50 | 11.29 | 9.66 | 8.87 |

Table 2: **Experiments on CIFAR-100.** Top-1 test accuracy of student networks for various student-teacher combinations and knowledge distillation methods. The pre-trained teacher models are taken from [32] and for all methods except ours we use the results reported in [32] for a fair comparison. We report average over 5 runs as in [32]

| Model | | Knowledge distillation methods | | | | | | | | | | | |
|------------|-----------|--------------------------------|--------|-------|-------|-------|-------|-------|-------|--------------|--------------|---------|---------|
| Teacher | Student | KD | FitNet | AT | AB | FSP | SP | VID | CRD | CRD+KD | Ours | teacher | student |
| WRN-40-2 | WRN-16-2 | 74.92 | 73.58 | 74.08 | 72.50 | 72.91 | 73.83 | 74.11 | 75.48 | 75.64 | 76.10 | 75.61 | 73.26 |
| WRN-40-2 | WRN-40-1 | 73.54 | 72.24 | 72.77 | 72.38 | - | 72.43 | 73.30 | 74.14 | 74.38 | 74.58 | 75.61 | 71.98 |
| ResNet56 | ResNet20 | 70.66 | 69.21 | 70.55 | 69.47 | 69.95 | 69.67 | 70.38 | 71.16 | 71.63 | 71.84 | 72.34 | 69.06 |
| ResNet110 | ResNet20 | 70.67 | 68.99 | 70.22 | 69.53 | 70.11 | 70.04 | 70.16 | 71.46 | 71.56 | 71.89 | 74.31 | 69.06 |
| ResNet110 | ResNet32 | 73.08 | 71.06 | 72.31 | 70.98 | 71.89 | 72.69 | 72.61 | 73.48 | 73.75 | 74.08 | 74.31 | 71.14 |
| ResNet32x4 | ResNet8x4 | 73.33 | 73.50 | 73.44 | 73.17 | 72.62 | 72.94 | 73.09 | 75.51 | 75.46 | 75.88 | 79.42 | 72.50 |
| VGG13 | VGG8 | 72.98 | 71.02 | 71.43 | 70.94 | 70.23 | 72.68 | 71.23 | 73.94 | 74.29 | 73.81 | 74.64 | 70.36 |

For all the experiments we use $\alpha = 1$, $\beta = 1$, and for ImageNet and CIFAR-100 $K = 4096$, $\tau = 0.2$ while for CIFAR-10 $K = 256$, $\tau = 0.005$. We apply our distillation loss on the feature maps of the last convolutional layer. For complete implementation details, see Sections B and C of the supplementary material.

ImageNet results In Table 1 we provide results for the ImageNet dataset, which contains 1.28M training images and 50K test images, over 1000 semantic classes. It is much more challenging than CIFAR-100 and CIFAR-10, which makes it the ultimate benchmark for evaluating distillation methods.

We evaluate our approach on it with two teacher-student combinations: ResNet34 to ResNet18 (same architecture design) and ResNet50 to MobileNet [14] (different architecture design). We observe that our distillation method reduces the Top-1 error of ResNet18 from 30.25% to 28.33% and MobileNet from 31.13% to 27.46%, which sets the new state-of-the-art on this challenging dataset.

Table 3: **Distillation between different architectures.** Top-1 test accuracy on CIFAR100 dataset of the student networks. The student models are learned with knowledge distillation from a teacher with different architecture. The table compares various distillation methods. Similar to Table 2, we use the pre-trained teacher networks provided by [32] and the results for other methods are also taken from [32]. Following the protocol of [32], we report average over 3 runs

