

Learning to Structure an Image with Few Colors

Yunzhong Hou Liang Zheng Stephen Gould
Australian National University
{firstname.lastname}@anu.edu.au

Abstract

Color and structure are the two pillars that construct an image. Usually, the structure is well expressed through a rich spectrum of colors, allowing objects in an image to be recognized by neural networks. However, under extreme limitations of color space, the structure tends to vanish, and thus a neural network might fail to understand the image. Interested in exploring this interplay between color and structure, we study the scientific problem of identifying and preserving the most informative image structures while constraining the color space to just a few bits, such that the resulting image can be recognized with possibly high accuracy. To this end, we propose a color quantization network, ColorCNN, which learns to structure the images from the classification loss in an end-to-end manner. Given a color space size, ColorCNN quantizes colors in the original image by generating a color index map and an RGB color palette. Then, this color-quantized image is fed to a pre-trained task network to evaluate its performance. In our experiment, with only a 1-bit color space (i.e., two colors), the proposed network achieves 82.1% top-1 accuracy on the CIFAR10 dataset, outperforming traditional color quantization methods by a large margin. For applications, when encoded with PNG, the proposed color quantization shows superiority over other image compression methods in the extremely low bit-rate regime. The code is available at https://github.com/hou-yz/color_distillation.

1. Introduction

Color and structure are two important aspects of a natural image. The structure is viewed as a combination of shapes, textures, etc, and is closely related to colors. Particularly, the structure is well presented only when there exists a sufficient set of colors. In this paper, we are interested in how structure can be best presented under color constraints.

In literature, a closely related line of works is color quantization. Color quantization investigates how to preserve visual similarity in a restricted color space [17, 29]. This

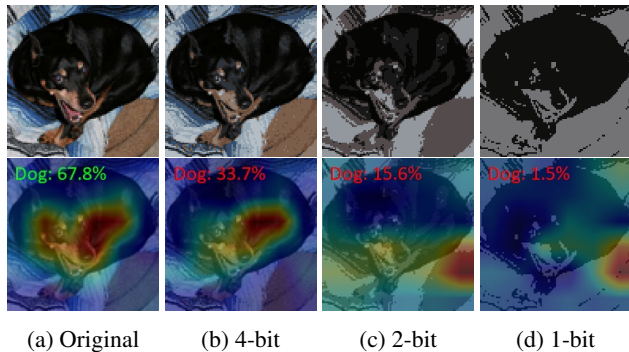


Figure 1: Color quantized images (top row) and class activation maps [50] with softmax probabilities for the ground truth classes (bottom row) in decreasing color space sizes. In “Original” (a), colors are described by 24 bits. In (b), (c) and (d), fewer bits are used. We use MedianCut [17] as the quantization method. When the color space is reduced, the focus of neural network deviates from the informative parts in (a) (correctly recognized (green)), resulting in *recognition failures* (red) in (b), (c), and (d).

problem is *human-centered*, as it usually focuses on the visual quality for human viewing. Particularly, most existing methods are designed under a relatively large color space, e.g., 8-bit, so that quantized images are still visually similar to the original images. In smaller color spaces, e.g., 2-bit or 1-bit, color quantization remains an open question.

Natural images usually contain rich colors and structures. When the color space is limited, their connection will compromise the structure. For example, in the first row of Fig. 1, structures vanish as the color space reduces. Moreover, we argue that the neural network trained on original natural images might be ineffective in recognizing quantized images with few colors. In fact, quantized colors and compromised structures drift the attention of the neural network, which might result in recognition failures. For example, in the second row of Fig. 1, a neural network trained on original images finds the head and body most critical for recognizing the dog. When the color space is reduced gradually, the neural network first fails to attend to the head, and then the body, thus leading to recognition failures.

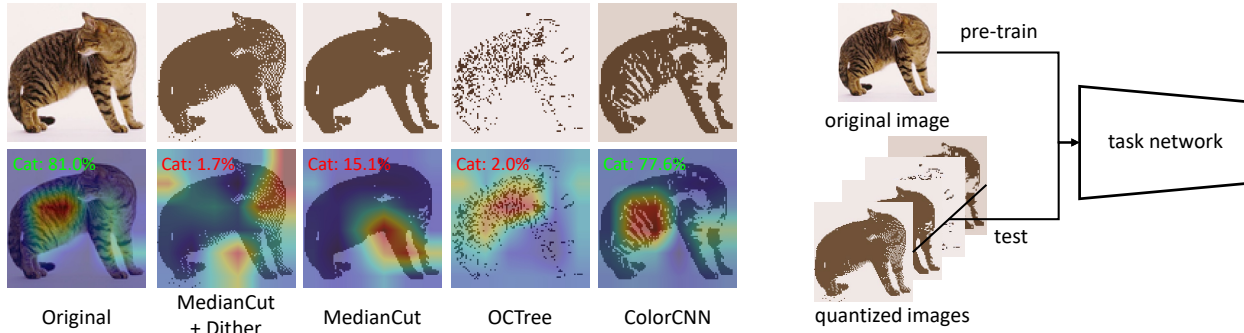


Figure 2: 1-bit color quantization results and evaluation. **Quantization results:** “MedianCut + Dither” [17, 12], “MedianCut” [17], and “OCTree” [14] are traditional color quantization methods. “ColorCNN” is the proposed method. **Evaluation:** We first pre-train a classification network on the original images. Then, we use this network to evaluate the quantized images. Traditional methods quantize the original image only based on the color, and hence may lose important shapes and textures. When feeding the traditional method results, the pre-trained classifier fails to attend to the informative parts, and make mistakes. In comparison, by learning to structure, the image quantized by our method keeps a most similar activation map and thus is successfully recognized by the pre-trained classifier.

