## Hands-on with Artificial Intelligence on the Edge

### University of Bologna Cesena Campus

**Giacomo Bartoli** Machine Learning lab



# 1.Edge Al 2.CNNs as base model **3.Deep Architectures for Object Detection** 4.Google AIY Vision kit 5. Audi Autonomous Driving Cup



### **Edge Al** Running inference directly on the device

### Privacy

#### Images and videos are personal No one should access them

### Latency

Real time tasks are not possible using Machine Learning as a service



12,554 views | Apr 4, 2018, 08:48am

### Why AI At The Edge Is The Next Goldmine



9

in

Dan Woods Contributor (i)



Simon Crosby, CTO of SWIM.AI PRNEWSFOTO/SPOKEN In a fascinating article in MIT Technology Review, How the AI Cloud Could Produce the Richest Companies Ever, Peter Burrows argues that the next wave of AI will be a winner-take-all land grab. The company that creates what he calls the AI operating system in the cloud will win a huge prize, and the rest of the entrants will fall behind.







### **ARTIFICIAL INTELLIGENCE AT THE** EDGE

November 29, 2017

Imagine being able to...

- ... have a camera-enabled assistant monitor your aging parents to make sure they are alert and healthy
- ... autonomously watch for product imperfections in factories without human interference
- ... identify and locate lost hikers by using vision-enhanced drones to automatically send help
- ... automatically recognize your petsitter and let him or her into your house

These are just a few examples of how artificial intelligence (AI) at the edge, combined with connected



Naveen Rao

Vice President & General Manager, Artificial Intelligence **Products Group** 





### Machine Learning on the Edge



In a few years, the world will be filled with billions of small, connected, intelligent de homes, our cities, our vehicles, and our factories. Some of these devices will be carri of small computing devices will disrupt every industrial sector and play a key role in

MIPS) Almost all of them will use a variety of sensors to monitor their surroundings

Forbes



### The Paradigm-Changing Effects Of Al Innovation At The Edge



Intel AI BRANDVOICE

f 🎽 in

#### By Jason Compton

henever an innovation catches on in the data center, it's just a matter of time before it branches out to devices at the edge of computing. Business applications, storage and data processing have all grown in power and popularity on edge devices even as

their cloud and data center counterparts have continued to evolve.

The latest wave is the emergence of on-device artificial intelligence (AI). Instead of relying entirely on the cloud for AI insights, a new wave of specialized algorithms and chips is delivering deep insights wherever work is done. According to ABI Research, shipments of devices with edge AI capabilities will grow fifteenfold by 2023, to 1.2 billion units. The share of AI

## Edge Al



Fig: identifying bottlenecks by counting how many vehicles are on the road.

Al at the edge makes sense particularly when data is being generated at the same physical location where the decision needs to be made.



### Ideas for some cool projects:

- petsitter..).
- interference
- car system to detect when the driver is getting tired
- census of wild animals living in a remote area
- a smart traffic light
- eye tracking for ads insights
- where is Wally?

## Edge Al

Al at the edge makes sense particularly when data is being generated at the same physical location where the decision needs to be made.

automatically authorise access to your home (ex: babysitter, maid,

• autonomously watch for product imperfections in industries without human

drones able to distinguish between an adult and a child





Google AIY Vision Kit

> Amazon Deeplens

aws

e

### Edge Al Smart devices for developers:



Nvidia TEAL Drone

## Edge Al workflow



Performances must converge

### CNNs as base model

#### **Convolutional Neural Networks**



#### **Convolution and Pooling create a hierarchical representation of the image**

### From Image Classification to Object Detection

Object Detection is both classification and regression.



**Ex: R-CNN for Object Detection** 

## **Comparing Deep Architectures for Object Detection**

	mAP	Speed (FPS)	Speed (s/img)
R-CNN	62.4	.05 FPS	20 s/img
Fast R-CNN	70	.5 FPS	2 s/img
Faster R-CNN	78.8	7 FPS	140 ms/img
YOLO	63.7	45 FPS	22 ms/img
SSD	74.3	59 FPS	29 ms/img



#### https://aiyprojects.withgoogle.com/vision/

### **Google Vision Kit** Do-it-yourself intelligent camera



However, the model must be compiled for this specific device.