| Model | | Knowledge distillation methods | | | | | | | | | | |
|------------|--------------|--------------------------------|--------|-------|-------|-------|-------|-------|--------------|--------------|---------|---------|
| Teacher | Student | KD | FitNet | AT | AB | SP | VID | CRD | CRD+KD | Ours | teacher | student |
| VGG13 | MobileNetV2 | 67.37 | 64.14 | 59.4 | 66.06 | 66.3 | 65.56 | 69.73 | 69.94 | 68.79 | 74.64 | 64.6 |
| ResNet50 | MobileNetV2 | 67.35 | 63.16 | 58.58 | 67.20 | 68.08 | 67.57 | 69.11 | 69.54 | 69.81 | 79.34 | 64.6 |
| ResNet50 | VGG8 | 73.81 | 70.69 | 71.84 | 70.65 | 73.34 | 70.30 | 74.30 | 74.58 | 75.17 | 79.34 | 70.36 |
| ResNet32x4 | ShuffleNetV1 | 74.07 | 73.59 | 71.73 | 73.55 | 73.48 | 73.38 | 75.11 | 75.12 | 76.28 | 79.42 | 70.50 |
| ResNet32x4 | ShuffleNetV2 | 74.45 | 73.54 | 72.73 | 74.31 | 74.56 | 73.4 | 75.65 | 76.05 | 77.09 | 79.42 | 71.82 |
| WRN-40-2 | ShuffleNetV1 | 74.83 | 73.73 | 73.32 | 73.34 | 74.52 | 73.61 | 76.05 | 76.27 | 76.75 | 75.61 | 70.50 |

For ResNet34 to ResNet18, we use the results reported in CRD [32] for all other methods. Our approach outperforms all including CRD+KD which is second to ours with 0.30 higher error rate. For ResNet50 to MobileNet, we follow OFD [11]. We outperform OFD with 1.29% lower error rate, which is a very significant improvement. Note that OFD extensively studied distillation with feature regression and carefully picks the location for feature regression, uses a modified ReLU (margin-ReLU) activation function, and a more robust (partial L_2) distance loss. Nevertheless, our simple approach outperforms this well-engineered direct feature regression method.

Note that, due to the time-consuming nature of the ImageNet experiments, we did not try to tune the hyper-parameters of our method in this case (instead we reused the ones chosen for CIFAR-100). As a result, a further reduction of the error rates of the student might very well be possible by a more proper adjustment of the hyper-parameters.

CIFAR-100 results The CIFAR-100 dataset is one of the most commonly used dataset for evaluating knowledge distillation methods. It consists of small 32×32 resolution images and 100 semantic classes. It has 50K images in the training set and 10K in the test set which are evenly distributed across the semantic classes.

In Tables 2 and 3 we provide an exhaustive evaluation of our method under many different network architectures. Also, we compare against several prior methods, i.e., KD [13], FitNet [28], AT [39], AB [12], FSP [36], SP [34], VID [1], and CRD [32]. For network sizes and compression rates see §D of supplementary.

Knowledge transfer between networks with the same architecture design. In Table 2, we compare knowledge distillation methods with teacher and student networks that have the same architecture design but different depth or width (width refers to the number of channels per layer). For the network architectures, we evaluate knowledge distillation between WideResNet [38], ResNet [10] and VGG [29].

Table 4: **CIFAR-10 experiments.** Top-1 error of student networks on CIFAR-10. We use the results reported in [34] and following it, use median of 5 runs

| Model | | Knowledge Distillation methods | | | | | |
|-------------------|-------------------|--------------------------------|-------------|-------------|-------------|---------|---------|
| Teacher (#params) | Student (#params) | KD | AT+KD | SP | Ours | teacher | student |
| WRN-40-1 (0.56M) | WRN-16-1 (0.17M) | 8.48 | 8.30 | 8.13 | 8.02 | 6.51 | 8.74 |
| WRN-16-2 (0.69M) | WRN-16-1 (0.17M) | 7.94 | 8.28 | 7.52 | 7.55 | 6.07 | 8.74 |
| WRN-40-2 (2.24M) | WRN-16-2 (0.69M) | 6.00 | 5.89 | 5.52 | 5.56 | 5.18 | 6.07 |
| WRN-16-8 (11M) | WRN-16-2 (0.69M) | 5.62 | 5.47 | 5.34 | 5.06 | 4.24 | 6.07 |
| WRN-16-8 (11M) | WRN-40-2 (2.24M) | 4.86 | 4.47 | 4.55 | 4.48 | 4.24 | 5.18 |

