

BDD100K: Images 100K Dataset 📁 100000 📁 12 👁 33172[Self-Driving, Benchmark](#) [Instance Segmentation](#) [Made in 2020](#) [Custom](#)[❤](#) [📁 5 GB](#) [📁](#) [Explore](#)

Introduction

Fisher Yu, Haofeng Chen, Xin Wang [et al.](#)

This is the **Images 100K** part of the **Berkeley Deep Drive Dataset (BDD100K): A Diverse Driving Dataset for Heterogeneous Multitask Learning**, which is the largest driving video dataset with 100K videos and 10 tasks, providing a comprehensive evaluation platform for image recognition algorithms in autonomous driving. The dataset boasts geographic, environmental, and weather diversity, enhancing the robustness of trained models. Through BDD100K, the authors establish a benchmark for heterogeneous multitask learning, demonstrating the need for specialized training strategies for existing models to handle such diverse tasks, thereby opening avenues for future research in this domain.


Currently, the following datasets are presented on a DatasetNinja platform:

- **BDD100K: Images 100K** (current)
- **BDD100K: Images 10K** ([available on DatasetNinja](#))




The dataset includes a diverse set of driving videos under various weather conditions, time, and scene types. The dataset also comes with a rich set of annotations: scene tagging, object bounding box, lane marking, drivable area, full-frame semantic and instance segmentation, multiple object tracking, and multiple object tracking with segmentation.

The authors construct BDD100K as a diverse and large-scale dataset of visual driving scenes. This dataset overcomes limitations by collecting over 100K diverse video clips, covering various driving scenarios, and capturing a broad range of appearance variations and pose configurations. The benchmarks encompass ten tasks, including image tagging, lane detection, drivable area segmentation, road object detection, semantic segmentation, instance segmentation, multi-object detection tracking, multi-object segmentation tracking, domain adaptation, and imitation learning. These diverse tasks enable the study of heterogeneous multitask learning, and the authors conduct extensive evaluations of existing algorithms on the new benchmarks, shedding light on the challenges of designing a single model for multiple tasks.



Geographical distribution of the data sources. Each dot represents the starting location of every video clip. The videos are from many cities and regions in the populous areas in the US.



Instance statistics of our object categories. (a) Number of instances of each category, which follows a long-tail distribution. (b) Roughly half of the instances are occluded. (c) About 7% of the instances are truncated.

To provide a large-scale diverse driving video dataset, the authors utilize a crowdsourcing approach with contributions from tens of thousands of drivers, supported by Nexar. BDD100K includes over 100K driving videos, each 40 seconds long, collected from more than 50K rides across locations like New York and the San Francisco Bay Area, offering diverse scene types and weather conditions. The dataset is split into training, validation, and testing sets, with annotations for image tasks at the 10th second in each video and the entire sequences used for tracking tasks.

Tasks

The Image tagging involves annotations for weather conditions, scene types, and times of day. Object detection includes bounding box annotations for 10 categories, while lane marking involves labeling with eight main categories, continuity, and direction attributes. Drivable area detection distinguishes between directly and alternatively drivable areas. Semantic instance segmentation provides pixel-level annotations for 40 object classes in images randomly sampled from the dataset.

The authors present their dataset as an essential resource for advancing research in street-scene understanding, offering unparalleled diversity and complexity for evaluating algorithms in the domain of autonomous driving.

[Homepage](#) [Research Paper](#) [GitHub](#)


Summary

Berkeley Deep Drive Dataset (BDD100K): A Diverse Driving Dataset for Heterogeneous Multitask Learning (Images 100K) is a dataset for instance segmentation, semantic segmentation, object detection, and identification tasks. It is used in the automotive industry.

The dataset consists of 100000 images with 2221128 labeled objects belonging to 12 different classes including car, drivable area, lane, and other: traffic sign, traffic light, person, truck, bus, bike, rider, motor, and train.









