SEG - Cityspace

07:00 July 25, 2023

1. Image Features

1.1. General Statistics

	Train	**Validation**
Images	**2968**	**500**
Classes	**20**	**20**
Classes in use	**20**	**20**
Annotations	**241553**	**46644**
Annotations per images	**81.39**	**93.29**
Images with no annotations	**0**	**0**
Median image resolution	**1024x2048**	**1024x2048**
Smallest annotation	**9**	**9**
Largest annotation	**2094081**	**2094081**
Most annotations in an image	**291**	**244**
Least annotations in an image	**5**	**6**

1.2. Image Width and Height Distribution

These histograms depict the distributions of image height and width. It's important to note that if certain images have been rescaled or padded, the histograms will represent the size after these operations.

1.3. Color Distribution

Here's a comparison of RGB or grayscale intensity intensity (0-255) distributions across the entire dataset, assuming RGB channel ordering.
It can reveal discrepancies in the image characteristics between the two datasets, as well as potential flaws in the augmentation process.
E.g., a notable difference in the mean value of a specific color between the two datasets may indicate an issue with the augmentation process.

1.4. Image Brightness Distribution

This graph shows the distribution of the brightness levels across all images.
This may for instance uncover differences between the training and validation sets, such as the presence of exclusively daytime images in the training set and nighttime images in the validation set.

2. Segmentation Features
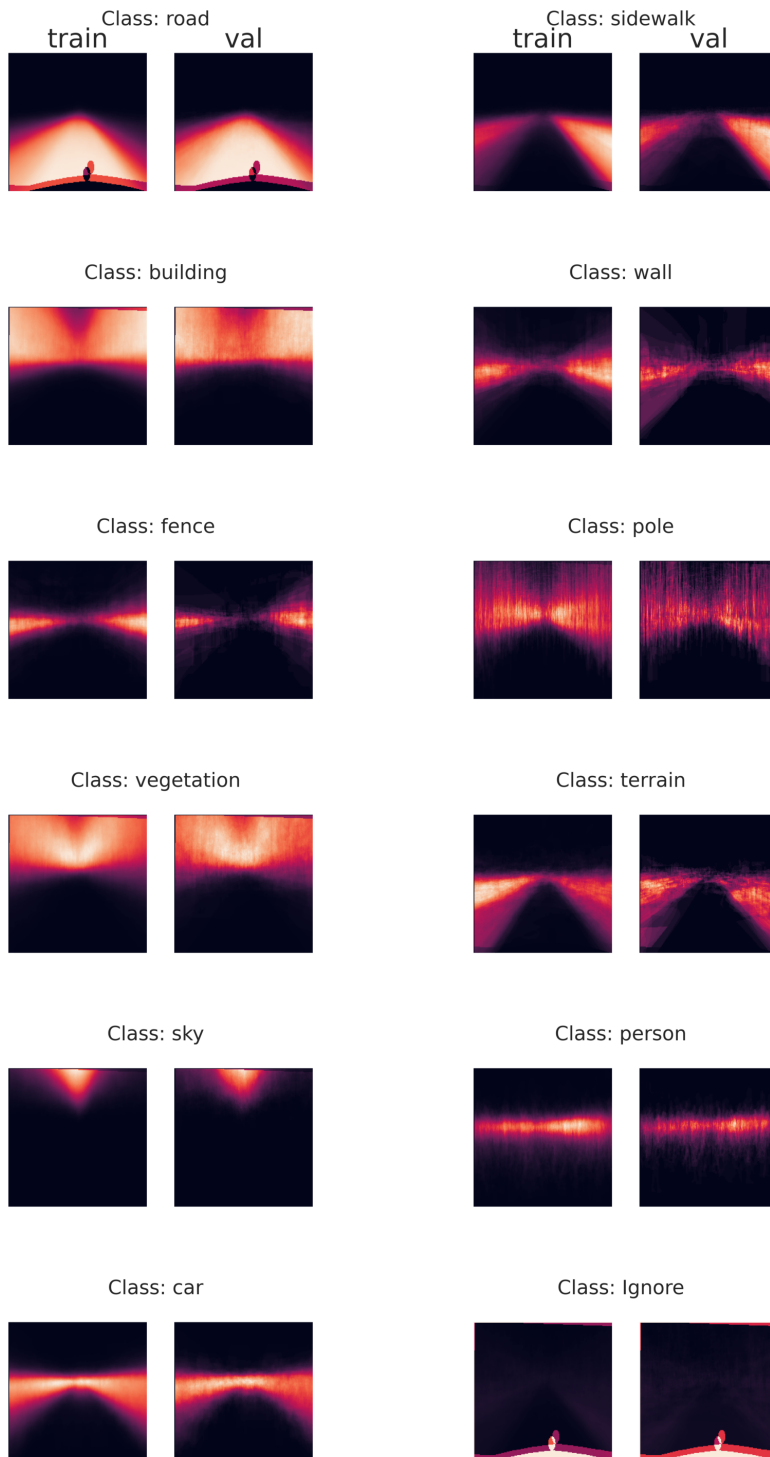
2.1. Visualization of Samples

train

val

The sample visualization feature provides a visual representation of images and labels. This visualization aids in understanding of the composition of the dataset.

> ⓘ **Notice:** Only 9 random samples are shown.
> You can increase the number of images by changing `n_cols` and `n_rows` in the configuration file.

2.2. Objects Density

The heatmap represents areas of high object density within the images, providing insights into the spatial distribution of objects. By examining the heatmap, you can quickly detect whether objects are predominantly concentrated in specific regions or if they are evenly distributed throughout the scene. This information can serve as a heuristic to assess if the objects are positioned appropriately within the expected areas of interest.

> **Notice:** Only the 12 classes with highest density are shown.
> You can increase the number of classes by changing `n_cols` and `n_rows` in the configuration file.

2.3. Class Frequency

This bar plot represents the frequency of appearance of each class. This may highlight class distribution gap between training and validation splits. For instance, if one of the class only appears in the validation set, you know in advance that your model won't be able to learn to predict that class.

2.4. Distribution of Class Frequency per Image

This graph shows how many times each class appears in an image. It highlights whether each class has a constant number of appearances per image, or whether there is variability in the number of appearances from image to image.

2.5. Distribution of Objects per Image

These graphs shows how many different objects appear in images.
This can typically be valuable to know when you observe a very high number of objects per image, as some models include a parameter to filter the top k results.

2.6. Distribution of Object Width and Height

These heat maps illustrate the distribution of objects width and height per class.
Large variations in object size can affect the model's ability to accurately recognize objects.

2.7. Distribution of Object Area

This graph shows the frequency of each class's appearance in the dataset. This can highlight distribution gap in object size between the training and validation splits, which can harm the model's performance.
Another thing to keep in mind is that having too many very small objects may indicate that your are downsizing your original image to a low resolution that is not appropriate for your objects.

2.8. Object Convexity

This graph depicts the convexity distribution of objects in both training and validation sets.
Higher convexity values suggest complex structures that may pose challenges for accurate segmentation.

Notice: To better understand how to tackle the data issues highlighted in this report, explore our comprehensive course on analyzing computer vision datasets.