

CSC420: Tutorial on object detection

Bin Yang
17.11.2016

Outline

- What's object detection? And its challenges?
- Face detection (demo)
- General object detection
- 3D object detection (demo)

Outline

- What's object detection? And its challenges?
- Face detection (demo)
- General (multi-class) object detection
- 3D object detection (demo)

Object detection

- Problem definition: Where? What?
- Challenges: Speed issue; large appearance variance.

Outline

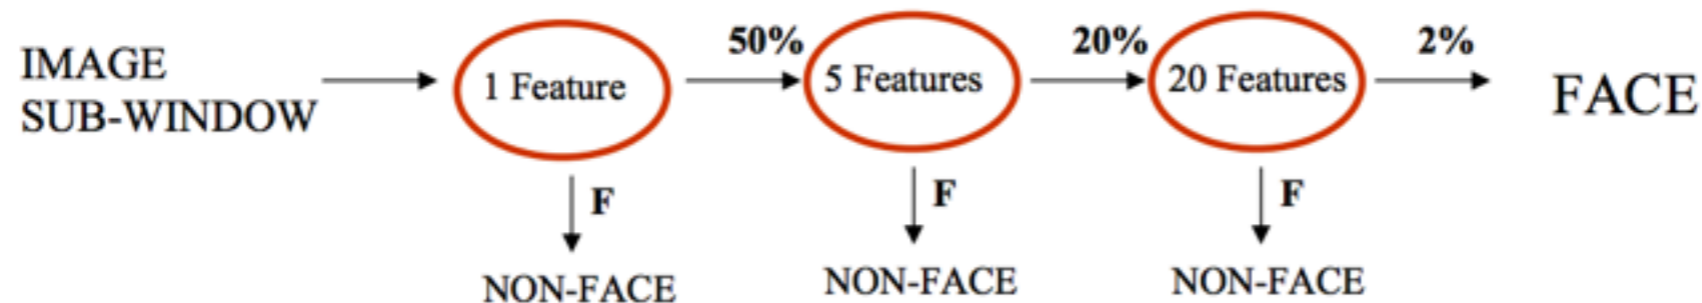
- What's object detection? And its challenges?
- Face detection (demo)
- General (multi-class) object detection
- 3D object detection (demo)

Face detection

- Binary classification: face / non-face
- Speed matters.




Viola/Jones face detector (2001, The Longuet-Higgins Prize in 2011)

Cascaded Classifier



- A 1 feature classifier achieves 100% detection rate and about 50% false positive rate.
- A 5 feature classifier achieves 100% detection rate and 40% false positive rate (20% cumulative)
 - using data from previous stage.
- A 20 feature classifier achieve 100% detection rate with 10% false positive rate (2% cumulative)

