

Hang Zhao

Massachusetts Institute of Technology



Bolei Zhou
MIT



Hang Zhao
MIT



Xavier Puig
MIT



Zhiding Yu
CMU




Sanja Fidler
University of Toronto



Antonio Torralba
MIT

Last year's challenge at ECCV



Scene Parsing Challenge 2016

- Each pixel of the image is classified into some classes

Scene parsing

semantic mask

class label

sky
road
sidewalk
escalator
building
car
bus
streetlight


ADE Dataset



- ceiling
- wall
- wall
- window (arch)
 - ↳ pane (glass)
 - ↳ figurine
- door frame
- double door
 - ↳ door
 - ↳ handle
- tray
- figurine
- refrigerator (crop)
- cabinet
 - ↳ door
 - ↳ knob
 - ↳ door
 - ↳ knob
 - ↳ door
 - ↳ knob
- jar
- cabinet
 - ↳ door
 - ↳ knob
- cabinet
 - ↳ door
 - ↳ knob
- cabinet
 - ↳ door
 - ↳ knob
- cabinet
 - ↳ door
 - ↳ knob
- microwave
 - ↳ door
 - ↳ window
 - ↳ button (door release)
- outlet
- pot
- sink
 - ↳ faucet
- soap_dispenser
- spice_rack
- coffee_maker
- knife_set
- knife_set
- range
 - ↳ button_panel
 - ↳ dial
 - ↳ dial
 - ↳ dial
 - ↳ dial
 - ↳ screen_time
 - ↳ stove
 - ↳ burner
 - ↳ burner
 - ↳ burner
 - ↳ burner
 - ↳ oven
 - ↳ door
 - ↳ handle
- toaster
- blender
- pot
- box
- worktop
- cabinet
 - ↳ drawer
 - ↳ knob
- jar
- salt_cellar
- worktop
- paper_towels
- dishwasher
- cabinet
- cabinet
- bottle_rack
- napkin_rack
- kitchen_island
- glass (wine)
- glass (wine)
- coasters
- bowl
- bowl
- trash_can
- dog_dish
- dog_dish
- chair
 - ↳ back
 - ↳ seat (fabric)
 - ↳ leg
 - ↳ leg
 - ↳ leg
- chair
 - ↳ back
 - ↳ seat (fabric)
 - ↳ leg
 - ↳ leg
 - ↳ leg
- chair
 - ↳ back
 - ↳ seat (fabric)
 - ↳ leg
 - ↳ leg
 - ↳ leg
- chair
 - ↳ back
 - ↳ seat (fabric)
 - ↳ leg
 - ↳ leg
 - ↳ leg
- side_table (crop)
- rug
- sofa (crop)
- cushion
- cushion
- cushion
- floor (tile)
- carpet
- bowl
- light_switch
- picture (map)

New tasks this year...


