

# Do we train on test data? Purging CIFAR of near-duplicates

Björn Barz

Joachim Denzler

Computer Vision Group  
Friedrich Schiller University Jena  
Ernst-Abbe-Platz 2, 07743 Jena, Germany

`bjoern.barz@uni-jena.de`

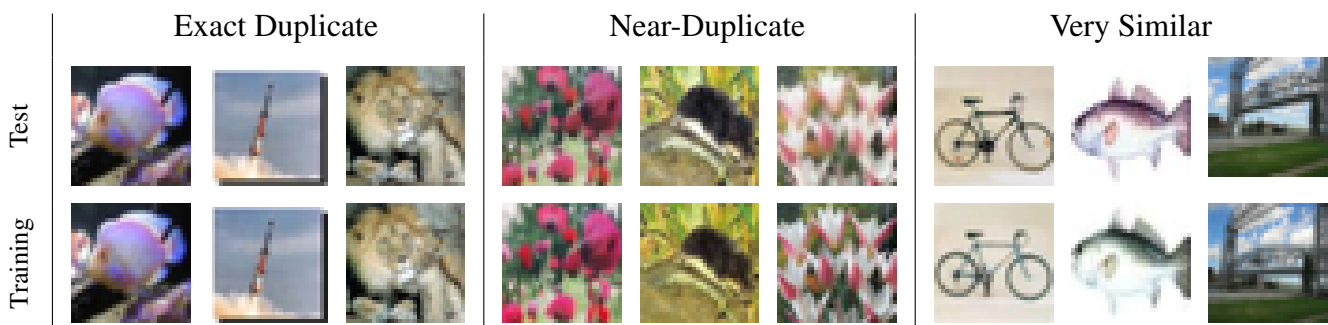


Figure 1: Examples for different types of duplicates between the CIFAR-100 test and training set. The top row shows images from the test set and the bottom row shows their nearest neighbors from the training set in a CNN feature space. Please see the main text for a description of the three categories of duplicates.

## Abstract

We find that 3.3% and 10% of the images from the CIFAR-10 and CIFAR-100 test sets, respectively, have duplicates in the training set. This may incur a bias on the comparison of image recognition techniques with respect to their generalization capability on these heavily benchmarked datasets.

To eliminate this bias, we provide the “fair CIFAR” (ciFAIR) dataset, where we replaced all duplicates in the test sets with new images sampled from the same domain. The training set remains unchanged, in order not to invalidate pre-trained models.

We then re-evaluate the classification performance of various popular state-of-the-art CNN architectures on these new test sets to investigate whether recent research has overfitted to memorizing data instead of learning abstract concepts. Fortunately, this does not seem to be the case yet.

The ciFAIR dataset and pre-trained models are available at <https://cvjena.github.io/cifair/>, where we also maintain a leaderboard.

## 1. Introduction

Almost ten years after the first instantiation of the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [15], image classification is still a very active field of research. The majority of recent approaches belongs to the domain of *deep learning* with several new architectures of *convolutional neural networks* (CNNs) being proposed for this task every year and trying to improve the accuracy on held-out test data by a few percent points [7, 22, 21, 8, 6, 13, 3]. A key to the success of these methods is the availability of large amounts of training data [12, 17]. The world wide web has become a very affordable resource for harvesting such large datasets in an automated or semi-automated manner [4, 11, 9, 20].

A problem of this approach is that there is no effective automatic method for filtering out near-duplicates among the collected images. When the dataset is split up later into a training, a test, and maybe even a validation set, this might result in the presence of near-duplicates of test images in the training set. Usually, the post-processing with regard to duplicates is limited to removing images that have exact

pixel-level duplicates [11, 4]. However, many duplicates are less obvious and might vary with respect to contrast, translation, stretching, color shift etc. These are variations that can easily be accounted for by data augmentation, so that these variants will actually become part of the augmented training set.

