



Finding Four-Leaf Clovers: A Benchmark for Fine-Grained Object Localization


Gustavo Pérez*, Laura Bravo*, Alejandro Pardo*, Pablo Arbeláez

Finding Four-Leaf Clovers


Goal: to create a reliable benchmark for **fine-grained object localization** problems

Other fine-grained object localization problems:

Cancer diagnosis
(Medicine)



Identifying cells in mitosis
(Medicine)



Detecting infected crops
(Biology)



Finding a specific person in a crowd
(Industry)

Finding Four-Leaf Clovers

High intra-class variability



Low inter-class variability



Finding Four-Leaf Clovers



Highly un-balanced

Finding Four-Leaf Clovers



Finding Four-Leaf Clovers



The Four-Leaf Clover Dataset









Examples of the level of detail in segmentation annotations of the FLC dataset

The Four-Leaf Clover Dataset

General statistics	Trainval set	Test set
Total positive images	1,000	500
Total negative images	51,637	51,670
Total images	52,637	52,170
4-leaf clover instances	1,530	747
4-leaf clover leaves	6420	3,128
4-leaf clover pixels	1.0511%	1.2431%
4-leaf clover boundary pxls.	0.0608%	0.0719%


FLC dataset statistics. 4-leaf clover pixels and 4-leaf clover boundary pixels refer to the rate of the total of positive pixels over the total of pixels in the FLC dataset.

Challenges of the FLC Dataset

	Lighting and color variations
	Cast shadows
	Occlusion
	Leaf shape
	Leaf 3D orientation
	Different clover species

Tasks


Object Detection



Semantic Segmentation



Instance Segmentation



Object Parsing



Semantic Boundary Detection

Experiments

Object Detection						
Dataset	Method	mAP@.5(%)	mAP@[.5, .95] (%) all	mAP@[.5, .95] (%) small	mAP@[.5, .95] (%) medium	mAP@[.5, .95] (%) large
FLC test (Positives)	Mask R-CNN ResNet50+FPN	56.4	35.6	0.0	29.6	35.8
FLC test (Full)	Mask R-CNN ResNet50+FPN	6.20	4.10	0.0	0.00	8.0



(Top). results using Mask R-CNN for the task of detection.

(Left). This example shows one false positive (red), two false negatives (yellow) and a true positive (blue).

Semantic Segmentation

Dataset	Method	mIoU (%)
FLC (Positives)	MaskR-CNN R50	32.71
FLC	MaskR-CNN R50	7.71

(Left). Results for semantic segmentation.

(Bottom-left). Qualitative result on the test set.

(Bottom-right). Annotations of same image.



Experiments

Instance Segmentation

Dataset	Method	mAP@[.5,.95](%)
FLC (Positives)	MaskR-CNN R50	39.9
FLC (Full)	MaskR-CNN R50	4.7



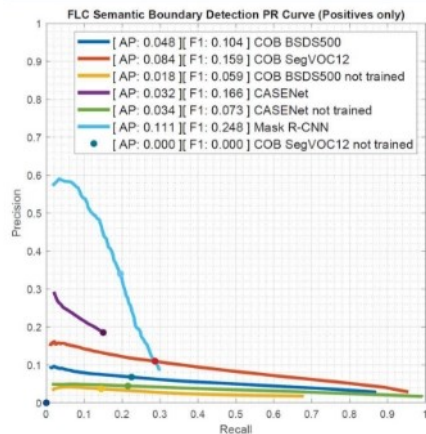
Object Parsing

Dataset	Method	mAP@[.5,.95] (%)
FLC (Positives)	MaskR-CNN R50	41.73
FLC (Full)	MaskR-CNN R50	1.3




Experiments

Semantic Boundary Detection



(Right). Semantic Boundary detection task PR curves.

(Right). (A) boundary annotation of a four-leaf clover. (B) Result of Mask R-CNN. (C) Result of the CASENet model. (D) Result of the trained COB model.





Thank you!

Comparison to Other Datasets

Dataset	Classification	Detection	Semantic Sem.	Instances	Boundaries	Parsing
Imagenet [1]	△	△	×	×	×	×
PASCAL [2]	△	△	△	△	△	×
MS-COCO [3]	△	△	×	△	×	×
DAVIS [4]	△	△	×	△	×	×
ADE20K [5]	×	△	△	△	×	△
CityScapes [6]	×	△	△	△	×	×
iNaturalist [7]	♣	△	×	×	×	×
Cats & Dogs [8]	♣	△	△	×	×	×
CUB-200 [9]	♣	△	△	×	×	×
CompCars [10]	♣	△	×	×	×	×
VegFru [11]	♣	×	×	×	×	×
CDVCE [12]	♣	△	×	×	×	×
FLC	×	♣	♣	♣	♣	♣

Comparison of FLC to major visual recognition datasets. Club (♣) indicates that a dataset allows to study a recognition problem at a fine-grained level, triangle (△) indicates that the version of the problem is not fine-grained, and (×) indicates that a dataset does not allow to study a problem.