Images in the BDD100K: Images 100K dataset have pixel-level instance segmentation annotations. Due to the nature of the instance segmentation task, it can be automatically transformed into a semantic segmentation (only one mask for every class) or object detection (bounding boxes for every object) tasks. There are 20137 (20% of the total) unlabeled images (i.e. without annotations). There are 3 splits in the dataset: train (70000 images), test (20000 images), and val (10000 images). Additionally, every image contains tags with weather, scene and timeofday, while objects contain dictionary with useful meta-information about their attributes (area type, occlusion, truncation, etc.). The dataset was released in 2020 by the [UC Berkeley, USA](#), [Cornell University, USA](#), [UC San Diego, USA](#), and [Element, Inc.](#)

Here is the visualized example grid with animated annotations:



Explore

BDD100K: Images 100K dataset has 100000 images. Click on one of the examples below or open "Explore" tool anytime you need to view dataset images with annotations. This tool has extended visualization capabilities like zoom, translation, objects table, custom filters and more. Hover the mouse over the images to hide or show annotations.



Have a look at 100000 images

View images along with annotations and tags, search and filter by various parameters

Explore

Class Balance

There are **12 annotation classes** in the dataset. Find the general statistics and balances for every class in the table below. Click any row to preview images that have labels of the selected class. Sort by column to find the most rare or prevalent classes.

Q Search for classes...

< Rows 1-10 of 12 >

Class	Images <i>i</i>	Objects <i>i</i>	Count on image <i>i</i> average	Area on image <i>i</i> average
car any	78951	815717	10.33	9.86 %
drivable area any	76352	144352	1.89	17.05 %
lane any	76059	604379	7.95	0.93 %
traffic sign any	65375	274594	4.2	0.57 %
traffic light any	44890	213002	4.74	0.28 %
person any	25296	104611	4.14	1.26 %
truck any	21579	34216	1.59	4.77 %
bus any	10235	13269	1.3	5.02 %
bike any	4921	8217	1.67	0.93 %
rider any	4101	5166	1.26	0.87 %

Co-Occurrence Matrix

Co-occurrence matrix is an extremely valuable tool that shows you the images for every pair of classes: how many images have objects of both classes at the same time. If you click any cell, you will see those images. We added the tooltip with an explanation for every cell for your convenience, just hover the mouse over a cell to preview the description.

	bike	bus	car	drivable at...	lane	motor	person	rider	traffic light	traffic sign	train	truck
bike	4921	988	4865	4685	4809	449	3637	2764	3809	4360	7	1715
bus	988	10235	10151	9610	9941	464	4779	825	6735	8643	15	3627
car	4865	10151	78951	75503	75199	2593	25033	4054	44431	64562	117	21435
drivable area	4685	9610	75503	76352	73298	2500	23785	3901	42520	62591	112	20501
lane	4809	9941	75199	73298	76059	2532	24540	4023	43348	63015	117	20784
motor	449	464	2593	2500	2532	2618	1412	1457	1672	2202	5	655

person	3637	4779	25033	23785	24540	1412	25296	2672	18869	22020	46	8624
rider	2764	825	4054	3901	4023	1457	2672	4101	3080	3589	5	1345
traffic light	3809	6735	44431	42520	43348	1672	18869	3080	44890	39595	67	12498
traffic sign	4360	8643	64562	62591	63015	2202	22020	3589	39595	65375	100	17806
train	7	15	117	112	117	5	46	5	67	100	119	36
truck	1715	3627	21435	20501	20784	655	8624	1345	12498	17806	36	21579

Images

Explore every single image in the dataset with respect to the number of annotations of each class it has. Click a row to preview selected image. Sort by any column to find anomalies and edge cases. Use horizontal scroll if the table has many columns for a large number of classes in the dataset.


Search for images...