VJ face detection results



Today's image

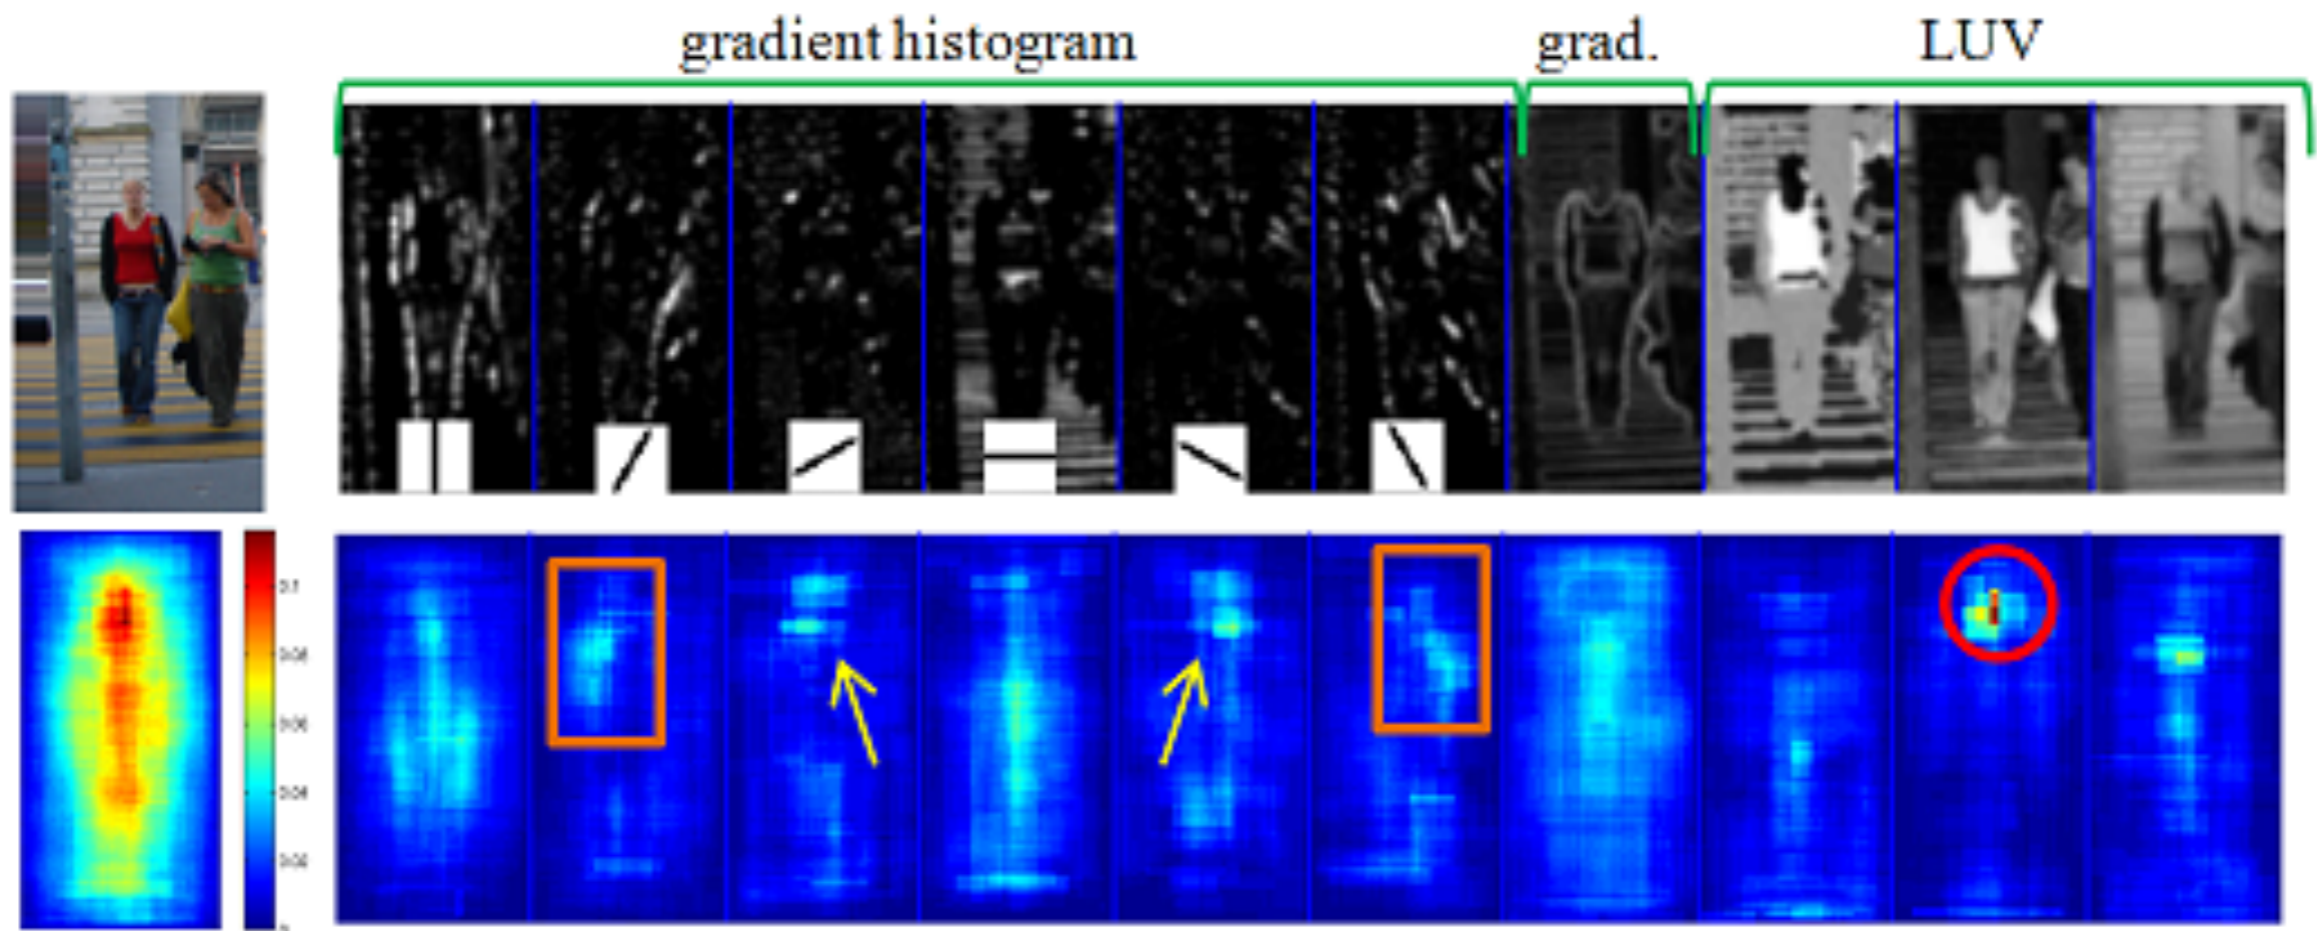


Resolution: 3815×2259

One example from AFW test set

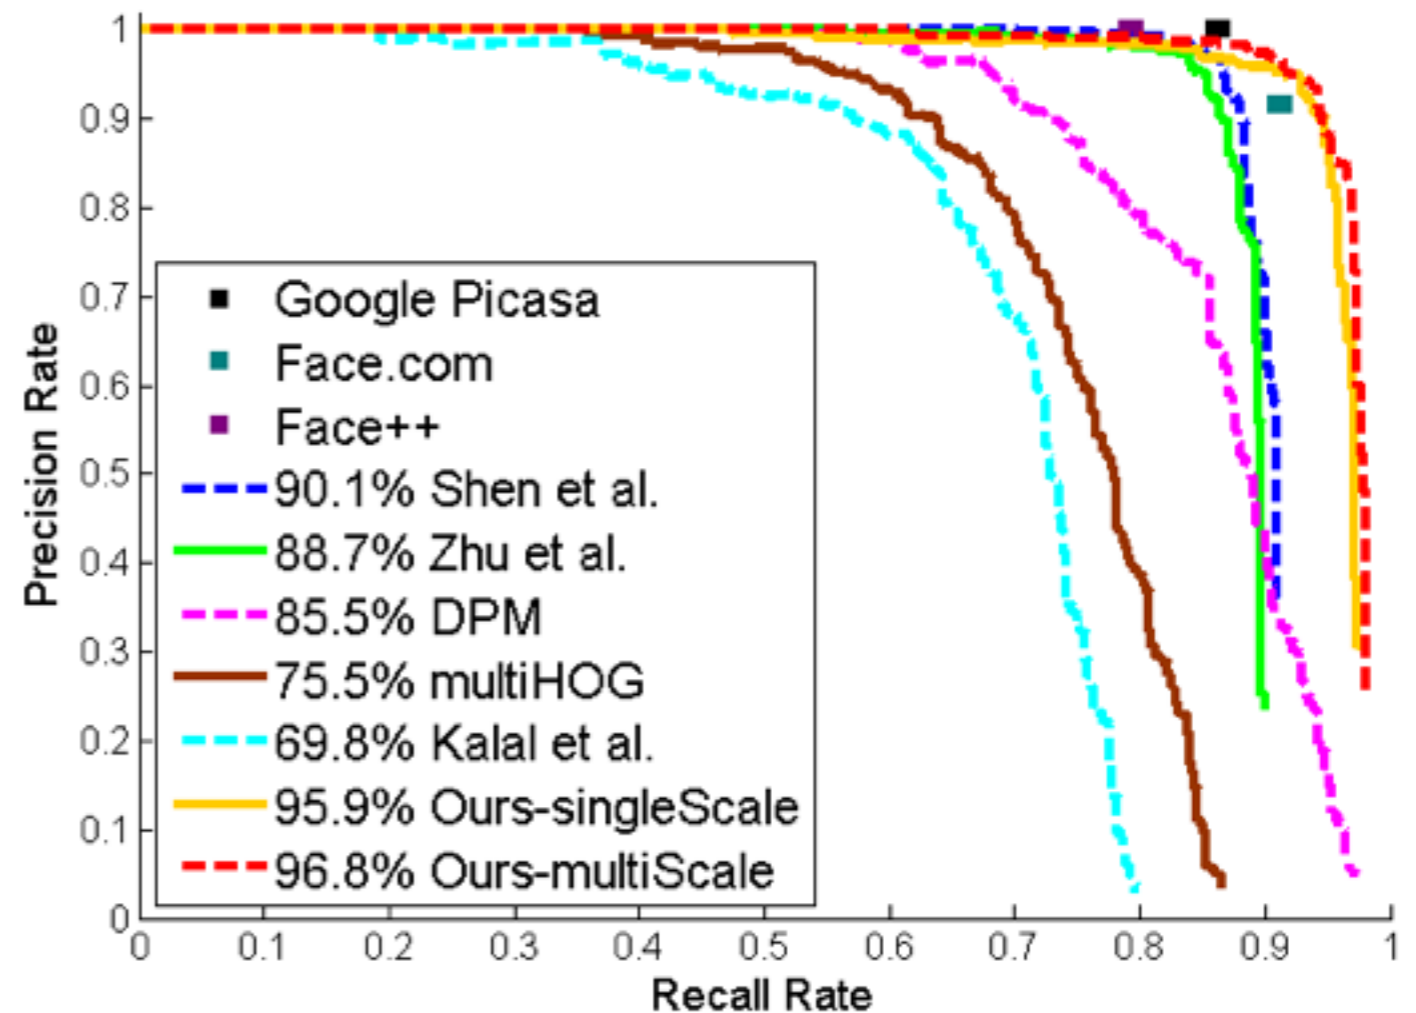

Aggregate channel features (P. Dollar, 2009)

- To handle more difficult scenario, we need more robust features with better representational power.




Aggregate channel features (B. Yang, 2014)

- Train view-specific detectors for each face view.



CNN based face detector (H. Qin, 2016)



Real-time demos

Reference: [http://www.cv-foundation.org/openaccess/
content_cvpr_2016/papers/
Qin_Joint_Training_of_CVPR_2016_paper.pdf](http://www.cv-foundation.org/openaccess/content_cvpr_2016/papers/Qin_Joint_Training_of_CVPR_2016_paper.pdf)


Outline

- What's object detection? And its challenges?
- Face detection (demo)
- General (multi-class) object detection
- 3D object detection (demo)

General object detection

- No longer binary classification.
- We need very good image representation.
- CNN+ImageNet provide us with that.

AlexNet on ImageNet (Alex, 2012)