In this work, we study a scientific problem: how to preserve the critical structures under an extremely small color space? This problem is orthogonal to traditional color quantization problems, as it is *task-centered*: neural network recognition accuracy is its major focus, instead of human viewing. As shown in Fig. 2, we evaluate the (quantized) images with a classifier pre-trained on original images. To optimize the recognition accuracy, we design a color quantization method, ColorCNN, which learns to structure the image in an end-to-end manner. Unlike traditional color quantization methods that solely rely on color values for the decision, ColorCNN exploits colors, structures, and semantics to spot and preserve the critical structures. It quantizes the image by computing a color index map and assigning color palette values. In Fig. 2, images quantized by ColorCNN enables the classifier to successfully focus on the cat tabby and forelimb. In this example, attending to the informative regions leads to correct recognition.

We demonstrate the effectiveness of ColorCNN quantization on classification tasks. With a few colors, we show that the proposed methods outperforms traditional ones by a large margin. Four datasets, including CIFAR10 [20], CIFAR100 [20], STL10 [8], and tiny-imagenet-200 [22], and three classification networks, including AlexNet [21], VGG [34], and ResNet [16], are used for testing. For applications, the ColorCNN image quantization can be used in extremely low bit-rate image compression.

2. Related Work

Color quantization. Color quantization [29, 10, 1, 11, 45] clusters the colors to reduce the color space while keeping visual similarity. Heckbert *et al.* [17] propose the popular MedianCut method. Later, Gervautz *et al.* [14] design another commonly-used color quantization method, OC-

Tree. Dithering [12], which removes visual artifacts by adding a noise pattern, is also studied as an optional step. The color quantized images can be represented as *indexed color* [30], and encoded with PNG [6].

Human-centered image compression. Based on heuristics, many image compression methods are designed for human viewers. These methods fall into two categories, lossless compression, *e.g.*, PNG [6], and lossy compression, *e.g.*, JPEG [41, 35] and *color quantization*.

Recently, deep learning methods are introduced to image compression problems. Both recurrent methods [28, 19, 39, 38] and convolutional methods [27, 23, 40, 3, 2, 37] are investigated. Ballé *et al.* [5] propose generalized divisive normalization (GDN) for image compression. Agustsson *et al.* [3] propose a multi-scale discriminator in a generative adversarial network (GAN) for low-bitrate compression. In [24, 18], researchers apply image compression methods to defend against adversarial attacks.

Task-centered image compression. Traditional or deep-learning based, the fore-mentioned image compression methods are human-centered. Liu *et al.* [25] points out that for segmentation, human-centered compression is not the best choice for 3D medical images. For 2D map data and 3D scene models, task-centered compression methods are designed for localization [43, 7].

Neural network recognition with compressed data. Some researchers work on certain tasks with compressed data. For example, solving action recognition problem with compressed video [44, 47, 48, 33]. Wang *et al.* [42] accelerate video object recognition by exploiting compressed data.

3. Motivation

In an extremely small color space, we find the traditional color quantization methods fail to preserve the critical struc-

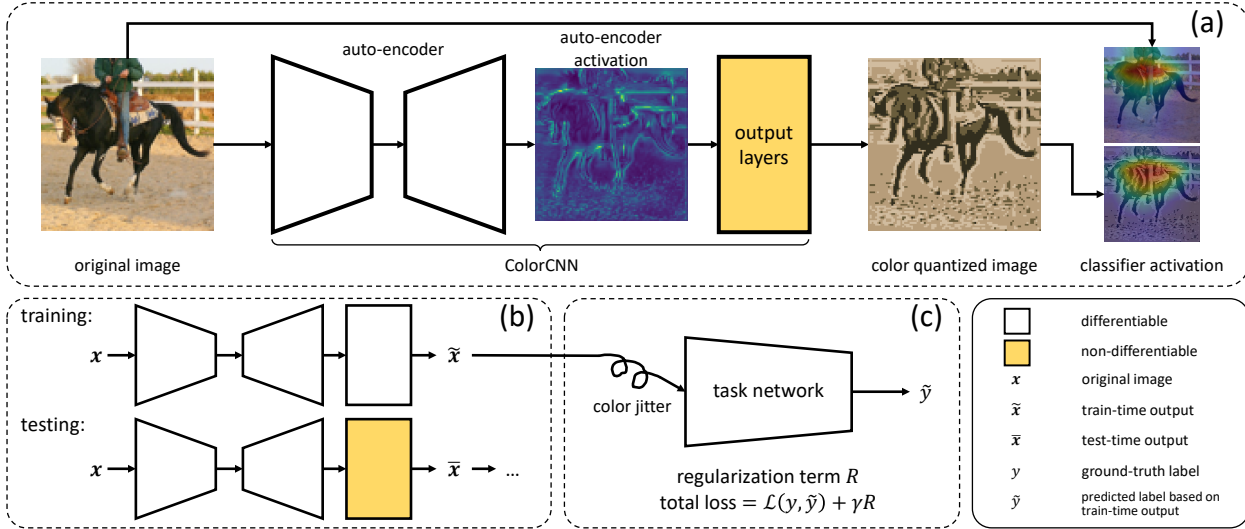


Figure 3: Overview of our color quantization method. (a): ColorCNN can identify (auto-encoder activation) and preserve (quantized image) the critical structures in the original image. Its output has a similar class activation map to the original image. (b): we replace the non-differentiable parts with approximations during training. (c): one regularization term is introduced to keep the approximation similar to the original network. Also, we add a color jitter to the quantized image to prevent premature convergence. The ColorCNN network is trained with classification loss in an end-to-end manner.

tures. This is because these methods usually take the color-value-only approach to cluster the colors, completely ignoring structures. However, as some [13, 15] suggest, critical shapes and textures with semantic meaning play an important role in neural network recognition. In fact, we witness attention drifts when the structure is not well preserved in Fig. 1 and Fig. 2. This further leads to recognition failure. Motivated by this failure, in the following sections, we further investigate how to effectively preserve the structures in an extremely small color space.