For a proper scientific evaluation, the presence of such duplicates is a critical issue: We actually aim at comparing models with respect to their ability of generalizing to unseen data. With a growing number of duplicates, however, we run the risk to compare them in terms of their capability of memorizing the training data, which increases with model capacity. This is especially problematic when the difference between the error rates of different models is as small as it is nowadays, *i.e.*, sometimes just one or two percent points. The significance of these performance differences hence depends on the overlap between test and training data. In some fields, such as fine-grained recognition, this overlap has already been quantified for some popular datasets, *e.g.*, for the Caltech-UCSD Birds dataset [19, 10].

In this work, we assess the number of test images that have near-duplicates in the training set of two of the most heavily benchmarked datasets in computer vision: CIFAR-10 and CIFAR-100 [11]. We will first briefly introduce these datasets in Section 2 and describe our duplicate search approach in Section 3. In a nutshell, we search for nearest neighbor pairs between test and training set in a CNN feature space and inspect the results manually, assigning each detected pair into one of four duplicate categories. We find that 3.3% of CIFAR-10 test images and a surprising number of 10% of CIFAR-100 test images have near-duplicates in their respective training sets.

Subsequently, we replace all these duplicates with new images from the Tiny Images dataset [18], which was the original source for the CIFAR images (see Section 4). To determine whether recent research results are already affected by these duplicates, we finally re-evaluate the performance of several state-of-the-art CNN architectures on these new test sets in Section 5.

Similar to our work, Recht et al. [14] have recently sampled a completely new test set for CIFAR-10 from Tiny Images to assess how well existing models generalize to truly unseen data. Furthermore, they note parenthetically that the CIFAR-10 test set comprises 8% duplicates with the training set, which is more than twice as much as we have found. As opposed to their work, however, we also analyze CIFAR-100 and only replace the duplicates in the test set, while leaving the remaining images untouched. Moreover, we distinguish between three different types of duplicates and publish a list of duplicates, the new test sets, and pre-trained models at <https://cvjena.github.io/cifair/>.

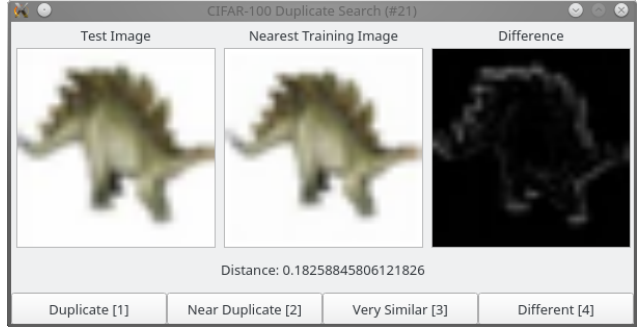


Figure 2: GUI for duplicate annotation.

## 2. The CIFAR Datasets

There exist two different CIFAR datasets [11]: CIFAR-10, which comprises 10 classes, and CIFAR-100, which comprises 100 classes. Both contain 50,000 training and 10,000 test images. Neither the classes nor the data of these two datasets overlap, but both have been sampled from the same source: the Tiny Images dataset [18].

In this context, the word “tiny” refers to the resolution of the images, not to their number. On the contrary, Tiny Images comprises approximately 80 million images collected automatically from the web by querying image search engines for approximately 75,000 synsets of the WordNet ontology [5]. However, all images have been resized to the “tiny” resolution of  $32 \times 32$  pixels.

Due to their much more manageable size and the low image resolution, which allows for fast training of CNNs, the CIFAR datasets have established themselves as one of the most popular benchmarks in the field of computer vision.

## 3. Hunting Duplicates

As we have argued above, simply searching for exact pixel-level duplicates is not sufficient, since there may also be slightly modified variants of the same scene that vary by contrast, hue, translation, stretching etc. Thus, we follow a content-based image retrieval approach [16, 2, 1] for finding duplicate and near-duplicate images: We train a lightweight CNN architecture proposed by Barz et al. [3] on the training set and then extract  $L^2$ -normalized features from the global average pooling layer of the trained network for both training and testing images. For each test image, we find the nearest neighbor from the training set in terms of the Euclidean distance in that feature space.