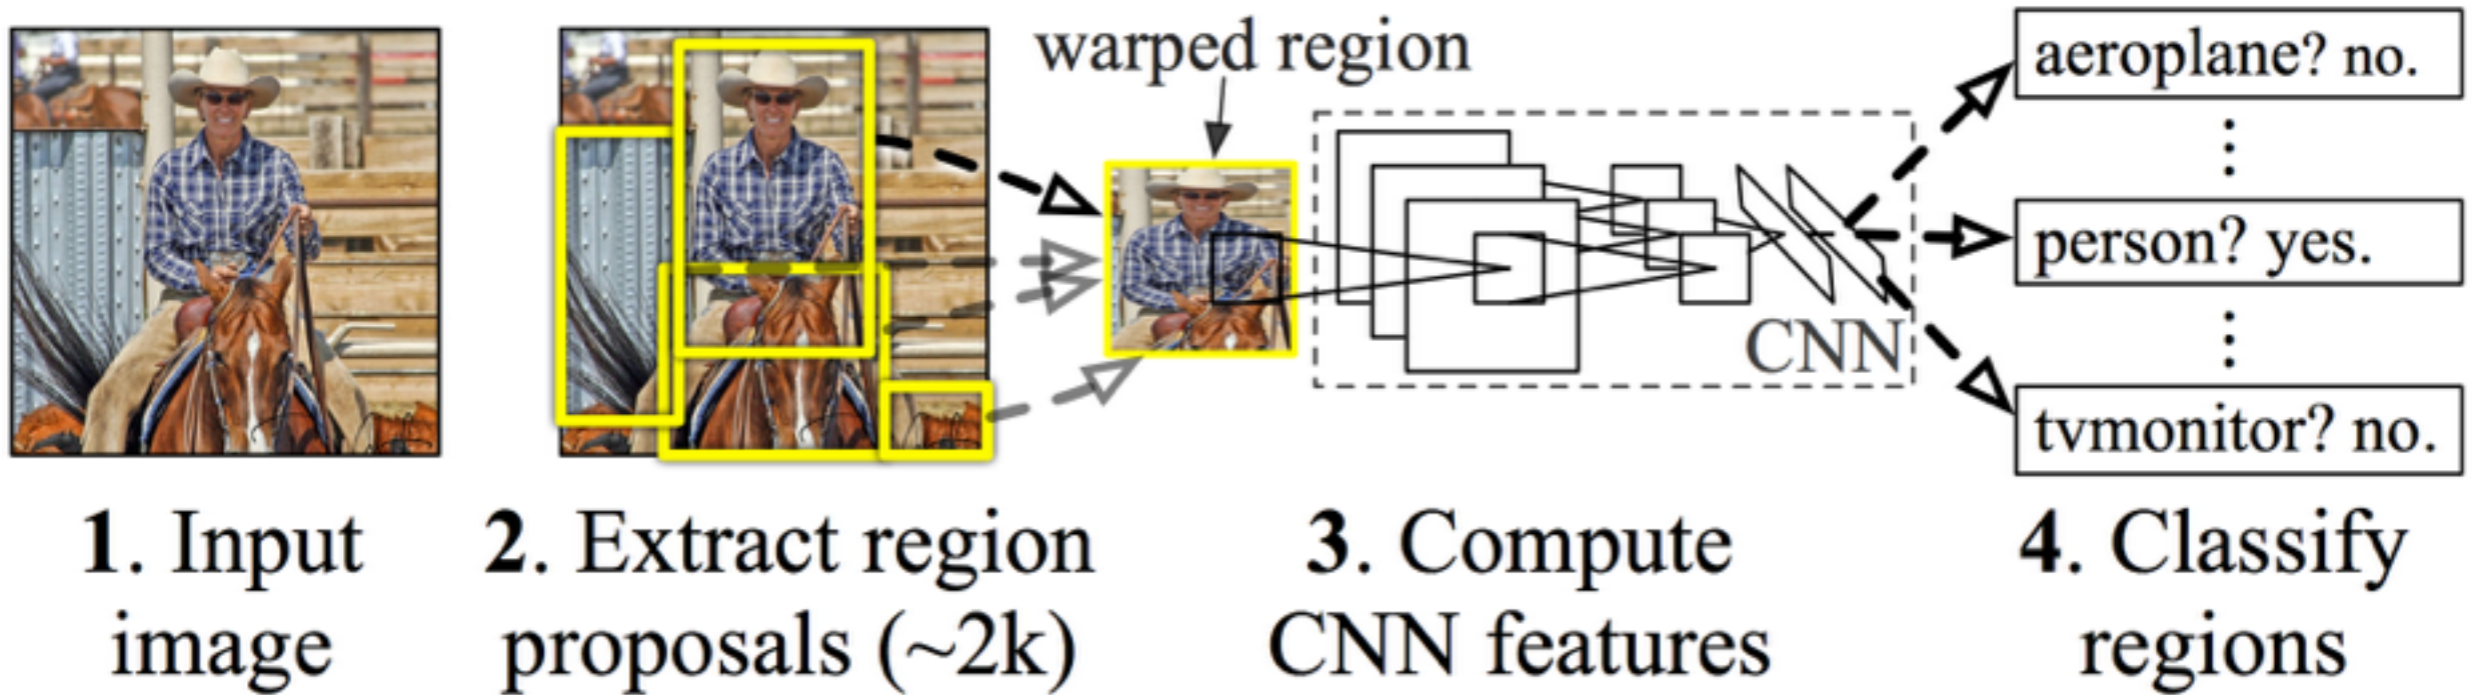


1000-class image classification:

