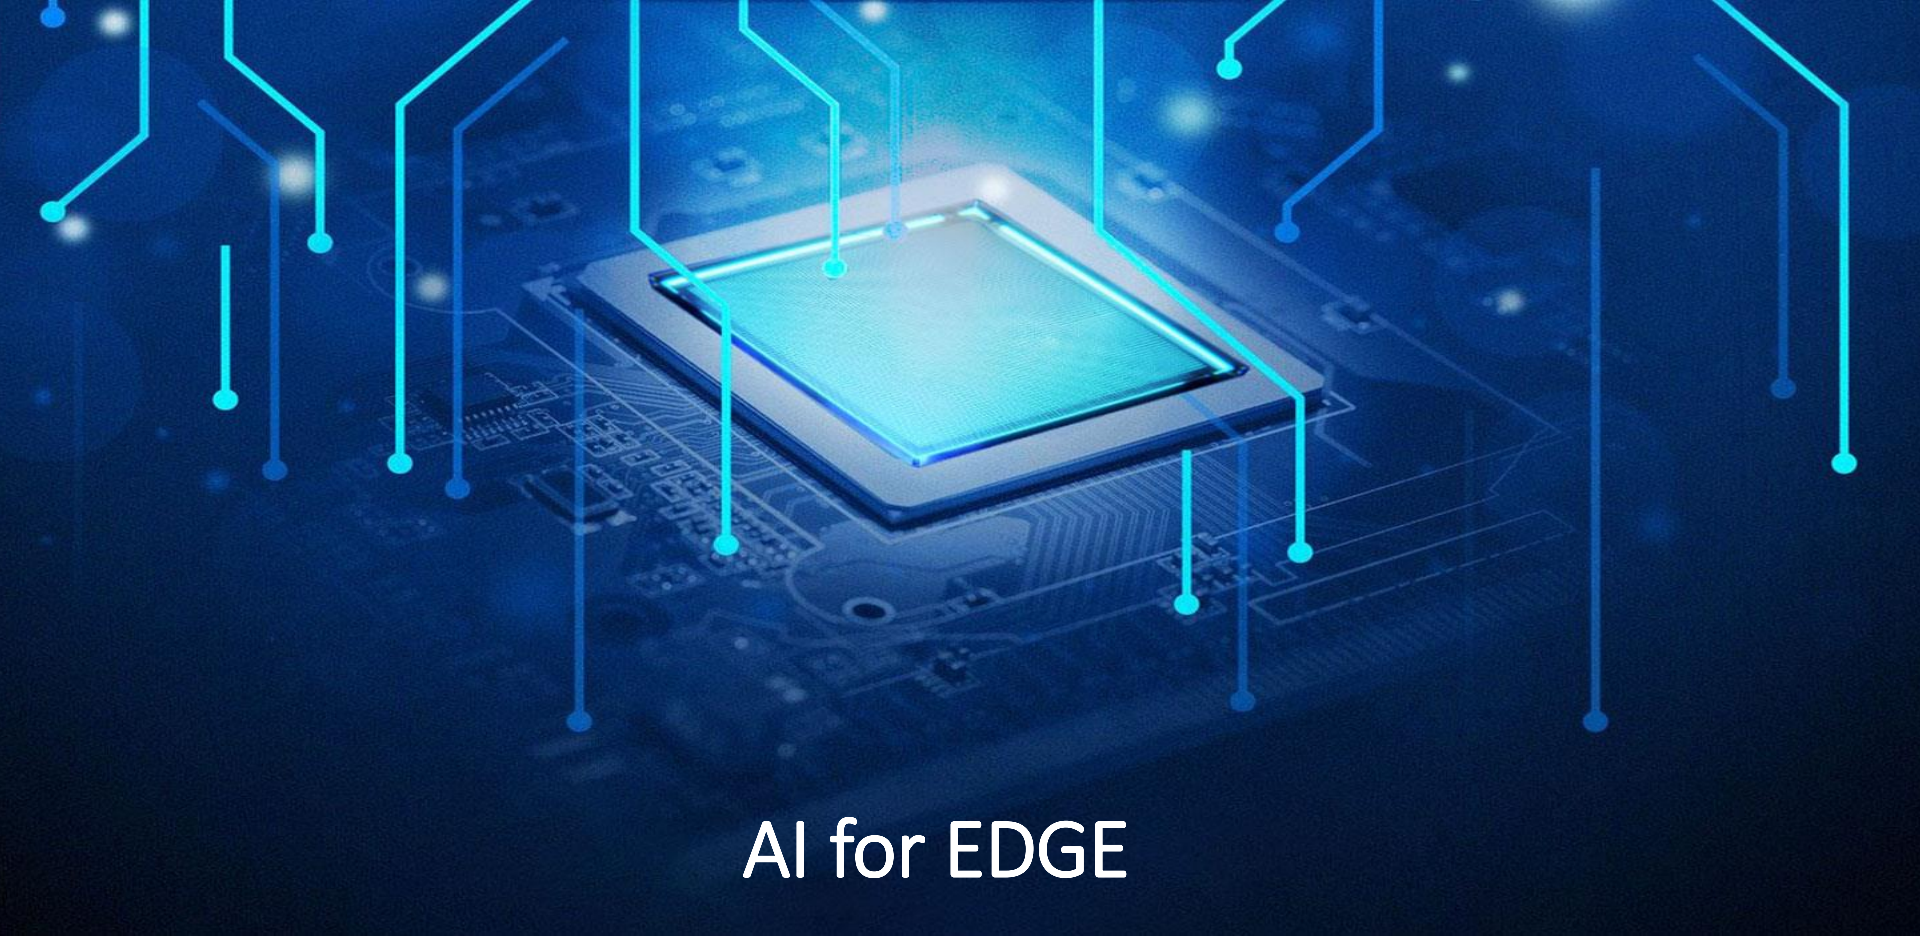

AI for EDGE

Object detection

Outline

- Computer vision
- Latest technological advances
- Object detection algorithms
- Object detection use cases

- Scientific field that deals with how computers can gain high-level understanding from digital images
- Understanding in this context means the transformation of visual images into descriptions of the world
- Sub-domains of computer vision:
  - object detection,
  - image classification,
  - video tracking,
  - motion estimation,
  - scene reconstruction,
  - 3D scene modeling, and
  - image restoration

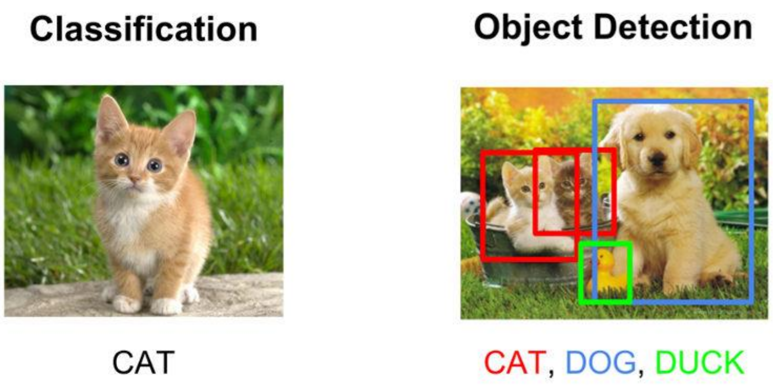

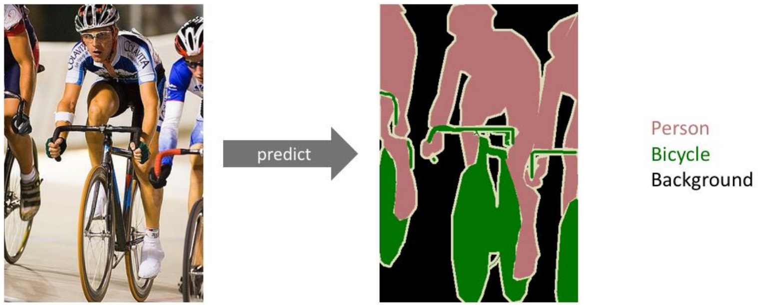

Computer vision

- Object detection is a key field in artificial intelligence, allowing computer systems to “see” their environments by detecting objects in images.
- Object detection applications include
  - Animal detection
  - Pedestrian detection
  - People counting
  - Vehicle detection
  - Face detection
  - Text detection
  - Number-plate recognition…

MediaPipe KNIFT: Template-based feature matching
https://google.github.io/mediapipe/solutions/knift.html

Latest advances

- In the last few years, the rapid advances of deep learning techniques have greatly accelerated the evolution of object detection.
- Cameras are smaller, cheaper and of higher quality
- Computing platforms moved toward parallelization through multi-core processing and GPU
- As a result, numerous real-world applications, such as healthcare monitoring, autonomous driving, video surveillance, anomaly detection
  are based on deep learning object detection.