4. Proposed Approach

To identify and preserve the critical structures in the original image, we design ColorCNN (see Fig. 3). In this section, we first formulate the learning-to-structure problem mathematically. Next, we present ColorCNN architecture. At last, we provide an end-to-end training method.

4.1. Problem Formulation

Without loss of generality, we can define a classification network with parameter θ as $f_\theta(\cdot)$. For an image label pair (\mathbf{x}, y) , this network estimate its label $\hat{y} = f_\theta(\mathbf{x})$. To train the network, we minimize its loss $\mathcal{L}(y, \hat{y})$ on dataset D ,

$$\theta^* = \arg \min_{\theta} \sum_{(\mathbf{x}, y) \in D} \mathcal{L}(y, f_\theta(\mathbf{x})), \quad (1)$$

where θ^* denotes the optimal parameter.

For color quantization, we design ColorCNN architecture. Given input image \mathbf{x} , it can output the color quantized

image $\bar{\mathbf{x}}$, by computing the color index map $M(\mathbf{x})$ and the color palette $T(\mathbf{x})$. We represent the forward pass for ColorCNN as one function $\bar{\mathbf{x}} = g_\psi(\mathbf{x})$ with parameter ψ .

Our objective is to structure an image with few colors such that a pre-trained classifier has a possibly high accuracy on the color quantized images. It is written as

$$\psi^* = \arg \min_{\psi} \sum_{(\mathbf{x}, y) \in D} \mathcal{L}(y, f_{\theta^*}(g_\psi(\mathbf{x}))) + \gamma R, \quad (2)$$

where ψ^* denotes the optimized parameter for ColorCNN, given the pre-trained classifier parameterized by θ^* . R is a regularization term and γ denotes its weight.

4.2. ColorCNN Architecture

We show the ColorCNN architecture in Fig. 4. Its first component is an U-net [31] auto-encoder that identifies the critical and semantic-rich structures (auto-encoder activation map in Fig. 3).

For the second component, two depth-wise (1×1 kernel size) convolution layers create a softmax probability map of each pixel taking one specific color. This results in a C -channel probability map $m(\mathbf{x})$ (softmax over C -channel).

Then, for each input image \mathbf{x} , the 1-channel color index map $M(\mathbf{x})$ is computed as the arg max over the C -channel probability map $m(\mathbf{x})$,

$$M(\mathbf{x}) = \arg \max_c m(\mathbf{x}). \quad (3)$$

The RGB color palette, $T(\mathbf{x})$, which is of shape $C \times 3$, is computed as average of all pixels that falls into certain

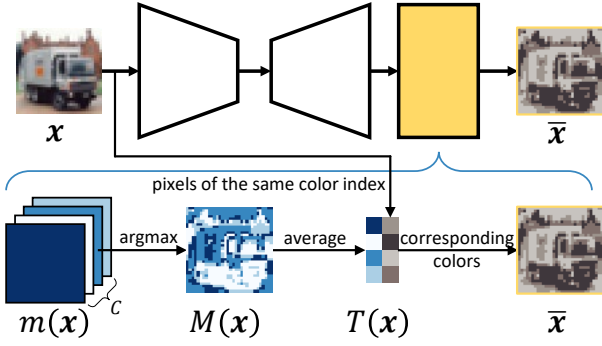


Figure 4: ColorCNN architecture (**test-time**). First, the convolutional layers output a C -channel probability map $m(\mathbf{x})$ for C colors. Next, a 1-channel color index map $M(\mathbf{x})$ is created via the arg max function. Then, the color palette $T(\mathbf{x})$ is computed as average of all pixels that are of the *same color index*. At last, the color quantized image $\bar{\mathbf{x}}$ is created via a *table look-up* session.

quantized color index,

$$[T(\mathbf{x})]_c = \frac{\sum_{(u,v)} [\mathbf{x}]_{u,v} \cdot \mathbb{I}([M(\mathbf{x})]_{u,v} = c)}{\sum_{(u,v)} \mathbb{I}([M(\mathbf{x})]_{u,v} = c)}, \quad (4)$$

where $[\cdot]_i$ denotes the element or tensor with index i . $\mathbb{I}(\cdot)$ is an indicator function. \bullet denotes the pointwise multiplication. For pixel (u, v) in a $W \times H$ image, $[\mathbf{x}]_{u,v}$ denotes the pixel and its RGB value in the input image, and $[M(\mathbf{x})]_{u,v}$ represents its computed color index. $[T(\mathbf{x})]_c$ denotes the RGB value for the quantized color c .

At last, the quantized image $\bar{\mathbf{x}}$ is created as

$$\bar{\mathbf{x}} = \sum_c [T(\mathbf{x})]_c \cdot \mathbb{I}(M(\mathbf{x}) = c). \quad (5)$$

By combining Eq. 3, 4, 5, we finish the ColorCNN forward pass $\bar{\mathbf{x}} = g_\psi(\mathbf{x})$.

4.3. End-to-End Learning

4.3.1 Differentiable Approximation

Fig. 5 shows the differentiable approximation in training. To start with, we remove the arg max 1-channel color index map $M(\mathbf{x})$ in Eq. 3. Instead, we use the C -channel softmax probability map $m(\mathbf{x})$.