Scene Parsing





Instance Segmentation



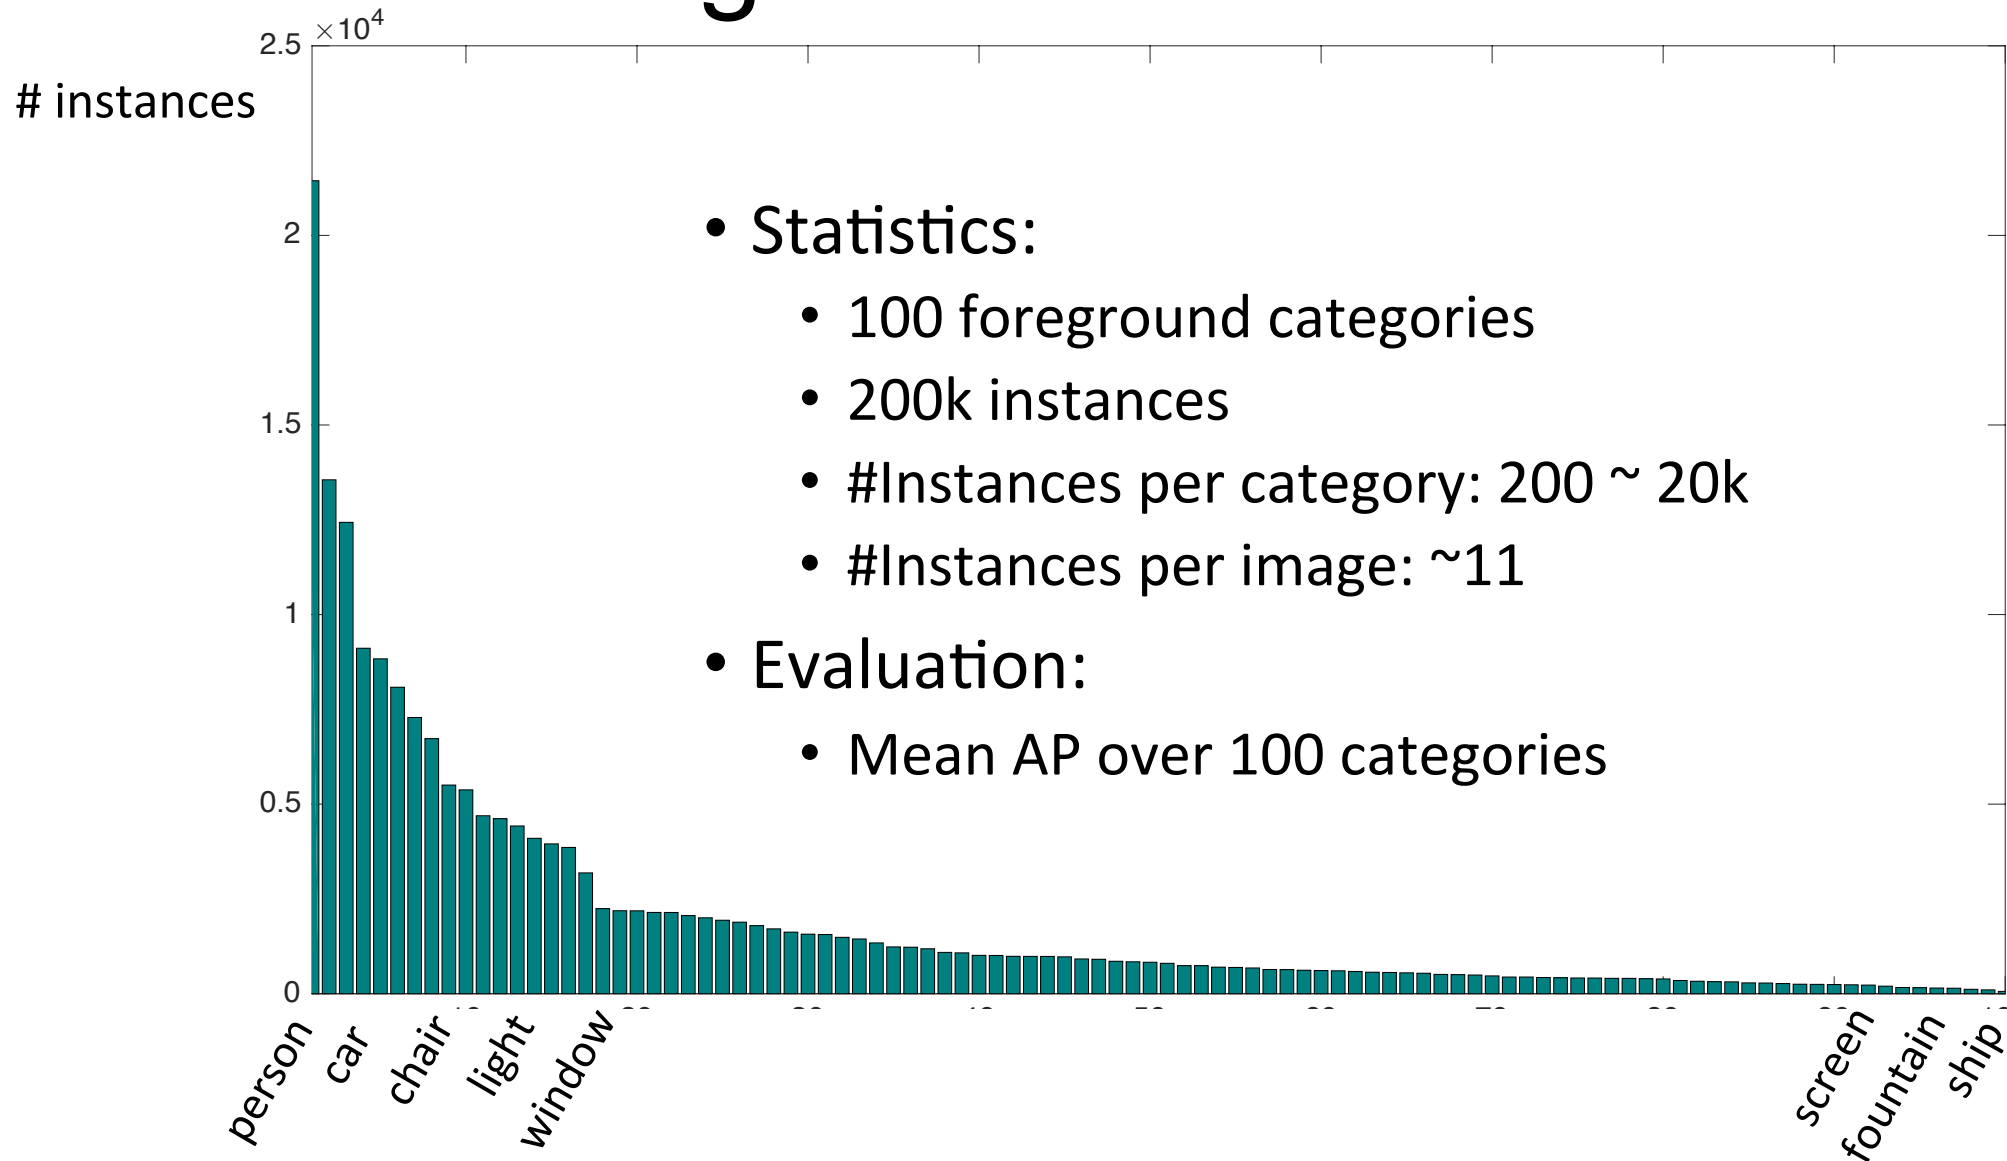
Semantic Boundary Detection



Instance Segmentation



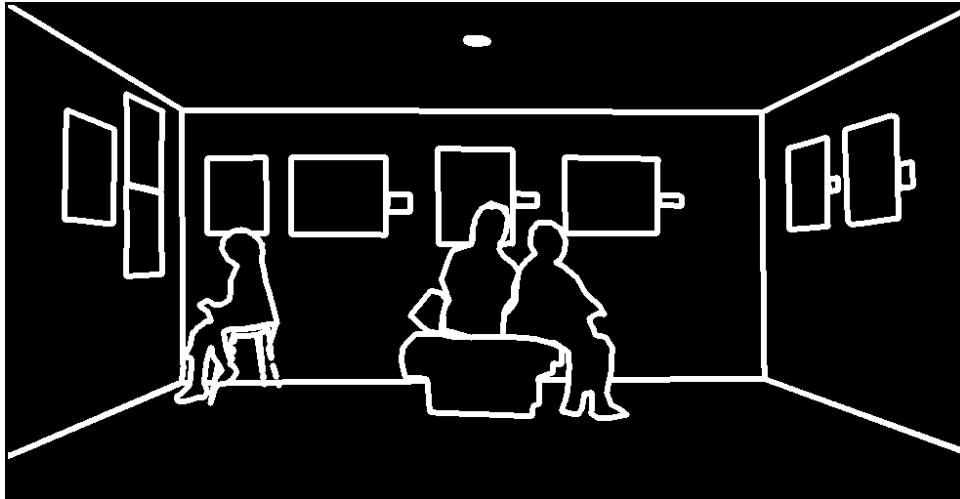
Instance Segmentation



Semantic Boundary Detection




Perceptual Edges




Semantic Boundary Detection




Binary Edges



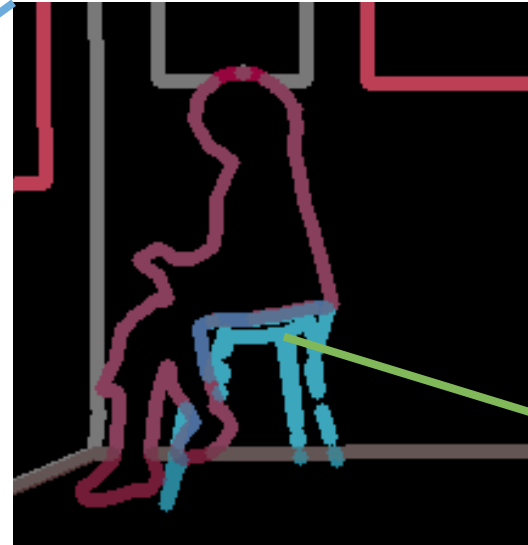
Category-aware Semantic Boundaries




Semantic Boundary Detection



Boundaries reveal object interactions





Boundary between:
person and chair



Boundary between:
person and sofa

Semantic Boundary Detection

- 150 categories
- Inter-class boundaries
- Intra-class boundaries
- Evaluation:
 - F-measure at optimal dataset scale (F-ODS)



Results – Scene Parsing

Final Score = (mean IoU + pixel accuracy) / 2

Team Name	Final Score
CASIA_IVA_JD	0.5547
WinterIsComing	0.5544