Our method outperforms all the other methods on all the different teacher-student combinations, except for the *VGG13 to VGG8* case where we are second only to CRD [32] a contemporary method to ours. On average we improve by an absolute 2.97% over students without distillation, *i.e.* relatively 16.9% more than CRD (2.54% over student) and 7.6% more than CRD+KD (2.76%). Moreover, in some cases our students achieve accuracy either very close to that of the teacher (*ResNet56 to ResNet20* and *ResNet110 to ResNet32*), or even exceeds it (*WRN-40-2 to WRN-40-1*), which further confirms the potency of our method.

Knowledge transfer between networks with different architectures. In Table 3, we evaluate the merit of our approach for distillation between different network designs. All of them but one have different spatial dimensions between feature maps of student and teacher. For example, MobileNetV2 at the penultimate layer has 2×2 feature maps while the teachers, VGG13 and ResNet50, have 4×4 . Similarly, ShuffleNetV1/V2 have 4×4 while ResNet32x4 and WRN-40-2 have 8×8 . As the soft-assignment maps \mathbf{p}_T and \mathbf{p}_S should have the same spatial dimension to apply our distillation loss, we do average pooling on the feature maps of teacher before quantization, as explained in section §3.2.

We observe that our approach outperforms every other method on all but one experiment. In the cases of the ResNet32x4 teacher, we notice an improvement of more than 1% against the other methods. Overall, the proposed approach improves by an average of 5.25% on the student without distillation. While the most competitive CRD and CRD+KD bring respectively an improvement of 4.59% and 4.85%. Note that, in terms of average gain over the student without distillation, we get a relative improvement of 14.29% and 8.21% compared to CRD and CRD+KD respectively.

CIFAR-10 results The CIFAR-10 dataset is similar to CIFAR-100 with the only difference that there are now 10 semantic classes instead of 100.

Comparison with prior work. In Table 4, we compare our approach with KD, AT and SP in terms of error rate on CIFAR-10 test set. We consider distillation between WideResNet student and teacher with different depth and/or width.

Table 5: **Results for transfer learning.** The teacher is pre-trained on ImageNet while the student is trained from scratch on a target dataset with cross-entropy loss and distillation loss. The results for other methods are reported from [1]

| (a) MIT-67, ResNet34 to ResNet18 | | | | | | (b) MIT-67, ResNet34 to VGG-9 | | | | | |
|----------------------------------|---|--------------|--------------|--------------|--------------|-------------------------------|---|--------------|--------------|--------------|--------------|
| | M | 80 | 50 | 25 | 10 | | M | 80 | 50 | 25 | 10 |
| student | | 48.13 | 37.69 | 27.01 | 14.25 | student | | 53.58 | 43.96 | 29.70 | 15.97 |
| fine-tuning | | 70.97 | 66.04 | 58.13 | 47.91 | fine-tuning | | 65.97 | 58.51 | 51.72 | 39.63 |
| LwF | | 63.43 | 51.79 | 41.04 | 22.76 | LwF | | 60.90 | 52.01 | 41.57 | 27.76 |
| FitNet | | 71.34 | 60.45 | 54.78 | 36.94 | FitNet | | 70.90 | 64.70 | 54.48 | 40.82 |
| AT | | 58.21 | 48.66 | 43.66 | 27.01 | AT | | 60.90 | 52.16 | 42.76 | 25.60 |
| NST | | 55.52 | 46.34 | 33.21 | 20.82 | NST | | 55.60 | 46.04 | 35.22 | 21.64 |
| VID-LP | | 67.91 | 58.51 | 47.09 | 31.94 | VID-LP | | 68.88 | 61.64 | 50.22 | 39.25 |
| VID-I | | 71.34 | 63.66 | 60.07 | 50.97 | VID-I | | 72.01 | 67.01 | 59.33 | 45.90 |
| LwF+FitNet | | 70.97 | 60.37 | 54.48 | 38.73 | LwF+FitNet | | 70.52 | 64.10 | 54.63 | 40.15 |
| VID-LP+VID-I | | 71.87 | 65.75 | 61.79 | 50.37 | VID-LP+VID-I | | 71.72 | 66.49 | 58.96 | 45.89 |
| Ours | | 73.18 | 69.40 | 62.71 | 50.92 | Ours | | 71.92 | 67.79 | 60.10 | 47.99 |