Given this, it would be easy to capture the majority of duplicates by simply thresholding the distance between these pairs. However, such an approach would result in a high number of false positives as well. Therefore, we inspect the detected pairs manually, sorted by increasing distance. In a graphical user interface depicted in Fig. 2, the annotator can




Figure 3: Number of duplicates per duplicate type between test and training set (blue) and within the test set (orange).

inspect the test image and its duplicate, their distance in the feature space, and a pixel-wise difference image. The pair is then manually assigned to one of four classes:

**Exact Duplicate** Almost all pixels in the two images are approximately identical.

**Near-Duplicate** The content of the images is exactly the same, *i.e.*, both originated from the same camera shot. However, different post-processing might have been applied to this original scene, *e.g.*, color shifts, translations, scaling etc.

**Very Similar** The contents of the two images are different, but highly similar, so that the difference can only be spotted at the second glance.

**Different** The pair does not belong to any other category.

Figure 1 shows some examples for the three categories of duplicates from the CIFAR-100 test set, where we picked the 10<sup>th</sup>, 50<sup>th</sup>, and 90<sup>th</sup> percentile image pair for each category, according to their distance. In the remainder of this paper, the word “duplicate” will usually refer to any type of duplicate, not necessarily to exact duplicates only.

We used a single annotator and stopped the annotation once the class “Different” has been assigned to 20 pairs in a row. In addition to spotting duplicates of test images in the training set, we also search for duplicates within the test set, since these also distort the performance evaluation. Note that we do not search for duplicates within the training set. In the worst case, the presence of such duplicates biases the weights assigned to each sample during training, but they are not critical for evaluating and comparing models.

We found 891 duplicates from the CIFAR-100 test set in the training set and another set of 104 duplicates within the test set itself. In total, 10% of test images have duplicates. The situation is slightly better for CIFAR-10, where

| (a) CIFAR-10 |    | (b) CIFAR-100 |    |
|--------------|----|---------------|----|
| frog         | 59 | cockroach     | 33 |
| automobile   | 55 | orange        | 29 |
| airplane     | 47 | lawn mower    | 28 |
| deer         | 40 | apple         | 26 |
| bird         | 26 | skunk         | 26 |
| horse        | 20 | keyboard      | 23 |
| dog          | 19 | leopard       | 23 |
| ship         | 17 | wolf          | 21 |
| truck        | 14 | worm          | 20 |
| cat          | 11 | whale         | 17 |
|              |    | rocket        | 17 |
|              |    | skyscraper    | 17 |
|              |    | dinosaur      | 17 |
|              |    | maple tree    | 17 |

Table 1: The top 14 classes with the most duplicates.

we found 286 duplicates in the training and 39 in the test set, amounting to 3.25% of the test set. This is probably due to the much broader type of object classes in CIFAR-10: We suppose it is easier to find 5,000 different images of birds than 500 different images of maple trees, for example.

The vast majority of duplicates belongs to the category of near-duplicates, as can be seen in Fig. 3. It is worth noting that there are no exact duplicates in CIFAR-10 at all, as opposed to CIFAR-100. This might indicate that the basic duplicate removal step mentioned by Krizhevsky et al. [12] has been omitted during the creation of CIFAR-100.

Table 1 lists the top 14 classes with the most duplicates for both datasets. The only classes without any duplicates in CIFAR-100 are “bowl”, “bus”, and “forest”.

On the subset of test images with duplicates in the training set, the ResNet-110 [7] models from our experiments in Section 5 achieve error rates of 0% and 2.9% on CIFAR-10 and CIFAR-100, respectively. This verifies our assumption that even the near-duplicate and highly similar images can be classified correctly much to easily by memorizing the training data.

## 4. The Duplicate-Free ciFAIR Test Dataset

To create a fair test set for CIFAR-10 and CIFAR-100, we replace all duplicates identified in the previous section with new images sampled from the Tiny Images dataset [18], which was also the source for the original CIFAR datasets.