Those advances enabled a key architectural concept called Edge AI.

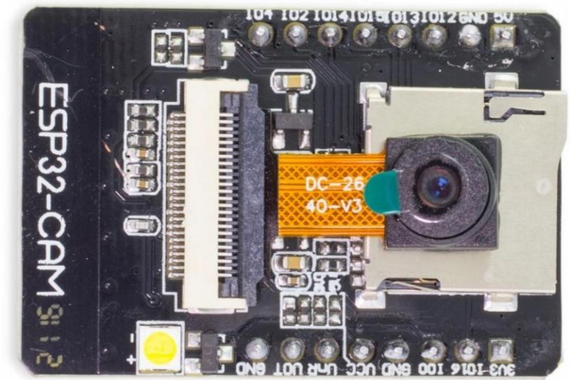

Object detection algorithms

Object detection can be performed using traditional(1) or modern(2) techniques:

- (1) image processing techniques – don't require historical data for training and are unsupervised in nature
- (2) deep learning methods – depend on supervised learning, the performance is limited by computation power (CPUs or GPUs)

Deep learning object detection is widely accepted by researchers and adopted to build commercial products.

python detect.py --weights yolov7.pt --conf 0.25 --img-size 640 --source inference/images/horses.jpg

Object detection algorithms

- Before 2014 – traditional object detection
  - Viola-Jones detector (2001) the pioneering work that started the development of traditional detection methods
  - HOG detector (2006) a popular feature descriptor for object detection
  - DPM (2008) first introduction of bounding box regression

- After 2014 – deep learning detection
  - Two-stage algorithms
    - R-CNN and SPPNet (2014)
    - Mask R-CNN (2017)
    - Pyramid Networks/FPN (2017)
    - G-RCNN (2021)
  - One-stage algorithms
    - YOLO (2016)
    - SSD (2016)
    - RetinaNet (2017)
    - YOLOv3 (2018)
    - YOLOv4 (2020)
    - YOLOR (2021)
    - YOLOv7 (2022)

Object detection algorithms

One-stage vs two-stage detectors

- Object detector solves two subsequent tasks:
  1. Find an arbitrary number of objects
  2. Classify every single object and estimate its size with a bounding box
- Single-stage detectors combine both tasks into one step (higher performance at the cost of accuracy)
- The main advantage of single-stage is that those algorithms are generally faster than multi-stage detectors and structurally simpler
- The most popular one-stage detector is YOLO

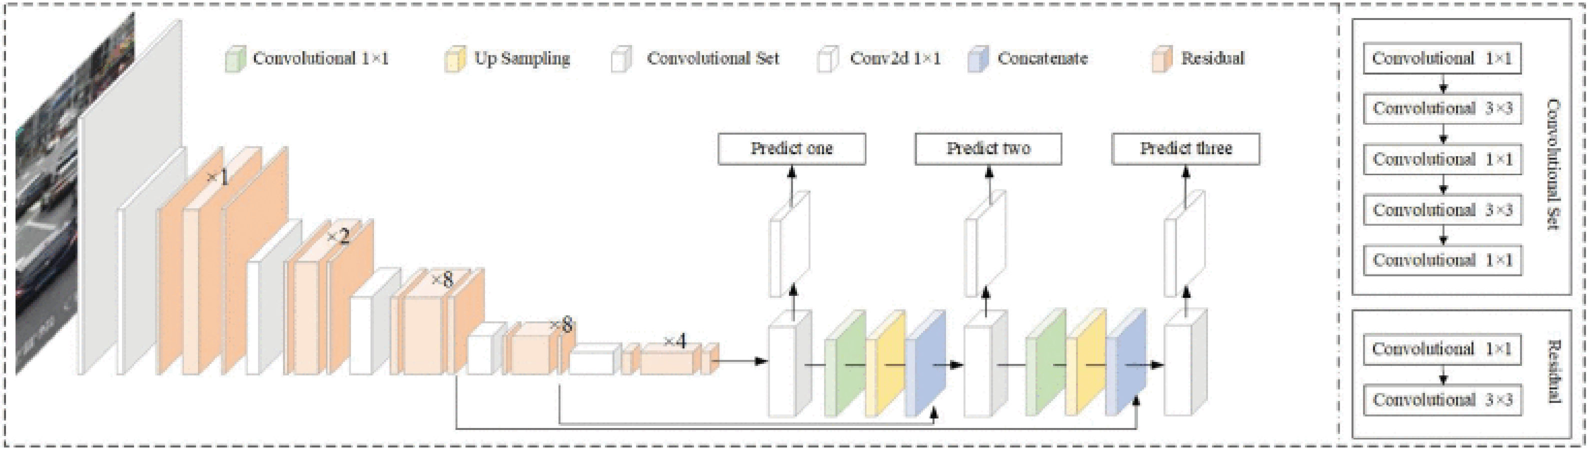

Mini-YOLOv3: Real-Time Object Detector for Embedded Applications, IEEE Access vol.7, September 2019