Next, we change the color palette design that follows. For each quantized color, instead of averaging the pixels of the *same color index*, we set its RGB color value $[t(\mathbf{x})]_c$ as the weighted average over *all* pixels,

$$[t(\mathbf{x})]_c = \frac{\sum_{(u,v)} [\mathbf{x}]_{u,v} \cdot [m(\mathbf{x})]_{u,v,c}}{\sum_{(u,v)} [m(\mathbf{x})]_{u,v,c}}, \quad (6)$$

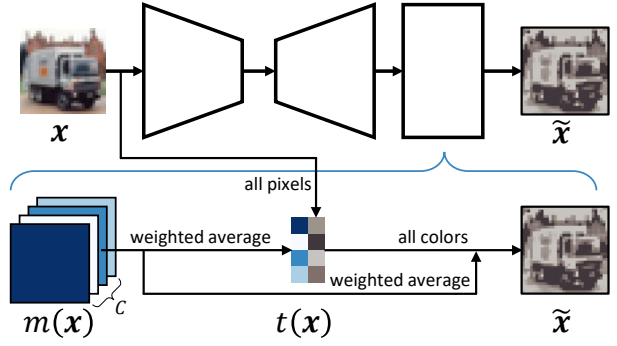


Figure 5: The differentiable approximation (**train-time**). The C -channel probability map $m(\mathbf{x})$ is used instead of the arg max color index map $M(\mathbf{x})$. Next, the color palette $t(\mathbf{x})$ is adjusted as weighted average over *all* pixels. At last, instead of table look-up, the quantized image $\tilde{\mathbf{x}}$ is computed as the weighted average of *all* colors in the color palette.

Here, the C -channel probability distribution $[m(\mathbf{x})]_{u,v}$ is used as the contribution ratio of pixel (u, v) to C colors. This will result in a slightly different color palette $t(\mathbf{x})$.

In the end, we change the table look-up process from the original forward pass into a weighted sum. For quantized color with index c , we use $[m(\mathbf{x})]_c$ as the intensity of expression over entire image. Mathematically, the train-time quantized image $\tilde{\mathbf{x}}$ is computed as

$$\tilde{\mathbf{x}} = \sum_c [t(\mathbf{x})]_c \cdot [m(\mathbf{x})]_c. \quad (7)$$

By combining Eq. 6, 7, the forward pass for ColorCNN during training can be formulated as $\tilde{\mathbf{x}} = \tilde{g}_\psi(\mathbf{x})$. At last, we substitute $g_\psi(\cdot)$ with $\tilde{g}_\psi(\cdot)$ in Eq. 2 for end-to-end training.

Even though the two forward pass $g_\psi(\cdot)$ and $\tilde{g}_\psi(\cdot)$ use the same parameter ψ , they behave very differently. See Fig. 6 for a side-by-side comparison between the outputs. The test-time output $\bar{\mathbf{x}}$ only has C colors, whereas the train-time output $\tilde{\mathbf{x}}$ has more than C colors. The main reason for this mismatch boils down to the difference between one-hot and softmax vectors. As shown in Fig. 4, the one-hot approach allows influence only from *some* pixels to one quantized color, and from *one* color to any quantized pixel. On the other hand, in Fig. 5, with softmax function, *all* pixels influence all colors in the palette, and *all* colors in the palette contribute to each pixel in the output image.

4.3.2 Regularization

During training, the softmax probability distribution in each pixel can deviate far from a one-hot vector, which can lead to overfitting.

In order to minimize this difference, we propose a regularization term. It encourages the probability distribution in

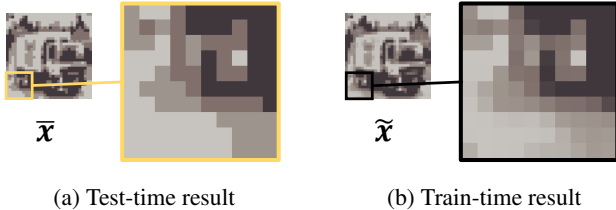


Figure 6: Comparison between test-time result \bar{x} and train-time result \tilde{x} . Each pixel in \tilde{x} is a weighted average of all colors in its palette. Thus, more colors are introduced.

each pixel to be more like a one-hot vector. For pixel (u, v) , we maximize the largest value of the probability distribution $[m(\mathbf{x})]_{u,v}$. The regularization term is designed as

$$R = \log_2 C \times \left(1 - \frac{1}{C} \times \sum_c \max_{(u,v)} [m(\mathbf{x})]_{u,v} \right). \quad (8)$$

We take the minus of the summation, since it is desired to minimize this regularization term R . We also offset the regularization term by 1 in order to make it positive.

4.3.3 Color Jitter

We train the proposed ColorCNN with a pre-trained classifier $f_{\theta^*}(\cdot)$ (see Eq. 2). During training, as long as ColorCNN can provide barely satisfactory results, the pre-trained classifier will have a good chance in making the correct decision. However, given more freedom in the train-time results, the test-time results might still be struggling when the network converges. We refer to this phenomenon as premature convergence.

In order to prevent this premature convergence, during training, we add a jitter $\xi \times n$ to the color quantized image \tilde{x} after normalization. The noise n is sampled from a Gaussian distribution $\mathcal{N}(0, 1)$. ξ denotes its weight. **First**, higher variance in training output delays convergence (more difficult to fit data with higher variance), allowing a better-trained network. **Second**, feature distribution of training output with color jitter (higher variance) may have more overlap with that of the testing output. Since the classifier can recognize the color jittered training output, it may perform better during testing.

5. Discussion

Larger color space is not the focus of this work. When the color space is larger, the richer colors naturally support more structures, making structure preserving less of a scientific problem. Moreover, many traditional methods study this problem, and their quantization results achieve very good accuracy on pre-trained classifiers. In fact, in a 6-bit color space, accuracy of quantized images only fall behind original images marginally (see Section 6.2).



(a) Original (b) MedianCut (c) OCTree (d) ColorCNN

Figure 7: 6-bit color quantization results. Traditional methods formulate color quantization as a clustering problem. On the contrary, we formulate this problem as per-pixel classification. When a bigger number-of-cluster (color space size C) is assumed, classification-based method (ColorCNN) cannot compete with clustering-based methods. For more details, see Section 5.