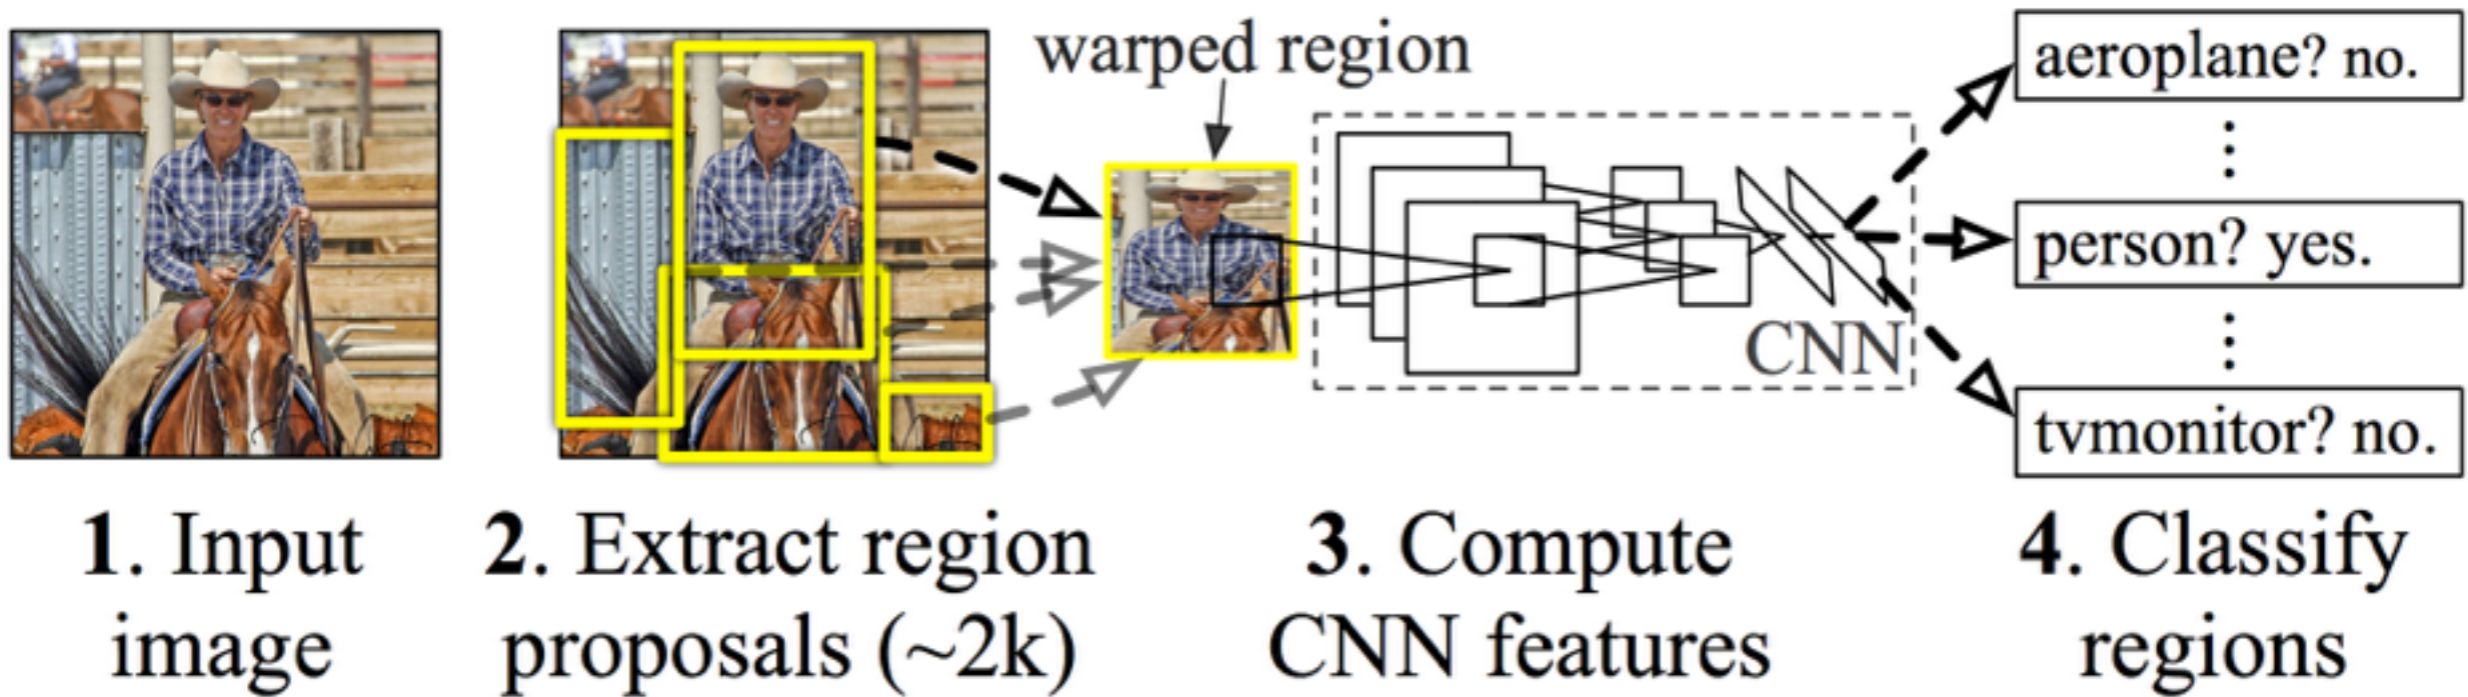
Top1 accuracy: 57.1%

Top5 accuracy: 80.2%

Region-CNN (R. Girshick, 2013)

