

Kitti 2d Object Detection Factsheet

Overview

Task: Object Detection (7 classes)

Total dataset: 7481 images

Training set: 5984 images

Curated training set (90%): 5386 images

Validation set: 1497 images

Evaluation Method

- Detectron2 (<https://github.com/facebookresearch/detectron2>)
- Faster RCNN using a ResNet-50, trained for 46k steps
- Varied parameter training subset size: 75%, 90%, 100%
- Varied parameter sampler: random, coreset

Results Using Lightly

As depicted in the plot, Lightly Coreset achieves overall better mAP than random sampling. Furthermore, the mAP of random sampling increases monotonously with a larger training set size. However, Lightly Coreset reaches its optimum by using only 90% of the training data. The reason for this result is redundancies in the training data, which negatively affect the model when used for training it. While random sampling introduces a proportional amount of redundancies to the sample size, Lightly Coreset is able to identify them and avoids including them in the training data.

About Kitti

The [KITTI](#) (Karlsruhe Institute of Technology and Toyota Technological Institute) dataset consists of hours of traffic scenarios recorded with a variety of sensors, and it is one of the most

popular within the field of autonomous driving. The KITTI dataset offers itself for various tasks, but in our case, we focused on object detection. The [KITTI object detection benchmark](#) consists of 7'481 training images and 7'518 test images, comprising 80'256 labeled objects. All images are in color and saved as png.

Because there are no labels available for the official test set we randomly split the training set into a 80% training and 20% validation set.

Experiments

For our experiments, we trained the same model and varied the number of training examples used for training. We evaluate two sampling methods: (1) random and (2) coreset. Random subsampling is a common baseline as it does not change the statistical distribution of the underlying data. The coreset method focuses on diversifying the dataset. Coreset requires embeddings of the images to perform the selection process. In order to get the embeddings we use our [self-supervised learning framework](#) and train a SimCLR model for 200 epochs on the training set. After sampling the training data (either using coreset or random), we train the object detection model (Faster RCNN with a ResNet-50 backbone) and measure the mean average precision (mAP) on the validation data.

