

The banner features a grid of 18 small images from the CIFAR-10 dataset, each with a label above it: 'plane:a-plane', 'dog:dog', 'cat:deer', 'ship:frog', 'truck:car', 'bird:deer', 'cat:cat', 'truck:truck', 'frog:frog', 'truck:truck', 'horse:deer', 'frog:frog', 'truck:truck'. A navigation bar at the bottom contains the following tabs: Overview, Data (highlighted with a blue underline), Kernels, Discussion, Leaderboard, and Rules.

Overview

Description

Description

Evaluation

The learning task in this exercise is to create an image classifier using Convolutional Neural Network (CNN) architectures for the CIFAR-10 Dataset of images [1,2,6]. The data consists of 60,000 images of objects from 10 distinct categories (airplane, car, bird, cat, deer, dog, frog, horse, ship, truck). The images are 32x32 pixels with three RGB color channels [6]. You should create an image classifier by developing a Convolutional Neural Network (CNN) and training it using the packages and approaches we have been discussing in lecture. See the course website for example codes. For the test dataset, your output should be a prediction of the class of each image. This will be evaluated by the percentage of predictions that are in the correct category.

You should experiment with different choices of architectures, hyper-parameters, and training protocols to get familiar with how convolutional neural networks work and how they can be tuned to improve image classification. You are welcome to experiment with other CNN architectures, with stacking additional convolutional layers, by adjusting the width, kernel size, and depth of filters, or with performing data augmentation or by making other modifications. For inspiration, see the papers below for discussions of both historic and modern techniques [1-5].

References

- [1] Hand written digit recognition with a back-propagation network, Y. LeCun, In Proc. Advances in Neural Information Processing Systems, (1990).
- [2] ImageNet Classification with Deep Convolutional Neural Networks, Krizhevsky, A., Sutskever, I. and Hinton, G. E., NIPS: Neural Information Processing Systems, (2012).
- [3] Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position, K. Fukushima, Biological Cybernetics, 36, 4, April (1980).

[4] Receptive fields of single neurons in the cat's striate cortex, Hubel, D.H.; Wiesel, T.N., J Physiol. 148 (3), (1959).

[5] Deep learning, LeCun, Y., Bengio, Y., and Hinton, G., Nature, 521, May, (2015).


[6] CIFAR10 Dataset: Learning Multiple Layers of Features from Tiny Images, Alex Krizhevsky, 2009. Images collected by Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton of the Canadian Institute For Advanced Research, available at: <https://www.cs.toronto.edu/~kriz/cifar.html>.

Author of Competition

This competition was setup for a data science course by Paul J. Atzberger, <http://atzberger.org/>

The image shows three navigation panels in a light gray theme. The first panel is titled 'Leaderboard' and contains a list of numbers from 1 to 8, each followed by a hyphen. The second panel is titled 'Kernels' and has a white rectangular area in the center. The third panel is titled '0 discussion topics' and also has a white rectangular area in the center. Each panel has a right-pointing chevron icon.

A horizontal progress bar is shown. On the left, it says 'Launch' followed by '15 minutes ago' and a solid blue circle. On the right, it says 'Close' followed by '3 months' and an open circle. A thin gray line connects the two circles.





InClass Prediction Competition

CIFAR-10: Image Classification Exercise

Convolutional Neural Networks (CNNs) and Image Classification.

Overview

Data

Kernels

Discussion

Leaderboard

Rules

Data Description

Data Description

This is the CIFAR-10 Dataset consists of a collection of 60,000 images of 32x32 pixels with three 8-bit RGB channels. Note that in some machine learning frameworks this data set can be downloaded directly from into data structures useful for processing in the CNN. You are welcome to use those APIs to access the data set as an alternative to manually downloading the data posted here. Be sure to keep the last 10,000 images as your test dataset for submitting predictions.

To read in Python a batch of images from the data files posted here, please use the following approach. For python 2, you can use the cPickle module to the load the batch of data as

```
import cPickle
```

```
def loadImageBatch(filename):  
    fid = open(filename, 'rb');  
    data_batch = cPickle.load(fid);  
    return data_batch;
```

For python 3, you can use pickle module to load the batch of data as

Data (163 MB)

API

kaggle competitions download -c atz-neural-netwo...



Download All



Data Sources

About this file

cifar-10-python.tar.gz

No description yet

cifar-10-batche... 8 files



InClass Prediction Competition

CIFAR-10: Image Classification Exercise

Convolutional Neural Networks (CNNs) and Image Classification.

Overview

Data

Kernels

Discussion

Leaderboard

Rules

Data Description

For python 3, you can use pickle module to load the batch of data as

```
import pickle
```

```
def loadImageBatch(filename):
```

```
    fid = open(filename, 'rb');
```

```
    data_batch = pickle.load(fid,encoding='bytes');
```

```
    return data_batch;
```

The data batches are numpy arrays of dimension 10000x3072 of data type uint8. Each image can be obtained by reshaping the 3072 vector into a 32x32x3 array encoding the RGB image. More information can also be found on the CIFAR-10 dataset website, <http://www.cs.toronto.edu/~kriz/cifar.html>.

File descriptions

- cifar-10-python.tar.gz - the data_batchX with X in the range 1-5 are the training data images and test_batch is the test data images on which to make predictions.

Data (163 MB)

API

kaggle competitions download -c atz-neural-netwo...



Download All



Data Sources

About this file

▼ cifar-10-python.tar.gz

No description yet

➤ cifar-10-batche... 8 files