xdliang	0.5447
lijc15	0.5294
...	...
baseline-DilatedNet	0.4567

CASIA_IVA_JD
Jun Fu, Jing Liu, Longteng Guo, Haijie Tian, Fei Liu, Yong Li, Yongjun Bao, Weipeng Yan, Hanqing Lu
CASIA IVA and JD.com

WinterIsComing
Riwei Chen, Qi Chen, Xinglong Wu, Yifan Lu, Yudong Jiang, Linfu Wen
Toutiao AI Lab, ByteDance

* We received 30 submissions from 10 teams in this task. Full ranking can be found on places challenge website.


Sample Results

Scene Parsing

Image



Ground Truth




Prediction



Image



Ground Truth



Prediction



Results - Instance Segmentation

Team Name	mean AP
Megvii (Face++)	0.2977
G-RMI	0.2415
BlueSky	0.1551
...	...
baseline	0.200

Megvii (Face++)
*Tete Xiao, *Ruixuan Luo, *Borui Jiang,
Shuai Shao, Yuning Jiang, Yadong Mu,
Jieqi Shi, Chi Zhang, Jian Sun
Megvii Research and Peking University

G-RMI
Alireza Fathi, Nori Kanazawa, Kai Yang,
Kevin Murphy
Google Research Machine Intelligence

* We received 15 submissions from 5 teams in this task. Full ranking can be found on places challenge website.

Sample Results

Instance Segmentation


Sample Results

Semantic Boundary Detection



Image



Ground Truth




Prediction






Image



Ground Truth



Prediction



Awards

Winner of Scene Parsing:

CASIA_IVA_JD

Jun Fu, Jing Liu, Longteng Guo, Haijie Tian, Fei Liu,
Yong Li, Yongjun Bao, Weipeng Yan, Hanqing Lu

CASIA IVA and JD.com

Prize: Titan Xp GPU
Sponsored by: NVIDIA



Awards

Winner of Instance Segmentation:

Megvii (Face++)

*Tete Xiao, *Ruixuan Luo, *Borui Jiang, Shuai Shao,
Yuning Jiang, Yadong Mu, Jieqi Shi, Chi Zhang, Jian Sun

Megvii Research and Peking University

**Prize: Titan Xp GPU
Sponsored by: NVIDIA**