ColorCNN versus clustering-based color quantization. Using a per-pixel classification formulation, ColorCNN cannot provide a competitive result to clustering-based methods, when the color space is large (Fig. 7). This is very normal, given that per-pixel classification treats each pixel individually, whereas clustering consider all pixels to enforce intra-cluster similarity to make a collective decision. The feature of each pixel has global info, but this is insufficient: making decisions globally and collectively is needed. With a competent end-to-end neural network clustering method, ColorCNN can possibly out-perform traditional methods, even in a large color space. Why not incorporating clustering in ColorCNN, one may ask. In fact, using neural networks to solve the clustering problem is non-trivial, and stands as a challenging problem itself. Some pioneer works in neural network clustering investigate end-to-end feature update, or feature dimension reduction [46, 4]. Still, they must rely on k-means or other clustering methods during testing. In fact, neural network clustering itself is a different line of work, and is beyond the scope of this paper.

6. Experiment

6.1. Experiment Setup

Datasets. We test ColorCNN performance on multiple datasets. CIFAR10 and CIFAR100 datasets [20] include 10 and 100 classes of general objects, respectively. Similar to CIFAR10, STL10 dataset [8] also contains 10 image classes. However, the images in STL10 have a higher 96×96 resolution. We also evaluate on the tiny-imagenet-200 dataset [22], which is a subset of ImageNet dataset [9]. It has 200 classes of 64×64 general object images. We compare the four datasets in Table 1.

Evaluation. For evaluation, we report top-1 classification accuracy on the mentioned datasets.

Classification networks. We choose AlexNet [21], VGG16 [34] and ResNet18 [16] for classification network. All classifier networks are trained for 60 epochs with a batch size of 128, except for STL10, where we set the batch

	CIFAR10	CIFAR100	STL10	Tiny200
#class	10	100	10	200
Train images	50,000	50,000	5,000	100,000
Test images	10,000	10,000	8,000	10,000 [†]
Resolution	32 × 32	32 × 32	96 × 96	64 × 64

Table 1: Datasets comparison. “Tiny200” denotes tiny-imagenet-200 dataset. [†] indicates we use validation set for testing, given that the online test set is not available.

size to 32. We use an SGD optimizer with a momentum of 0.5, L2-normalization of 5×10^{-4} . We choose the 1cycle learning rate scheduler [36] with a peak learning rate at 0.1.

ColorCNN. We train ColorCNN on top of the original image pre-trained classifier. We set the hyper-parameters as follows. For the regularization and color jitter weight, we set $\gamma = 1$ and $\xi = 1$. We also normalize the quantized image by $4 \times$ the default variance of the original images, so as to prevent premature convergence. For training, we run the gradient descent for 60 epochs with a batch size of 128. Similar to classification networks, we also reduce the batch size to 32 on the STL10 dataset. The SGD optimizer for ColorCNN training is the same as for classifier networks. For the learning rate scheduler, we choose Cosine-Warm-Restart [26] with a peak learning rate at 0.01, minimal learning rate at 0, and uniform restart period of 20.

We finish all the experiment on one RTX-2080TI GPU.

6.2. Evaluation of ColorCNN

Classification network performance. We report the top-1 test accuracy of the classification networks in Table 2. There is a consistent accuracy increase going from AlexNet to VGG16 to ResNet18 on all four datasets.

Visualization for ColorCNN low-bit quantization. As shown in Fig. 8, ColorCNN effectively preserves the shapes, textures and other structures. For instance, airplane wings, vehicle tires and windshield, and bird cheek and belly. In column (d), we find the feature maps extracted via auto-encoder show high activation on the informative edges, details, and textures. We further show the accuracy increase from these critical structures in Fig. 9.

Low-bit color quantization performance. We report top-1 classification accuracy with color quantized images on all four datasets using three networks in Fig. 9. We choose MedianCut [17], OCTree [14], and MedianCut with dithering [12] for comparisons.

First, the proposed ColorCNN method brings a consistent and significant improvement over the traditional quantization method in a **small color space**. Using AlexNet as the classification network, 1-bit quantization results of ColorCNN reach 82.1%, 24.8%, 52.3%, and 26.0% on four datasets, respectively. This translates into 37.6%, 8.6%, 9.6%, and 14.5% absolute accuracy increases and 92.9%,

	CIFAR10	CIFAR100	STL10	Tiny200
AlexNet [21]	86.8	62.5	73.8	50.2
VGG16 [34]	93.5	73.1	79.8	62.6
ResNet18 [16]	94.6	76.3	84.3	69.1

Table 2: Top-1 test accuracy (%) of classification networks.



(a) Original (b) MedianCut (c) ColorCNN (d) Activation

Figure 8: Example of 1-bit color quantization results. Visualized based on its auto-encoder output, column (d) shows the critical structures identified by ColorCNN.

56.5%, 66.5%, and 75.3% relative accuracy increases over the traditional color quantization methods. Due to more demanding task in the 100-way classification compared to the 10-way one (on same data), the improvements (absolute and relative) are smaller on CIFAR100 compared to CIFAR10. Nonetheless, these non-trivial accuracy increases still prove the effectiveness of the proposed method. We point out that classifiers trained on original images 1) have significantly lower performance on low-bit images, 2) but still manage to classify some low-bit images. The color quantized images and original images can be regarded as two different domains. The two domains differ significantly in appearance, but also share the same set of classes and some appearance similarities. Their feature distributions thus are very different but still have some level of overlap, which accounts for both sides of the phenomenon.

Second, ColorCNN is usually inferior to the traditional methods under a **large color space**. As discussed in Section 5, ColorCNN does not formulate color quantization as a clustering problem. This will naturally lead to inferior results under a larger number-of-cluster (color space size), since the per-pixel approach in ColorCNN cannot enforce intra-cluster similarity.