Again, our method achieves state-of-the-art results on CIFAR-10. Specifically, we outperform the other methods on two settings (*WRN-40-1* to *WRN-16-1* and *WRN-16-8* to *WRN-16-2*), while we achieve almost the same results on the other three settings with statistically negligible difference of less than 0.04%. Our approach achieves an average reduction in error rate of 0.82% compared to the student without distillation, while the most competitive SP method gets 0.75% followed by AT with 0.48%.

4.2 Transfer learning to small-sized datasets

In this section, we evaluate our method for transfer learning. In transfer learning, the objective is to train a student on a small target dataset with the aid of a teacher which is pre-trained on a large dataset. For our experimental evaluation, we use ResNet34 pre-trained on ImageNet as the teacher and, ResNet18 and VGG-9 [29] as the student network. We use MIT-67 [27] as the target dataset. It contains 15,620 indoor scene images, classified into 67 classes. Following the evaluation protocol of [1], we sub-sample the training set with $M = \{80, 50, 25, 10\}$ images per classes. This is to assess the performance at various levels of availability of the training data. The student is trained from scratch on target data with cross-entropy loss and distillation loss. For all the transfer learning experiments we use $K = 4096$ visual words which is learned on ImageNet, where the teacher network was pre-trained.

In case of transfer learning we found that it is better to apply the proposed loss at the last two layers (layer3 and layer4) of ResNet34. For student network ResNet18 we use the same, layer3 and layer4, and for VGG-9 we use the last two max-pool layers. Exactly the same two feature levels are being used by the competing methods in this section.

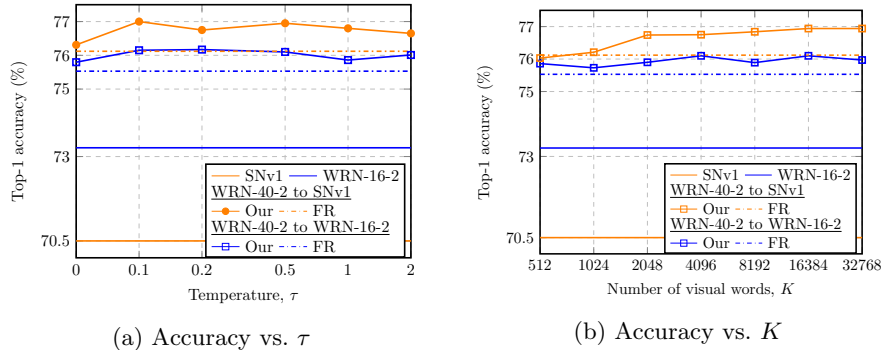


Fig. 3: **Effect of varying temperature τ and visual words K .** Top-1 accuracy on CIFAR-100 of students trained with proposed distillation loss with varying hyperparameters τ and K . The graphs show performance of two students, WRN-16-2 and ShuffleNetv1 (SNv1) with WRN-40-2 as teacher. (a) Effect of temperature τ . (b) Effect of varying number of visual words. The straight lines refer to the student trained without distillation (solid line) or with feature regression (FR, dash-dotted line)

The results for transfer learning experiments are given in Table 5. The table compares our method to several distillation approaches including FitNet, AT, NST VID. In the table, LwF refers to learning without forgetting [20], VID-LP and VID-I are VID [1] loss on logits and on intermediate features respectively, while VID-LP+VID-I uses both. Our proposed method outperforms all the methods on 3 out of 4 configurations of M for both the students. While being second only to VID-I with a statistically insignificant difference of less than 0.1%.

