



OPEN

MAGES


CHALLENGE 2018




Closing remarks



Open Images V4 dataset



Open Images V4 dataset



Dataset paper coming soon

8

First Author, Second Author



Fig. 5 Examples of large number of annotated boxes in an image (348, 386, and 743, respectively). *GroupOf* could have been used in many of these cases, but nevertheless they still have interest in practice.



Fig. 6 Examples boxes labeled with the five available attributes: *GroupOf*, *Occluded*, *Depiction*, *Truncated*, and *Inside*.




Fig. 7 Number of boxes per class. The horizontal axis is each class sorted position, but represented in logarithmic scale for the sake of readability.

note that the boxes are annotated only with the *leafmost* label, e.g. men are not annotated as *person* also.)

At the other end of the spectrum, Open Images has 517 classes with a number of instances above that of the most infrequent one in COCO (198 instances), and 417 classes in the case of ILSVRC (502 instances).

Interactions between different classes are also a reflection of the richness of the visual world. To measure this, Figure 8 (left) counts the percentage of images with boxes coming from different number of object classes. We can see that Open Images and COCO have a much richer distribution of

images with co-occurring classes compared to ILSVRC and PASCAL, which are more biased to a single class per image.





Fig. 8 Number of co-occurring classes per image: Normalized histogram of the number of distinct classes per image.

Figure 8 (right) shows the unnormalized plot (number of images), where we can see that Open Images has at least one order of magnitude more boxes in all number of different classes per image than COCO. As an example, Open Images has around 1,000 images with exactly 14 different categories, while COCO has 20, ILSVRC stops at eleven classes per image and PASCAL at four.


Figure 9 displays two images with a large number of classes annotated (11 and 7, respectively), to illustrate the variety and granularity that this entails.

[Jasper]: May need to rewrite or remove after adding Visual Relationship Detection (which is on more or less the same topic). To further

Future addition: object segmentation annotations?



Future addition: object segmentation annotations?



Congratulations to the winners again!

Object Detection track

1st place: **kivajok**  **Baidu** 百度

2d place: **PFDet**  **Preferred Networks**

3d place: **Avengers**  **Baidu** 百度

Visual Relationship Detection track

1st place: **Seiji**



2d place: **tito (individual)**

3d place: **Kyle (no disclosure)**

4th place: **toshif (individual)**

Challenge organizers



Vittorio
Ferrari



Alina
Kuznetsova



Jordi
Pont-Tuset



Matteo
Malloci



Jasper
Uijlings



Jake
Walker



Rodrigo
Benenson

All contributors to the Open Images Dataset

Project Lead & Coordination

Vittorio Ferrari
Tom Duerig
Victor Gomes

Image collection

Ivan Krasin
David Cai

Image-level labels

Neil Alldrin
Ivan Krasin
Shahab Kamali
Tom Duerig
Zheyun Feng
Anurag Batra
Alok Gunjan

Bounding boxes

Hassan Rom
Alina Kuznetsova
Jasper Uijlings
Stefan Popov
Matteo Mallocci
Sami Abu-El-Haija
Vittorio Ferrari

Website & visualizer

Jordi Pont-Tuset

Classes & hierarchy

Chen Sun
Kevin Murphy
Tom Duerig
Vittorio Ferrari

Challenge

Vittorio Ferrari
Alina Kuznetsova
Jordi Pont-Tuset
Matteo Mallocci
Jasper Uijlings
Jake Walker
Rodrigo Benenson

Advisers

Andreas Veit
Serge Belongie
Abhinav Gupta
Dhyanesh Narayanan
Gal Chechik

Thanks to ...



Common Visual Data Foundation

For hosting the data through AWS
Special thanks to Tsung-Yi Lin



Figure Eight

For hosting the data



Kaggle

For hosting the competition

Thanks to our sponsors