Third, preserving structures via dithering is found not useful. Dithering generates a noise pattern based on the quantization error of the color values. It can further remove

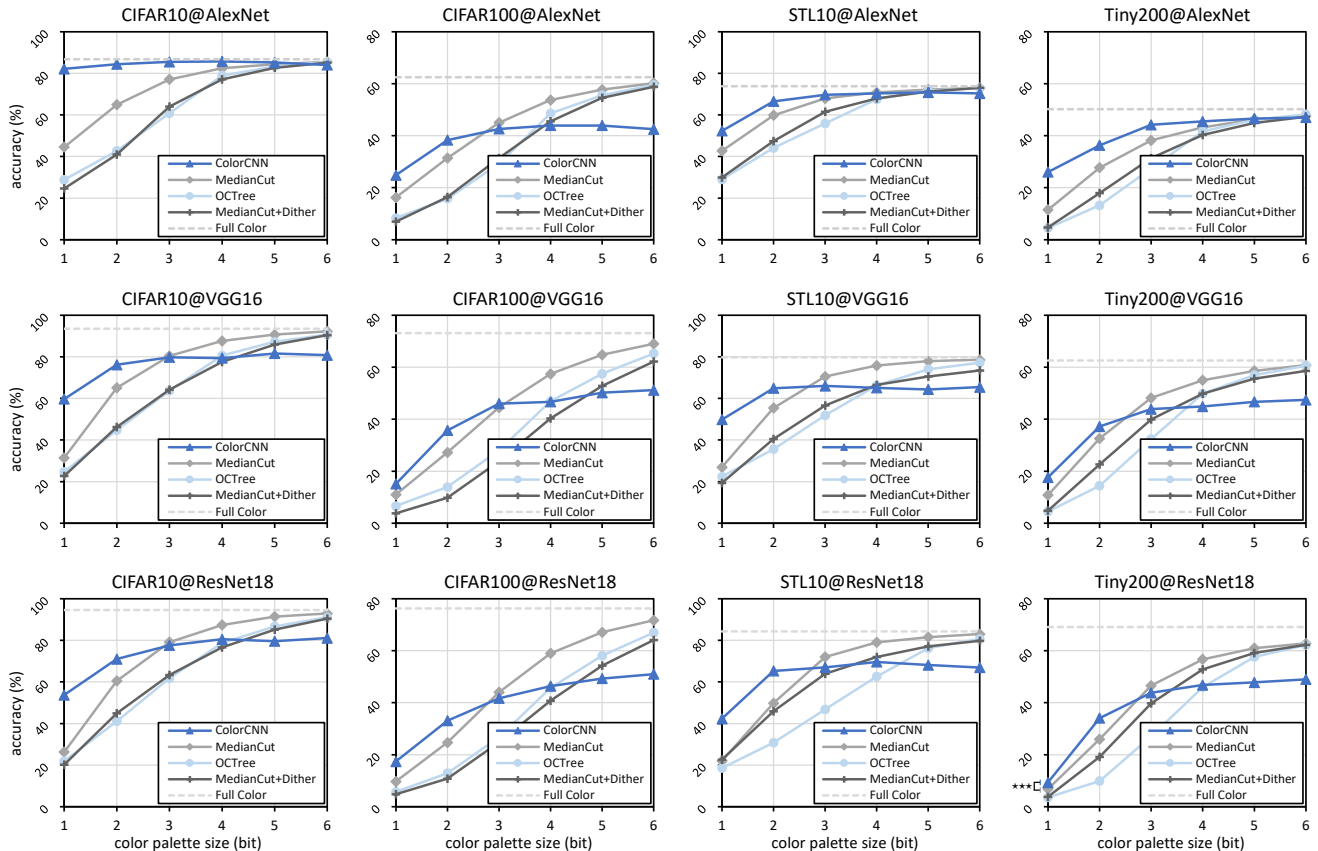


Figure 9: Top-1 classification accuracy of color quantized images on four datasets with three networks. We observe that ColorCNN is significantly superior to MedianCut, OCTree, and MedianCut+Dither under low-bit quantization. *** means that the accuracy difference between ColorCNN and MedianCut, is **statistically very significant** (*i.e.*, p -value < 0.001), in 1-bit color space, tiny-imagenet-200 dataset, and ResNet18 classifier. Note that quantization under the large color space is not the focus of this work. More discussion around this consideration is provided in Section 5.

the flat color region and false contours. Without consideration of the structure and semantic, dithering still fails to preserve the semantic-rich structures in a restricted color space, which further leads to inferior accuracy.

Impact of different classification networks. We compare the quantization performance with different classification networks in different rows of Fig. 9. It is found that stronger classifier has lower accuracy in an extremely small color space. A stronger classifier can add more transformation to the image data, extracting more expressive features, thus having higher accuracy. However, when the color space is limited, these colors and structures all diminish. This can lead to larger drifts in feature space from the stronger classifiers, since they add more transformation to the input, which ultimately leads to lower accuracy.

We also find that ColorCNN performance is not always higher when trained with stronger classifiers. In fact, stronger classifiers can easily classify the quantized image \tilde{x} during training, which can lead to earlier, less mature con-

vergence. Given the difference between train-time and test-time forward pass $\tilde{g}_\psi(\cdot)$ and $g_\psi(\cdot)$, this less mature convergence can result in more overfitting and lower accuracy.

Low-bit color quantization as image compression. In Fig. 10, as the color space size grows from 1-bit to 6-bit, the quantized images take a higher bit-rate when encoded with PNG, and have higher test accuracy. When compared to traditional color quantization methods, ColorCNN can reach higher test accuracy under a lower bit-rate. As shown in Fig. 7, even if 6-bit color space is allowed, ColorCNN only uses a few colors for the majority of the image. Not using all colors evenly will introduce a lower information entropy [32], which leads to a smaller bit-rate when compressed losslessly via PNG. This suggests ColorCNN can extract the key structure of an image more effectively. Moreover, under 0.2 bits per pixel, 1-bit ColorCNN quantization can even outperform JPEG compression by 13.2%, which has more than 2 colors. This clearly demonstrates the effectiveness of ColorCNN.