We also give results with the fine-tuning case, i.e., pre-training on ImageNet and then fine-tuning on MIT-67 without any distillation loss. Fine-tuning is a very strong baseline and standard method for transfer learning. However, it is not very practical since it requires pre-training the student on a large-size dataset (which is computationally expensive and needs access and storage of the dataset). In our experiments, the proposed method performs better than even fine-tuning.

4.3 Further analysis

Here we analyse several aspects of our method for the model compression scenario.

Impact of temperature value τ . In Figure 3a we plot how the temperature value τ of the soft-quantization of the teacher feature maps in Eq. (4) affects the performance of our method. We provide results for two teacher-student configurations, WRN-40-2 to WRN-16-2 and WRN-40-2 to ShuffleNetV1, and 5 different τ values, 0.1, 0.2, 0.5, 1.0, and 2.0. We also provide results for the hard-assignment case, which we denote with the $\tau = 0$ in the plot. We observe that choosing a small τ value, which means a more peaky soft-assignment (i.e., closer to

Table 6: **Distillation at different layers.** Comparing student networks with the proposed approach when applied at different layers. Block L corresponds to the penultimate layer (before classification layer) while Block $L - 1$ refers to the layer at a block before Block L . We report the Top-1 accuracy

| CIFAR-100 | | CIFAR-10 | |
|-------------------|--------------|-------------------|--------------|
| WRN-40-2→WRN-16-2 | | WRN-40-1→WRN-16-1 | |
| Block $L - 1$ | 74.16 | Block $L - 1$ | 91.64 |
| Block L | 76.10 | Block L | 91.98 |

hard-assignment), leads to better distillation performance. However, going to the extreme case of hard-assignment (i.e., $\tau = 0$) drops the distillation performance. Hence, very peaky soft-assignment achieves better knowledge transfer than the hard-assignment case.

Impact of vocabulary size K . In Figure 3b we plot the distillation performance of our method as a function of the vocabulary size K on CIFAR-100. We observe a relatively stable performance for $K \geq 2048$. In our CIFAR-10 experiments we noticed that K between 128 to 256 was leading to better performance³. Therefore, it seems that the number of visual words K depends on the number of classes, preferably it should be sufficiently larger than the number of classes.

At which feature level to apply the distillation loss? In Table 6 we measure the performance of our method when the distillation loss is applied (1) to the last feature level of the teacher-student networks (which is what we have used for all other experiments on the model compression scenario), and (2) to the feature level of the previous down-sampling stage (i.e., the output of the 2nd residual block in WRN-40-2/WRN-16-2 or WRN-40-1/WRN-16-1). We report results on CIFAR-100 and CIFAR-10. In all cases, switching to the feature maps of the previous down-sampling stage leads to a drop in performance. Therefore, our proposed distillation loss appears to perform best with the last feature level.

Comparison with direct feature regression. In Table 7, we explicitly compare our approach against the baseline method of directly regressing raw feature maps (Feature regression). To that end, we apply the Feature regression distillation loss on the same feature layer as in our method, i.e., the last convolutional layer. Also, we train the student network under the same conditions as in our method. In the upper part of the table, we evaluate distillation performance with different capacity networks. In this case, we add a 1×1 convolutional regression layer on top of the chosen layer for distillation of the student network. This is needed to match feature dimension of the teacher. Not surprisingly, our approach outperforms Feature regression for all the teacher-student pairs. Further, we also evaluate on distillation between same architecture and same capacity networks.

³ See Fig. 1 in Sec. E of supplementary for performance-vs.- K plot on CIFAR-10.