Object detection algorithms

YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors

Chien-Yao Wang, Alexey Bochkovskiy, Hong-Yuan Mark Liao

Performance (MS COCO)

Model	Test Size	Aptest	batch [1fps]	batch 32 average time
YOLOv7	640	51.4%	161 fps	2.8 ms
YOLOv7-X	640	53.1%	114 fps	4.3 ms
YOLOv7-W6	1280	54.9%	84 fps	7.6 ms
YOLOv7-E6	1280	56.0%	56 fps	12.3 ms
YOLOv7-D6	1280	56.6%	44 fps	15.0 ms
YOLOv7-E6E	1280	56.8%	36 fps	18.7 ms

Model	Parameters (million)	FPS	AP test (%)
YOLO7-Tiny	6.2	286	38.7
YOLOv7	36.9	161	51.4
YOLOv7-X	71.3	114	53.1
YOLOv7-W6	70.04	84	54.9
YOLOv7-E6	**97.2**	**56**	**56.0**
YOLOv7-D6	154.7	44	56.6
YOLOv7-E6E	151.7	36	56.8

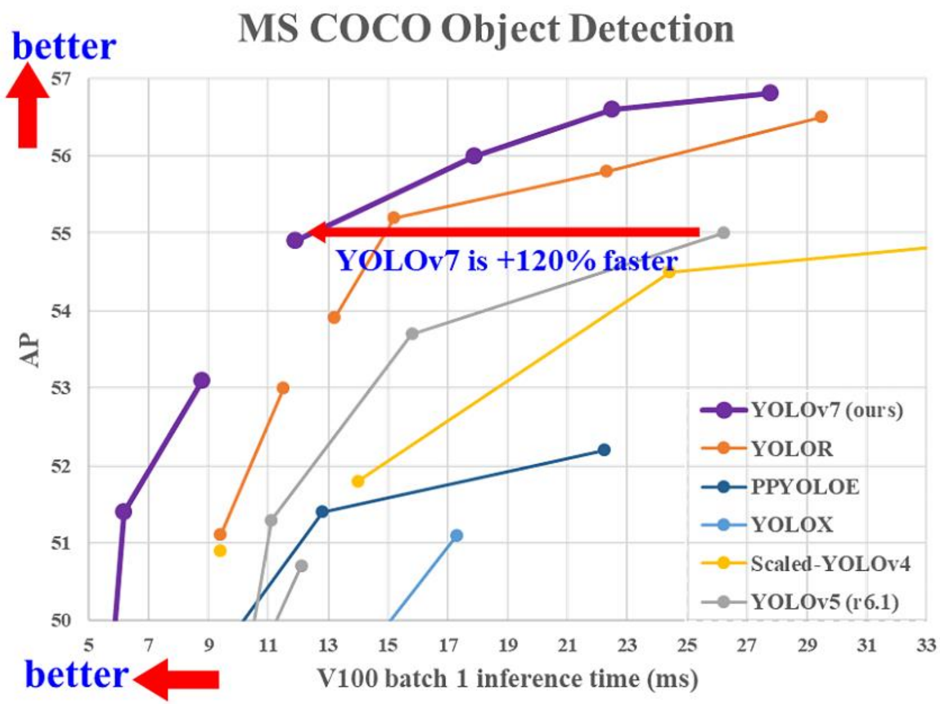

YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors

Chien-Yao Wang, Alexey Bochkovskiy, Hong-Yuan Mark Liao

Model	#Param.	FLOPs	Size	AP^{val}	AP_{50}^{val}	AP_{75}^{val}	AP_{S}^{val}	AP_{M}^{val}	AP_{L}^{val}
YOLOv4 [3]	64.4M	142.8G	640	49.7%	68.2%	54.3%	32.9%	54.8%	63.7%
YOLOR-u5 (r6.1) [81]	46.5M	109.1G	640	50.2%	68.7%	54.6%	33.2%	55.5%	63.7%
YOLOv4-CSP [79]	52.9M	120.4G	640	50.3%	68.6%	54.9%	34.2%	55.6%	65.1%
YOLOR-CSP [81]	52.9M	120.4G	640	50.8%	69.5%	55.3%	33.7%	56.0%	65.4%
YOLOv7	36.9M	104.7G	640	**51.2%**	**69.7%**	**55.5%**	**35.2%**	**56.0%**	**66.7%**
improvement	-43%	-15%	-	+0.4	+0.2	+0.2	+1.5	=	+1.3
YOLOR-CSP-X [81]	96.9M	226.8G	640	52.7%	**71.3%**	57.4%	36.3%	57.5%	68.3%
YOLOv7-X	71.3M	189.9G	640	**52.9%**	71.1%	**57.5%**	**36.9%**	**57.7%**	**68.6%**