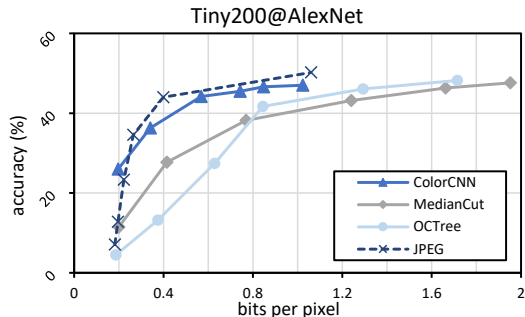


Figure 10: Classification accuracy under different bit-rate. Solid lines refer to color quantized image encoded via PNG. Dotted line refers to JPEG encoding as reference. For color quantization methods, bit-rate from low to high are quantized images under color space size from 1-bit to 6-bit.

6.3. Ablation Study

We set the regularization weight $\gamma = 0$ or color jitter weight $\xi = 0$, to see the ColorCNN performance without either of those. Results are shown in Table 3.

First, we find that removing the regularization term causes an accuracy drop. In fact, without the regularization, fewer colors are chosen during test-time. This is because the softmax color filling during training can introduce more colors in the image, as shown in Fig. 6. This difference in color filling leads to overfitting and a 2.2% accuracy drop.

Second, no color jitter also leads to an accuracy drop. Without color jitter, the train-time quantization can be too easy for the pre-trained classifier. This can lead to premature convergence, which further hurts accuracy by 1.9%.

Third, higher bit-rate does not necessarily lead to higher accuracy, what matters is preserving the critical structure. Under similar accuracy, ColorCNN without regularization actually has a smaller bit-rate than without color jitter. With both regularization and color jitter, the bit-rate becomes even higher. However, this time, since the introduced structures can help the recognition, ColorCNN achieves the highest accuracy. We find this coherent with the compression rate curves in Fig. 10, where ColorCNN achieves higher accuracy under lower bit-rate, since it preserve the more critical structures.

6.4. Variant Study

We compare the recognition accuracy of ColorCNN and its variants, including ones with different hyper-parameters, and one with different auto-encoder backbone (Fig. 11). ColorCNN performance is lower when the regularization weight is too small or too high. Similarly, too small or too large a color jitter can also result in a huge accuracy drop. This is because setting the weight too small or too large leads to either too small a influence, or completely overshadowing anything else. For both regularization and color

	Accuracy (%)	#color/image	#bit/pixel
ColorCNN	69.7	8.0	0.425
w/o regularization	67.5	5.1	0.323
w/o color jitter	67.8	8.0	0.390

Table 3: Test results under 3-bit color quantization, STL10 dataset, AlexNet classifier.

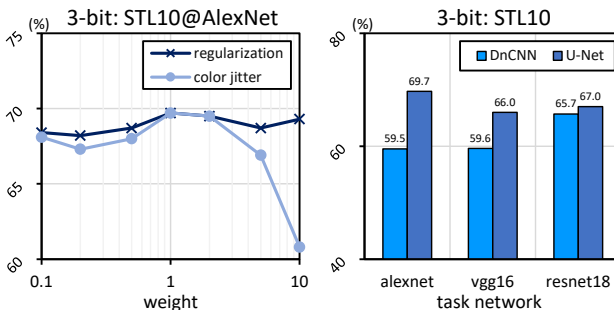


Figure 11: Performance comparison with different weights and different auto-encoder backbone.

jitter, we witness the highest accuracy when the weight is set to 1, which corresponds to our hyper-parameter setting.

When the auto-encoder backbone is replaced with DnCNN [49], the ColorCNN performance degrades under all classification networks. Unlike U-Net, DnCNN does not have bypasses to maintain the local structure. As a result, its quantization results might have structure misalignment, which hurts classification accuracy.

7. Conclusion

In this paper, we investigate the scientific problem of keeping informative structures under an extremely small color space. In such cases, traditional color quantization methods tend to lose both color and structure, making its output incomprehensible to neural networks. In order to maintain the critical structures in the quantized images so that they can be correctly recognized, we design the ColorCNN color quantization network. By incorporating multiple cues for a comprehensive quantization decision making, ColorCNN effectively identifies and preserves the informative structures, even in the extreme conditions. An end-to-end training method is also designed for ColorCNN, so as to maximize the performance of the quantized image on the neural network tasks. The effectiveness of the proposed method is demonstrated on four datasets with three classification networks.

8. Acknowledgements

Dr. Liang Zheng is the recipient of Australian Research Council Discovery Early Career Award (DE200101283) funded by the Australian Government.