We took care not to introduce any bias or domain shift during the selection process. To this end, each replacement candidate was inspected manually in a graphical user interface (see Fig. 4), which displayed the candidate image and the three nearest neighbors in the feature space from the existing training and test sets. We approved only those




Figure 4: GUI for replacement candidate selection.

samples for inclusion in the new test set that could not be considered duplicates (according to the category definitions in Section 3) of any of the three nearest neighbors.

Furthermore, we followed the labeler instructions provided by Krizhevsky et al. [11] for CIFAR-10. However, separate instructions for CIFAR-100, which was created later, have not been published. We found by looking at the data that some of the original instructions seem to have been relaxed for this dataset. For example, CIFAR-100 does include some line drawings and cartoons as well as images containing multiple instances of the same object category. Both types of images were excluded from CIFAR-10. Therefore, we also accepted some replacement candidates of these kinds for the new CIFAR-100 test set.

We term the datasets obtained by this modification as ciFAIR-10 and ciFAIR-100 (“fair CIFAR”). They consist of the original CIFAR training sets and the modified test sets which are free of duplicates. ciFAIR can be obtained online at <https://cvjena.github.io/cifair/>.

## 5. Re-evaluation of the State of the Art

Two questions remain: Were recent improvements to the state-of-the-art in image classification on CIFAR actually due to the effect of duplicates, which can be memorized better by models with higher capacity? Does the ranking of methods change given a duplicate-free test set?

To answer these questions, we re-evaluate the performance of several popular CNN architectures on both the CIFAR and ciFAIR test sets. Unfortunately, we were not able to find any pre-trained CIFAR models for any of the

architectures. Thus, we had to train them ourselves, so that the results do not exactly match those reported in the original papers. However, we used the original source code, where it has been provided by the authors, and followed their instructions for training (*i.e.*, learning rate schedules, optimizer, regularization etc.).

The results are given in Table 2. Besides the absolute error rate on both test sets, we also report their difference (“gap”) in terms of absolute percent points, on the one hand, and relative to the original performance, on the other hand.

On average, the error rate increases by 0.41 percent points on CIFAR-10 and by 2.73 percent points on CIFAR-100. The relative difference, however, can be as high as 12%. The ranking of the architectures did not change on CIFAR-100, and only Wide ResNet and DenseNet swapped positions on CIFAR-10.

## 6. Conclusions

In a laborious manual annotation process supported by image retrieval, we have identified a surprising number of duplicate images in the CIFAR test sets that also exist in the training set. We have argued that it is not sufficient to focus on exact pixel-level duplicates only. In contrast, slightly modified variants of the same scene or very similar images bias the evaluation as well, since these can easily be matched by CNNs using data augmentation, but will rarely appear in real-world applications. We hence proposed and released a new test set called ciFAIR, where we replaced all those duplicates with new images from the same domain.

A re-evaluation of several state-of-the-art CNN models for image classification on this new test set lead to a significant drop in performance, as expected. The relative ranking of the models, however, did not change considerably. This is a positive result, indicating that the research efforts of the community have not overfitted to the presence of duplicates in the test set. However, all models we tested have sufficient capacity to memorize the complete training data. Thus, a more restricted approach might show smaller differences.

We encourage all researchers training models on the CIFAR datasets to evaluate their models on ciFAIR, which will provide a better estimate of how well the model generalizes to new data. To facilitate comparison with the state-of-the-art further, we maintain a community-driven leaderboard at <https://cvjena.github.io/cifair/>, where everyone is welcome to submit new models. We will only accept leaderboard entries for which pre-trained models have been provided, so that we can verify their performance.