improvement	-36%	-19%	-	+0.2	-0.2	+0.1	+0.6	+0.2	+0.3
YOLOv4-tiny [79]	6.1	6.9	416	24.9%	42.1%	25.7%	8.7%	28.4%	39.2%
YOLOv7-tiny	6.2	5.8	416	**35.2%**	**52.8%**	**37.3%**	**15.7%**	**38.0%**	**53.4%**
improvement	+2%	-19%	-	+10.3	+10.7	+11.6	+7.0	+9.6	+14.2
YOLOv4-tiny-3l [79]	8.7	5.2	320	30.8%	47.3%	32.2%	**10.9%**	31.9%	51.5%
YOLOv7-tiny	6.2	3.5	320	**30.8%**	**47.3%**	**32.2%**	10.0%	**31.9%**	**52.2%**
improvement	-39%	-49%	-	=	=	=	-0.9	=	+0.7
YOLOR-E6 [81]	115.8M	683.2G	1280	55.7%	73.2%	60.7%	40.1%	**60.4%**	69.2%
YOLOv7-E6	97.2M	515.2G	1280	**55.9%**	**73.5%**	**61.1%**	**40.6%**	60.3%	**70.0%**
improvement	-19%	-33%	-	+0.2	+0.3	+0.4	+0.5	-0.1	+0.8
YOLOR-D6 [81]	151.7M	935.6G	1280	56.1%	73.9%	61.2%	**42.4%**	60.5%	69.9%
YOLOv7-D6	154.7M	806.8G	1280	56.3%	73.8%	61.4%	41.3%	60.6%	70.1%
YOLOv7-E6E	151.7M	843.2G	1280	**56.8%**	**74.4%**	**62.1%**	40.8%	**62.1%**	**70.6%**
improvement	=	-11%	-	+0.7	+0.5	+0.9	-1.6	+1.6	+0.7

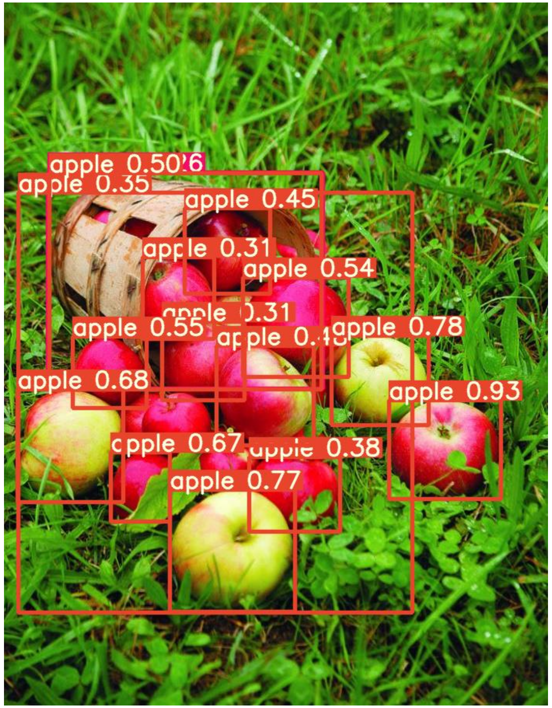

- Industrial PPE detection (personal protective equipment)
- Anomaly and defect detection in product assembly
- Autonomous driving
- Traffic monitoring and road maintenance
- People counting
- Parking occupancy
- Intrusion detection

Related topics

Solve the Mystery of Vehicle Detection Algorithm

https://www.mouser.mx/blog/mystery-of-vehicle-detection

Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors

https://huggingface.co/spaces/akhaliq/yolov7