< Rows 1-10 of 47198 >

Image 	Split folder name	Height	Width	Unlabeled area	bike objects count	bike covered area	bus objects count	bus covered area
c8d0d179-7492f586.jpg	val	720 px	1280 px	75.35 %	0	0 %	0	0 %
b2490dcb-7469cbe8.jpg	val	720 px	1280 px	72.49 %	1	0.21 %	0	0 %
c1496e0d-bb0887f7.jpg	val	720 px	1280 px	76.57 %	0	0 %	0	0 %
c324a11d-eaaeb859.jpg	val	720 px	1280 px	68.96 %	0	0 %	1	0.02 %
c20ae70f-1c40dc28.jpg	val	720 px	1280 px	82.31 %	0	0 %	0	0 %
bd9881c7-682619d1.jpg	val	720 px	1280 px	68.29 %	0	0 %	0	0 %
c630c103-34a53b18.jpg	val	720 px	1280 px	83.15 %	0	0 %	0	0 %
be42e5e6-33445d5b.jpg	val	720 px	1280 px	84.89 %	0	0 %	0	0 %
c13b9590-5105f7ad.jpg	val	720 px	1280 px	73.45 %	0	0 %	0	0 %
bb17665d-27525f20.jpg	val	720 px	1280 px	79.36 %	0	0 %	1	0.63 %

Object Distribution

Interactive heatmap chart for every class with object distribution shows how many images are in the dataset with a certain number of objects of a specific class. Users can click cell and see the list of all corresponding images.



Class Sizes

The table below gives various size properties of objects for every class. Click a row to see the image with annotations of the selected class. Sort columns to find classes with the smallest or largest objects or understand the size differences between classes.


Q Search for classes...

< Rows 1-10 of 12 >

Class	Object count	Avg area	Max area	Min area	Min height	Min height	Max height
car any	815717	1.04 %	66.64 %	0 %	1 px	0.14 %	720 px
traffic sign any	274594	0.14 %	99.86 %	0 %	1 px	0.14 %	719 px
traffic light any	213002	0.06 %	33 %	0 %	1 px	0.14 %	653 px
drivable area any	144352	8.94 %	58.58 %	0 %	3 px	0.42 %	691 px
lane any	126013	0.48 %	13.28 %	0 %	2 px	0.28 %	577 px
person any	104611	0.33 %	37.55 %	0 %	1 px	0.14 %	644 px
truck any	34216	3.03 %	90.23 %	0 %	2 px	0.28 %	720 px
bus any	13269	3.89 %	84.34 %	0 %	4 px	0.56 %	720 px


Class	Object count	Avg area	Max area	Min area	Min height	Min height	Max height
bike any	8217	0.65 %	29.77 %	0 %	2 px	0.28 %	564 px
rider any	5166	0.69 %	28.01 %	0 %	1 px	0.14 %	574 px

Another way to visualize class area sizes is using the following tree map chart:



Spatial Heatmap

The heatmaps below give the spatial distributions of all objects for every class. These visualizations provide insights into the most probable and rare object locations on the image. It helps analyze objects' placements in a dataset.



Objects

Table contains all **78238** objects. Click a row to preview an image with annotations, and use search or pagination to navigate. Sort columns to find outliers in the dataset.

< Rows 1-10 of 78238 >

Object ID	Class	Image name click row to open	Image size height x width	Height	Height	Width	Width	Area
1	traffic light any	bb17665d-27525f20.jpg	720 x 1280	36 px	5 %	16 px	1.25 %	0.06 %

Object ID 	Class	Image name click row to open	Image size height x width	Height	Height	Width	Width	Area
2	traffic light any	bb17665d-27525f20.jpg	720 x 1280	28 px	3.89 %	10 px	0.78 %	0.03 %
3	traffic light any	bb17665d-27525f20.jpg	720 x 1280	17 px	2.36 %	9 px	0.7 %	0.02 %
4	traffic light any	bb17665d-27525f20.jpg	720 x 1280	15 px	2.08 %	14 px	1.09 %	0.02 %
5	traffic sign any	bb17665d-27525f20.jpg	720 x 1280	31 px	4.31 %	27 px	2.11 %	0.09 %
6	traffic sign any	bb17665d-27525f20.jpg	720 x 1280	13 px	1.81 %	18 px	1.41 %	0.03 %
7	traffic sign any	bb17665d-27525f20.jpg	720 x 1280	53 px	7.36 %	38 px	2.97 %	0.22 %
8	traffic sign any	bb17665d-27525f20.jpg	720 x 1280	15 px	2.08 %	18 px	1.41 %	0.03 %
9	car any	bb17665d-27525f20.jpg	720 x 1280	48 px	6.67 %	84 px	6.56 %	0.44 %
10	bus any	bb17665d-27525f20.jpg	720 x 1280	64 px	8.89 %	91 px	7.11 %	0.63 %

License

Copyright ©2018. The Regents of the University of California (Regents). All Rights Reserved.