## References

- [1] A. Babenko and V. Lempitsky. Aggregating local deep features for image retrieval. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1269–1277, 2015.

| Model                         | CIFAR-10     |              |              |              | CIFAR-100     |               |              |               |
|-------------------------------|--------------|--------------|--------------|--------------|---------------|---------------|--------------|---------------|
|                               | CIFAR        | ciFAIR       | Gap          | Rel. Gap     | CIFAR         | ciFAIR        | Gap          | Rel. Gap      |
| Plain-11 [3]                  | 5.91%        | 6.43%        | 0.52%        | 8.80%        | 27.82%        | 31.34%        | 3.52%        | 12.65%        |
| ResNet-110 [7]                | 5.26%        | 5.77%        | 0.51%        | 9.70%        | 26.05%        | 29.24%        | 3.19%        | <b>12.25%</b> |
| WRN-28-10 [22]                | 3.89%        | 4.25%        | 0.36%        | 9.25%        | 18.95%        | 21.48%        | 2.53%        | 13.35%        |
| DenseNet-BC (L=190, k=40) [8] | 3.90%        | 4.20%        | <b>0.30%</b> | <b>7.69%</b> | 18.62%        | 21.02%        | 2.40%        | 12.89%        |
| ResNeXt-29 (8x64d) [21]       | <b>3.56%</b> | <b>3.95%</b> | 0.39%        | 10.96%       | 18.38%        | 20.84%        | 2.46%        | 13.38%        |
| PyramidNet-272-200 [6]        | 3.58%        | 4.00%        | 0.42%        | 11.73%       | <b>17.05%</b> | <b>19.38%</b> | <b>2.33%</b> | 13.67%        |

Table 2: Classification error rate of various CNN architectures on the original CIFAR test sets and the modified ciFAIR test sets. The best value in each column is highlighted in bold font.

- [2] A. Babenko, A. Slesarev, A. Chigorin, and V. Lempitsky. Neural codes for image retrieval. In *European Conference on Computer Vision (ECCV)*, pages 584–599. Springer, 2014.
- [3] B. Barz and J. Denzler. Deep learning is not a matter of depth but of good training. In *International Conference on Pattern Recognition and Artificial Intelligence (ICPRAI)*, pages 683–687. CENPARMI, Concordia University, Montreal, 2018.
- [4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255. IEEE, 2009.
- [5] C. Fellbaum. *WordNet*. Wiley Online Library, 1998.
- [6] D. Han, J. Kim, and J. Kim. Deep pyramidal residual networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6307–6315. IEEE, 2017.
- [7] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [8] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4700–4708, 2017.
- [9] M. J. Huiskes and M. S. Lew. The MIR Flickr retrieval evaluation. In *MIR '08: Proceedings of the 2008 ACM International Conference on Multimedia Information Retrieval*, New York, NY, USA, 2008. ACM.
- [10] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu. Spatial transformer networks. *arXiv preprint arXiv:1506.02025*, 2015.
- [11] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1097–1105, 2012.
- [13] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le. Regularized evolution for image classifier architecture search. *arXiv preprint arXiv:1802.01548*, 2018.
- [14] B. Recht, R. Roelofs, L. Schmidt, and V. Shankar. Do cifar-10 classifiers generalize to cifar-10? *arXiv preprint arXiv:1806.00451*, 2018.
- [15] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [16] A. W. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain. Content-based image retrieval at the end of the early years. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, (12):1349–1380, 2000.
- [17] C. Sun, A. Shrivastava, S. Singh, and A. Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *IEEE International Conference on Computer Vision (ICCV)*, pages 843–852. IEEE, 2017.
- [18] A. Torralba, R. Fergus, and W. T. Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 30(11):1958–1970, 2008.
- [19] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- [20] B. Wu, W. Chen, Y. Fan, Y. Zhang, J. Hou, J. Huang, W. Liu, and T. Zhang. Tencent ML-Images: A large-scale multi-label image database for visual representation learning. *arXiv preprint arXiv:1901.01703*, 2019.
- [21] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5987–5995. IEEE, 2017.
- [22] S. Zagoruyko and N. Komodakis. Wide residual networks. In E. R. H. Richard C. Wilson and W. A. P. Smith, editors, *British Machine Vision Conference (BMVC)*, pages 87.1–87.12. BMVA Press, September 2016.