A Vision Processing Unit designed by Intel and intended for machine vision in low-power environments.



Trade-off between accuracy and efficiency.

Depthwise convolutions to reduce OPS.



Joy Detector Dish Classifier Face Recognition Dog/Cat/Human Object Recognition



...

20.80% Orange Category: Edible fruit

STREET, STREET



### **Creating a custom model** Pikachu detector



**Based on embedded\_ssd\_mobilenet** SSD\_Mobilenet with depth. multiplier = 0.125



#### **Transfer Learning**

Starting from a pre-trained model on Pascal VOC



#### **Real time scenario**

detection in order of milliseconds





### **Creating a custom model Tensorflow Object Detection APIs**

**model**: it defines what type of model will be trained (ie. metaarchitecture, feature extractor).

train\_config: decides what parameters should be used to train model parameters.

eval\_config: determines what set of metrics will be reported for evaluation.

train\_input\_config: defines what dataset the model should be trained on.

eval\_input\_config: defines what dataset the model will be evaluated on. Typically this should be different than the training input dataset.

```
model {
(... Add model config here...)
train_config : {
(... Add train_config here...)
train_input_reader: {
(... Add train_input configuration here...)
eval_config: {
eval_input_reader: {
(... Add eval_input configuration here...)
```

#### https://github.com/tensorflow/models/tree/master/research/object\_detection





### **Visual Results Pikachu detector**



### **Creating a custom model** Pikachu detector on the smartphone



### **Based on SSD\_mobilenet** SSD\_Mobilenet with depth. multiplier = 1



#### **Transfer Learning**

Starting from a pre-trained model on COCO



#### **Real time scenario**

detection in order of milliseconds



### Visual Results Pikachu detector





### **Comparing performances** K-Fold Cross Validation

CONFIG	embedded_ssd_mobiler
Α	0,5692
В	0,5873
С	0,5902
D	0,5733
E	0,6084
avg	0,585 mAP ~ <b>59%</b>

ן	e	t
---	---	---

CONFIG	ssd_mobilenet
Α	0,7916
B	0,7675
С	0,7795
D	0,7982
E	0,7128
avg	0,769 mAP ~ <b>77%</b>



### Audi Autonomous Driving Cup 2018

Every year the famous car manufacturing company Audi organises the Audi Autonomous Driving Cup to test new technologies in the automotive field. This is the first year that the competition is open to teams coming from outside Germany. After a careful selection based on projects submitted from all European universities, only 10 teams were chosen to access the finals. Among these, there is also a team represented by the University of Bologna. The competition involves several challenges, but most of them require **Object Detection** for solving tasks such as avoiding pedestrians, recognising road signs, detecting zebra crossings, allowing emergency vehicles to pass.

https://www.audi-autonomous-driving-cup.com





### AADC 2018 Al-related driving tasks



#### **Pedestrian detection**

The car must be able to identify a pedestrian and stop before hitting him/her.



#### **Child detection**

If a child is detected, the speed shall be reduced and this shall be indicated by the brake lights. For adults, no actions are required.



#### **Cars detection**

Being able to distinguish a normal car from an emergency vehicle equipped with sirens.



#### **Emergency vehicle detection**

The car must yield the right-of-way to any emergency vehicle using a siren and emergency lights. Drive to the right edge of the road and stop until the emergency vehicle has passed.





### **One Neural Network to rule them all**

A single CNN, based on SSD\_Mobilenet, capable of detecting adults, children, emergency cars and regular vehicles.

### How to deal with emergency vehicles?

The idea is to train the neural network to detect flashing sirens. This way, the model will be able to distinguish normal cars. If the car is equipped with flashing sirens then it will be classified as emergency vehicle.



## AADC, Al-related driving tasks Performances







### **Inference time**

#### PascalBoxes\_PerformanceByCategory/AP@0.5IOU/e

Img	Processing tir
1	11,234 ms
2	11,265 ms
3	11,491 ms
4	11,538 ms
5	11,332 ms
avg	<b>11,372 ms</b>



### AADC, Al-related driving tasks Visual results available on Vimeo!



https://vimeo.com/giacomobartoli

## Questions?!