References

- [1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE transactions on pattern analysis and machine intelligence*, 34(11):2274–2282, 2012.
- [2] Eirikur Agustsson, Fabian Mentzer, Michael Tschannen, Lukas Cavigelli, Radu Timofte, Luca Benini, and Luc V Gool. Soft-to-hard vector quantization for end-to-end learning compressible representations. In *Advances in Neural Information Processing Systems*, pages 1141–1151, 2017.
- [3] Eirikur Agustsson, Michael Tschannen, Fabian Mentzer, Radu Timofte, and Luc Van Gool. Generative adversarial networks for extreme learned image compression. *arXiv preprint arXiv:1804.02958*, 2018.
- [4] Elie Aljalbout, Vladimir Golkov, Yawar Siddiqui, Maximilian Strobel, and Daniel Cremers. Clustering with deep learning: Taxonomy and new methods. *arXiv preprint arXiv:1801.07648*, 2018.
- [5] Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimized image compression. *arXiv preprint arXiv:1611.01704*, 2016.
- [6] Thomas Boutell. Png (portable network graphics) specification version 1.0. 1997.
- [7] Federico Camposco, Andrea Cohen, Marc Pollefeys, and Torsten Sattler. Hybrid scene compression for visual localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7653–7662, 2019.
- [8] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 215–223, 2011.
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [10] Yining Deng, Charles Kenney, Michael S Moore, and BS Manjunath. Peer group filtering and perceptual color image quantization. In *ISCAS’99. Proceedings of the 1999 IEEE International Symposium on Circuits and Systems VLSI (Cat. No. 99CH36349)*, volume 4, pages 21–24. IEEE, 1999.
- [11] Yining Deng and BS Manjunath. Unsupervised segmentation of color-texture regions in images and video. *IEEE transactions on pattern analysis and machine intelligence*, 23(8):800–810, 2001.
- [12] RW Floyd and L Steinberg. An adaptive technique for spatial grayscale. In *Proceedings of the Society of Information Display*, volume 17, pages 78–84, 1976.
- [13] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231*, 2018.
- [14] Michael Gervautz and Werner Purgathofer. A simple method for color quantization: Octree quantization. In *New trends in computer graphics*, pages 219–231. Springer, 1988.
- [15] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [17] Paul Heckbert. *Color image quantization for frame buffer display*, volume 16. ACM, 1982.
- [18] Xiaojun Jia, Xingxing Wei, Xiaochun Cao, and Hassan Foroosh. Comdefend: An efficient image compression model to defend adversarial examples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6084–6092, 2019.
- [19] Nick Johnston, Damien Vincent, David Minnen, Michele Covell, Saurabh Singh, Troy Chinen, Sung Jin Hwang, Joel Shor, and George Toderici. Improved lossy image compression with priming and spatially adaptive bit rates for recurrent networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4385–4393, 2018.
- [20] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Cite-seer, 2009.
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [22] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 2015.
- [23] Mu Li, Wangmeng Zuo, Shuhang Gu, Debin Zhao, and David Zhang. Learning convolutional networks for content-weighted image compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3214–3223, 2018.
- [24] Zihao Liu, Qi Liu, Tao Liu, Yanzhi Wang, and Wujie Wen. Feature distillation: Dnn-oriented jpeg compression against adversarial examples. *arXiv preprint arXiv:1803.05787*, 2018.
- [25] Zihao Liu, Xiaowei Xu, Tao Liu, Qi Liu, Yanzhi Wang, Yiyu Shi, Wujie Wen, Meiping Huang, Haiyun Yuan, and Jian Zhuang. Machine vision guided 3d medical image compression for efficient transmission and accurate segmentation in the clouds. *arXiv preprint arXiv:1904.08487*, 2019.
- [26] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [27] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Practical full resolution learned lossless image compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10629–10638, 2019.
- [28] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016.
- [29] Michael T Orchard and Charles A Bouman. Color quantization of images. *IEEE transactions on signal processing*, 39(12):2677–2690, 1991.

- [30] Charles Poynton. *Digital video and HD: Algorithms and Interfaces*. Elsevier, 2012.
- [31] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [32] Claude Elwood Shannon. A mathematical theory of communication. *Bell system technical journal*, 27(3):379–423, 1948.
- [33] Zheng Shou, Xudong Lin, Yannis Kalantidis, Laura Sevilla-Lara, Marcus Rohrbach, Shih-Fu Chang, and Zhicheng Yan. Dmc-net: Generating discriminative motion cues for fast compressed video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1268–1277, 2019.
- [34] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [35] Athanassios Skodras, Charilaos Christopoulos, and Touradj Ebrahimi. The jpeg 2000 still image compression standard. *IEEE Signal processing magazine*, 18(5):36–58, 2001.
- [36] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, volume 11006, page 1100612. International Society for Optics and Photonics, 2019.
- [37] Lucas Theis, Wenzhe Shi, Andrew Cunningham, and Ferenc Huszár. Lossy image compression with compressive autoencoders. *arXiv preprint arXiv:1703.00395*, 2017.
- [38] George Toderici, Sean M O’Malley, Sung Jin Hwang, Damien Vincent, David Minnen, Shumeet Baluja, Michele Covell, and Rahul Sukthankar. Variable rate image compression with recurrent neural networks. *arXiv preprint arXiv:1511.06085*, 2015.
- [39] George Toderici, Damien Vincent, Nick Johnston, Sung Jin Hwang, David Minnen, Joel Shor, and Michele Covell. Full resolution image compression with recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5306–5314, 2017.
- [40] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. In *Advances in neural information processing systems*, pages 4790–4798, 2016.
- [41] Gregory K Wallace. The jpeg still picture compression standard. *IEEE transactions on consumer electronics*, 38(1):xviii–xxxiv, 1992.
- [42] Shiyao Wang, Hongchao Lu, and Zhidong Deng. Fast object detection in compressed video. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7104–7113, 2019.
- [43] Xinkai Wei, Ioan Andrei Barsan, Shenlong Wang, Julieta Martinez, and Raquel Urtasun. Learning to localize through compressed binary maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10316–10324, 2019.
- [44] Chao-Yuan Wu, Manzil Zaheer, Hexiang Hu, R Manmatha, Alexander J Smola, and Philipp Krähenbühl. Compressed video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6026–6035, 2018.
- [45] Xiaolin Wu. Color quantization by dynamic programming and principal analysis. *ACM Transactions on Graphics (TOG)*, 11(4):348–372, 1992.
- [46] Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *International conference on machine learning*, pages 478–487, 2016.
- [47] Bowen Zhang, Limin Wang, Zhe Wang, Yu Qiao, and Hanli Wang. Real-time action recognition with enhanced motion vector cnns. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2718–2726, 2016.
- [48] Bowen Zhang, Limin Wang, Zhe Wang, Yu Qiao, and Hanli Wang. Real-time action recognition with deeply transferred motion vector cnns. *IEEE Transactions on Image Processing*, 27(5):2326–2339, 2018.
- [49] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017.
- [50] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.