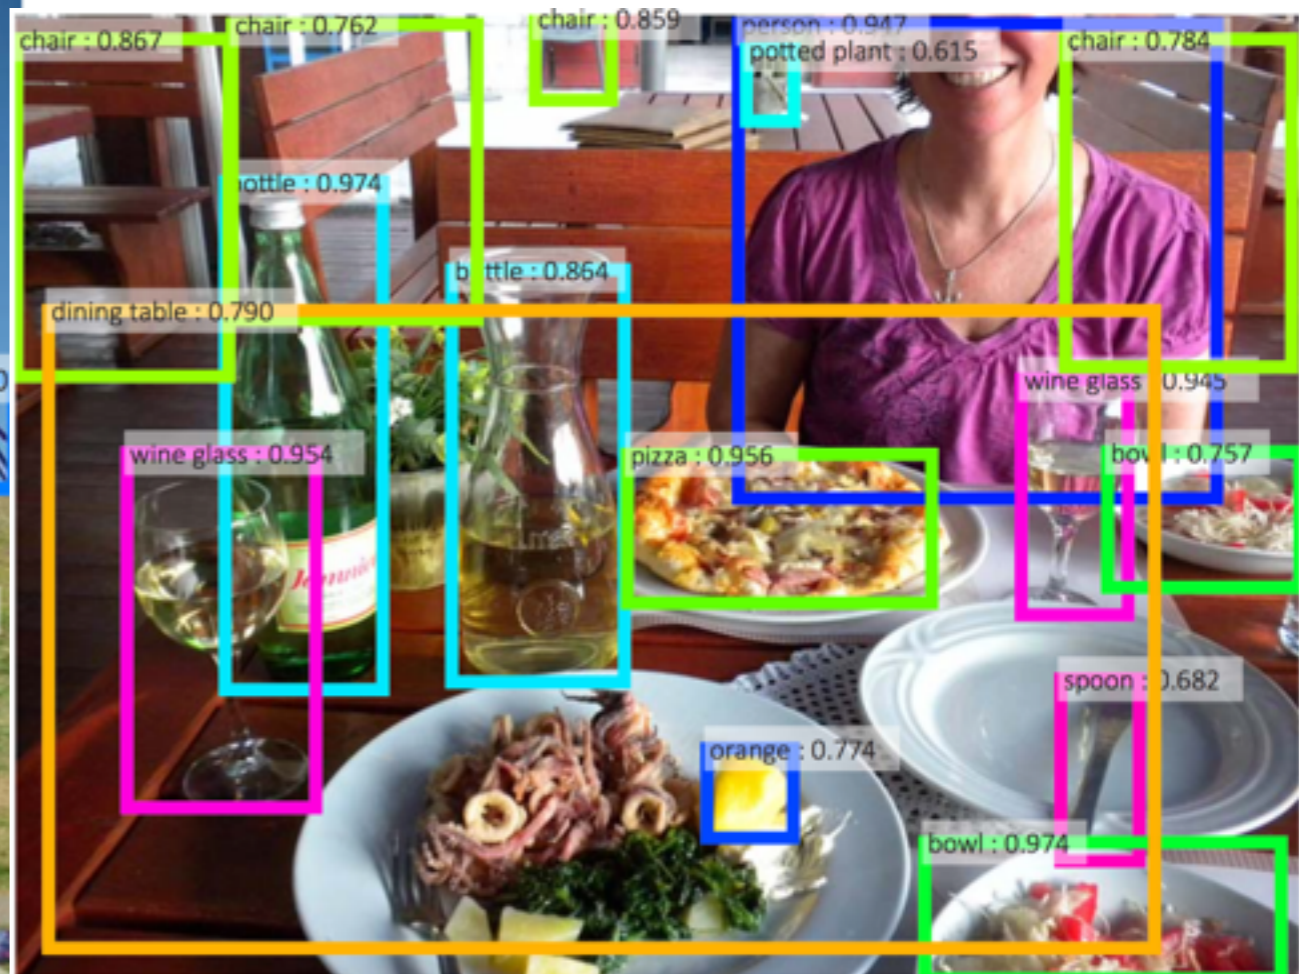
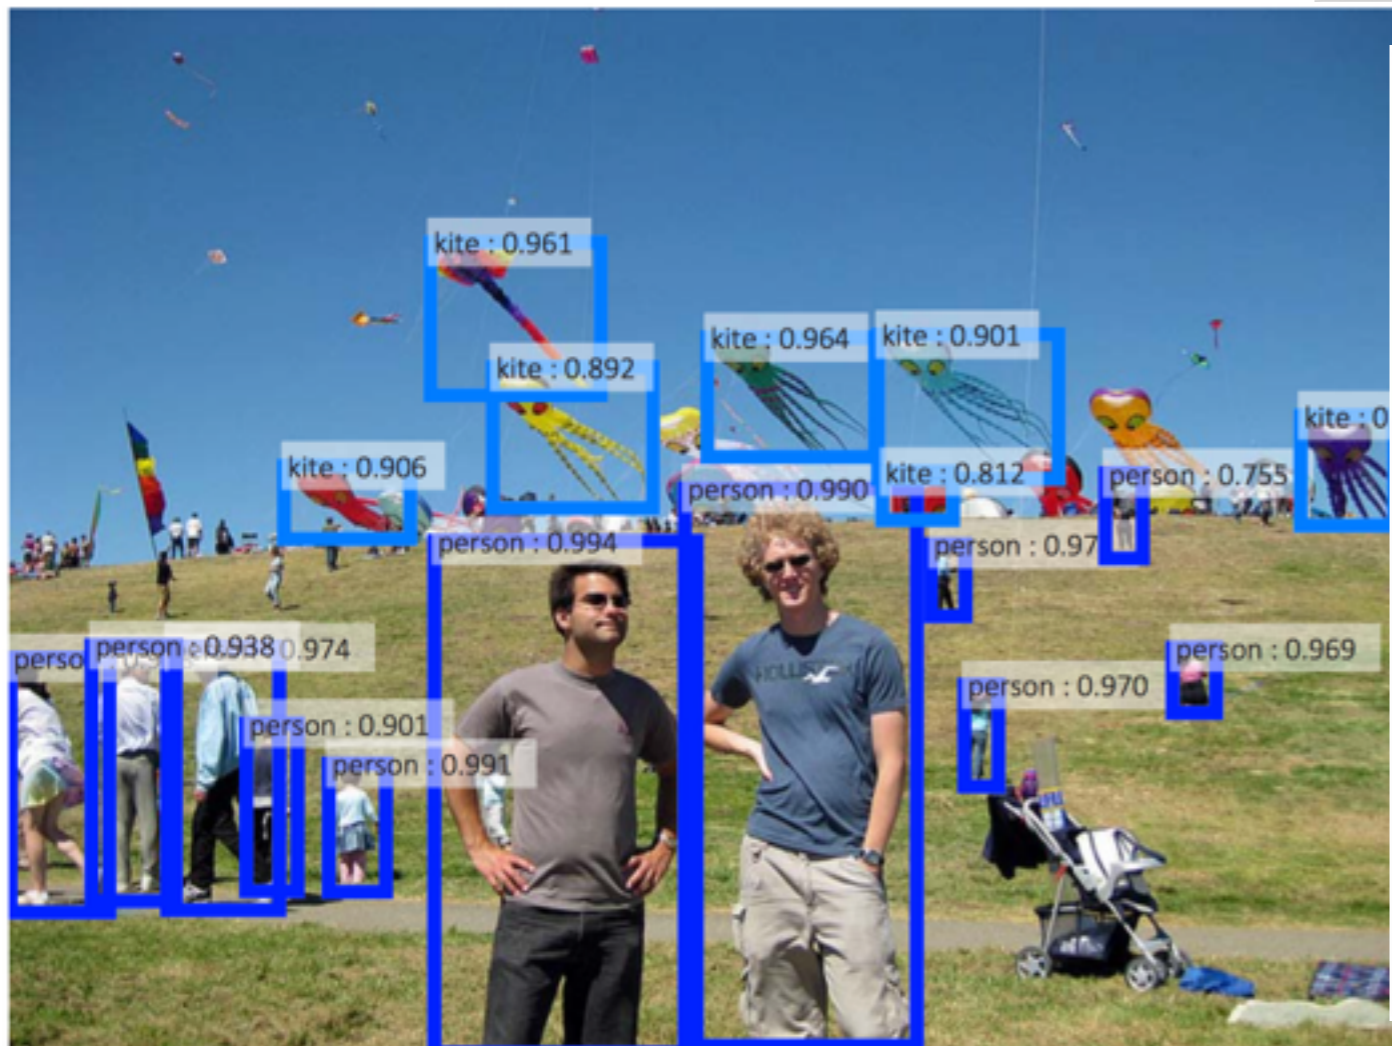



Region-CNN (R. Girshick, 2013)



Model	Architecture	Checkpoint	Top-1 Accuracy	Top-5 Accuracy
Inception-ResNet-v2	Code	inception_resnet_v2_2016_08_30.tar.gz	80.4	95.3
Inception V3	Code	inception_v3_2016_08_28.tar.gz	78.0	93.9
ResNet 152	Code	resnet_v1_152_2016_08_28.tar.gz	76.8	93.2
ResNet V2 200	Code	TBA	79.9*	95.2*

(*): Results quoted in ResNet paper.



Outline

- What's object detection? And its challenges?
- Face detection (demo)
- General (multi-class) object detection
- 3D object detection (demo)

Video demo

Link: http://3dimage.ee.tsinghua.edu.cn/files/XiaozhiChen/mono3d/cvpr16_demo.webm

Questions?

- What's object detection? And its challenges?
- Face detection (demo)
- General (multi-class) object detection
- 3D object detection (demo)