Table 7: **Comparison with direct feature regression on CIFAR-100.** Upper part of the table shows distillation performance between different capacity networks, thus we use a regression layer on student for the Feature regression method. In the lower part, distillation is between same architecture and same capacity networks thus we do not use any additional layer to regress. Relative gain refers to the relative improvement of our method over Feature regression

| Teacher | Student | student | Feature regression | Ours | relative gain (%) |
|------------|--------------|---------|--------------------|--------------|-------------------|
| ResNet50 | VGG8 | 70.36 | 71.9 | 75.17 | 4.55 |
| WRN-40-2 | WRN-16-2 | 73.26 | 75.53 | 76.10 | 0.75 |
| WRN-40-2 | ShuffleNetV1 | 70.50 | 75.89 | 76.75 | 1.13 |
| ResNet32x4 | ResNet8x4 | 72.50 | 74.12 | 75.88 | 2.37 |
| ResNet32x4 | ShuffleNetV2 | 71.82 | 75.58 | 77.09 | 2.00 |
| ResNet56 | ResNet20 | 69.06 | 71.56 | 71.84 | 0.39 |
| ResNet50 | ResNet50 | 79.34 | 79.12 | 80.58 | 1.85 |
| WRN40-2 | WRN40-2 | 75.61 | 78.26 | 77.96 | -0.38 |
| ResNet110 | ResNet110 | 74.31 | 75.52 | 76.15 | 0.83 |
| ResNet32x4 | ResNet32x4 | 79.42 | 80.57 | 80.53 | -0.05 |
| ResNet56 | ResNet56 | 72.34 | 73.64 | 74.77 | 1.53 |

This is an ideal case for Feature regression as there is no burden on student to mimic a higher capacity network and also no need for additional regression layer. Even in this ideal setup for Feature regression we outperform it for three architectures while being marginally lower on the other two; overall, our method achieves an average relative gain of 0.76% over Feature regression.

5 Conclusions

Our work deals with the important learning problem of knowledge distillation. The goal of knowledge distillation is to improve the accuracy of a student network by exploiting the learned knowledge of a pre-trained teacher network. To that end, we follow the common paradigm of transferring the knowledge encoded on the learned teacher features. However, instead of performing the distillation task in the context of the initial feature space of the teacher network, we transform it to a new quantized space.

Specifically, our distillation method first densely quantizes the teacher feature maps into visual words and then trains the student to predict this quantization based on its own feature maps. By solving this task the student is forced to align its feature maps with those of the teacher network while ignoring unimportant “feature” details, thus facilitating efficient knowledge transfer between the two networks. To demonstrate the effectiveness of our distillation method, we exhaustively evaluate it on two very common knowledge distillation scenarios, model compression (i.e., on ImageNet, CIFAR-100 and CIFAR-10 datasets) and transfer learning to small-sized datasets (i.e., ImageNet to MIT-67 datasets), and across a variety of deep network architectures. Despite its simplicity, our method manages to surpass prior work, achieving new state-of-the-art results on a variety of benchmarks.

References

1. Ahn, S., Hu, S.X., Damianou, A., Lawrence, N.D., Dai, Z.: Variational information distillation for knowledge transfer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2019) 1, 2, 3, 9, 11, 12
2. Arandjelovic, R., Gronat, P., Torii, A., Pajdla, T., Sivic, J.: NetVLAD: CNN architecture for weakly supervised place recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2016) 3
3. Buciluă, C., Caruana, R., Niculescu-Mizil, A.: Model compression. In: Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 535–541. ACM (2006) 3
4. Cho, J.H., Hariharan, B.: On the efficacy of knowledge distillation. In: Proceedings of the IEEE International Conference on Computer Vision (October 2019) 1
5. Csurka, G., Dance, C., Fan, L., Willamowski, J., Bray, C.: Visual categorization with bags of keypoints. In: ECCV Workshops (2004) 3
6. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 248–255. Ieee (2009) 7
7. Furlanello, T., Lipton, Z., Tschannen, M., Itti, L., Anandkumar, A.: Born-again neural networks. In: International Conference on Machine Learning. pp. 1602–1611 (2018) 1
8. Gidaris, S., Bursuc, A., Komodakis, N., Pérez, P., Cord, M.: Learning representations by predicting bags of visual words. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6928–6938 (2020) 3
9. Girdhar, R., Ramanan, D., Gupta, A., Sivic, J., Russell, B.: Actionvlad: Learning spatio-temporal aggregation for action classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 971–980 (2017) 3
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016) 9
11. Heo, B., Kim, J., Yun, S., Park, H., Kwak, N., Choi, J.Y.: A comprehensive overhaul of feature distillation. In: Proceedings of the IEEE International Conference on Computer Vision (2019) 1, 8, 9
12. Heo, B., Lee, M., Yun, S., Choi, J.Y.: Knowledge transfer via distillation of activation boundaries formed by hidden neurons. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 3779–3787 (2019) 1, 2, 9
13. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015) 1, 3, 4, 5, 7, 9
14. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017) 8
15. Huang, Z., Wang, N.: Like what you like: Knowledge distill via neuron selectivity transfer. arXiv preprint arXiv:1707.01219 (2017) 3
16. Jégou, H., Douze, M., Schmid, C., Pérez, P.: Aggregating local descriptors into a compact image representation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2010) 3
17. Kim, J., Park, S., Kwak, N.: Paraphrasing complex network: Network compression via factor transfer. In: Advances in Neural Information Processing Systems. pp. 2760–2769 (2018) 1, 8
18. Krizhevsky, A.: Learning multiple layers of features from tiny images (2009) 7

19. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: International Conference on Learning Representations (2017) 1
20. Li, Z., Hoiem, D.: Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence* **40**(12), 2935–2947 (2017) 1, 12
21. Liu, Y., Cao, J., Li, B., Yuan, C., Hu, W., Li, Y., Duan, Y.: Knowledge distillation via instance relationship graph. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (June 2019) 1
22. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018) 7
23. Park, W., Kim, D., Lu, Y., Cho, M.: Relational knowledge distillation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (June 2019) 3
24. Passalis, N., Tefas, A.: Learning deep representations with probabilistic knowledge transfer. In: Proceedings of the European Conference on Computer Vision. pp. 268–284 (2018) 1
25. Peng, B., Jin, X., Liu, J., Li, D., Wu, Y., Liu, Y., Zhou, S., Zhang, Z.: Correlation congruence for knowledge distillation. In: Proceedings of the IEEE International Conference on Computer Vision (October 2019) 3
26. Perronnin, F., Dance, C.: Fisher kernels on visual vocabularies for image categorization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2007) 3
27. Quattoni, A., Torralba, A.: Recognizing indoor scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 413–420. IEEE (2009) 11
28. Romero, A., Ballas, N., Kahou, S.E., Chassang, A., Gatta, C., Bengio, Y.: Fitnets: Hints for thin deep nets. In: International Conference on Learning Representations (2015), <https://arxiv.org/abs/1412.6550> 1, 3, 9
29. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014) 9, 11
30. Sivic, J., Zisserman, A.: Video google: Efficient visual search of videos. In: Toward category-level object recognition. Springer (2006) 3
31. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: Advances in neural information processing systems. pp. 1195–1204 (2017) 1
32. Tian, Y., Krishnan, D., Isola, P.: Contrastive representation distillation. In: International Conference on Learning Representations (2020) 3, 7, 8, 9, 10
33. Tolias, G., Avrithis, Y., Jégou, H.: To aggregate or not to aggregate: Selective match kernels for image search. In: Proceedings of the IEEE International Conference on Computer Vision (2013) 3
34. Tung, F., Mori, G.: Similarity-preserving knowledge distillation. In: Proceedings of the IEEE International Conference on Computer Vision (2019) 3, 9, 10
35. Yang, C., Xie, L., Qiao, S., Yuille, A.: Knowledge distillation in generations: More tolerant teachers educate better students. *arXiv preprint arXiv:1805.05551* (2018) 1
36. Yim, J., Joo, D., Bae, J., Kim, J.: A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4133–4141 (2017) 1, 2, 3, 9
37. Yuan, L., Tay, F.E., Li, G., Wang, T., Feng, J.: Revisit knowledge distillation: a teacher-free framework. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2020) 5

38. Zagoruyko, S., Komodakis, N.: Wide residual networks. In: Proceedings of the British Machine Vision Conference (2016) 9
39. Zagoruyko, S., Komodakis, N.: Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In: International Conference on Learning Representations (2017), <https://arxiv.org/abs/1612.03928> 1, 2, 3, 9