THIS SOFTWARE AND/OR DATA WAS DEPOSITED IN THE BAIR OPEN RESEARCH COMMONS REPOSITORY ON 1/1/2021

Permission to use, copy, modify, and distribute this software and its documentation for educational, research, and not-for-profit purposes, without fee and without a signed licensing agreement; and permission to use, copy, modify and distribute this software for commercial purposes (such rights not subject to transfer) to BDD and BAIR Commons members and their affiliates, is hereby granted, provided that the above copyright notice, this paragraph and the following two paragraphs appear in all copies, modifications, and distributions. Contact The Office of Technology Licensing, UC Berkeley, 2150 Shattuck Avenue, Suite 510, Berkeley, CA 94720-1620, (510) 643-7201, otl@berkeley.edu, <http://ipira.berkeley.edu/industry-info> for commercial licensing opportunities.

IN NO EVENT SHALL REGENTS BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING LOST PROFITS, ARISING OUT OF THE USE OF THIS SOFTWARE AND ITS DOCUMENTATION, EVEN IF REGENTS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

REGENTS SPECIFICALLY DISCLAIMS ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE SOFTWARE AND ACCOMPANYING DOCUMENTATION, IF ANY, PROVIDED HEREUNDER IS PROVIDED “AS IS”. REGENTS HAS NO OBLIGATION TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

[Source](#)

Citation

If you make use of the BDD100K Images 100K data, please cite the following reference:

```
@dataset{BDD100K Images 100K,
  author={Fisher Yu and Haofeng Chen and Xin Wang and Wenqi Xian and Yingying Chen and Fangchen Liu and Vashisht Madhavan and Trevor Darre},
  title={BDD100K: A Diverse Driving Dataset for Heterogeneous Multitask Learning (Images 100K)},
  year={2020},
  url={https://www.bdd100k.com/}
}
```

[Source](#)

If you are happy with Dataset Ninja and use provided visualizations and tools in your work, please cite us:

```
@misc{ visualization-tools-for-bdd100k-dataset,
  title = { Visualization Tools for BDD100K: Images 100K Dataset },
  type = { Computer Vision Tools },
  author = { Dataset Ninja },
  howpublished = { \url{ https://datasetninja.com/bdd100k } },
  url = { https://datasetninja.com/bdd100k },
}
```

```
journal = { Dataset Ninja },
publisher = { Dataset Ninja },
year = { 2026 },
month = { jan },
note = { visited on 2026-01-27 },
}
```

Download

Dataset **BDD100K: Images 100K** can be downloaded in [Supervisely format](#):

 Download • 5.39 GB  Download Sample • 88.03 MB

As an alternative, it can be downloaded with `dataset-tools` package:

```
pip install --upgrade dataset-tools
```

... using following python code:

```
import dataset_tools as dttools

dttools.download(dataset='BDD100K: Images 100K', dst_dir='~/dataset-ninja/')
```

Make sure not to overlook the [python code example](#) available on the Supervisely Developer Portal. It will give you a clear idea of how to effortlessly work with the downloaded dataset.

The data in original format can be [downloaded here](#).

• • •

Disclaimer

Our gal from the legal dep told us we need to post this:

Dataset Ninja provides visualizations and statistics for some datasets that can be found online and can be downloaded by general audience. Dataset Ninja is not a dataset hosting platform and can only be used for informational purposes. The platform does not claim any rights for the original content, including images, videos, annotations and descriptions. Joint publishing is prohibited.

You take full responsibility when you use datasets presented at Dataset Ninja, as well as other information, including visualizations and statistics we provide. You are in charge of compliance with any dataset license and all other permissions. You are required to navigate datasets homepage and make sure that you can use it. In case of any questions, get in touch with us at hello@datasetninja.com.