



Single-Shot Visual Object Detectors on Nano-Drones


Luca Bompani*, Lorenzo Lamberti*, Victor Javier Kartsch*,
Manuele Rusci^, Daniele Palossi^§, Luca Benini*§



DEMO

*University of bologna, ^Ku Leuven, †IDSIA USI-SUPSI, §ETH Zurich

Introduction




Running object detectors on nano-drones:

- Total power for computations <1W
- Real-time performance
- MCU class processors

Robotic platform

Nano-drone: weight 27g, diameter 10 cm.



- Total power budget 8 W
- STM32F405 MCU runs state estimation and actuation controls
- Flow deck: optical flow and height measurements
- Multi-ranger deck: line of sight distance measurements within [0,4] meters at 20 Hz.

L2 Memory [512 kB]

L1 Memory [64 kB]

Fabric Ctrl (FC)

Cluster (CL)

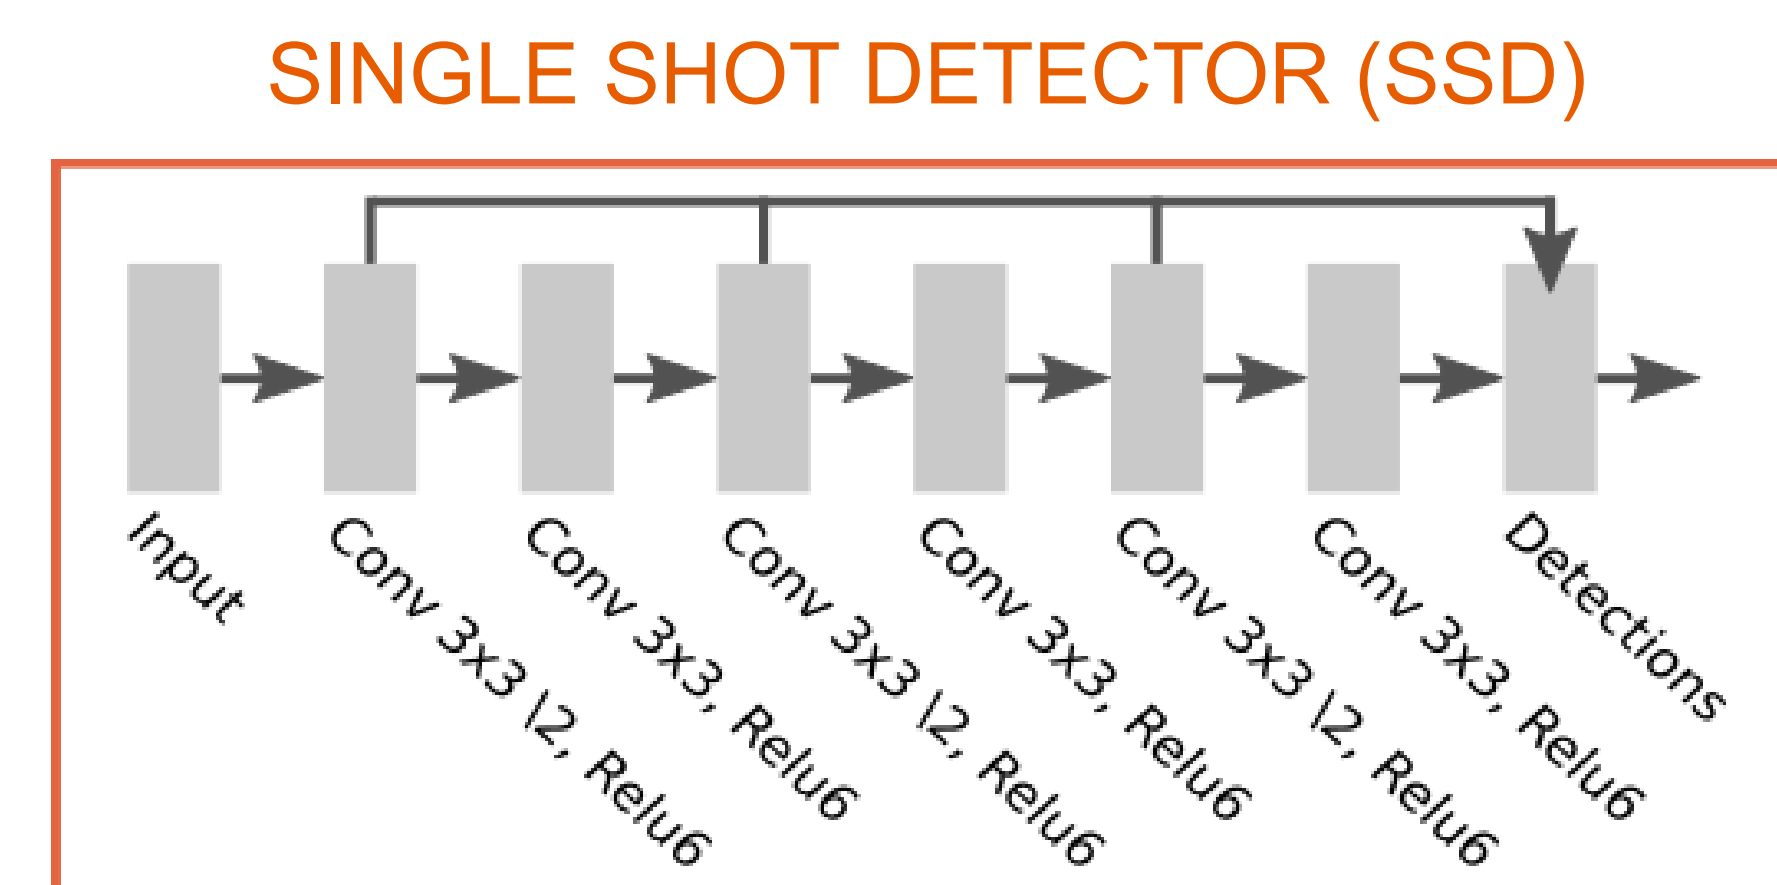
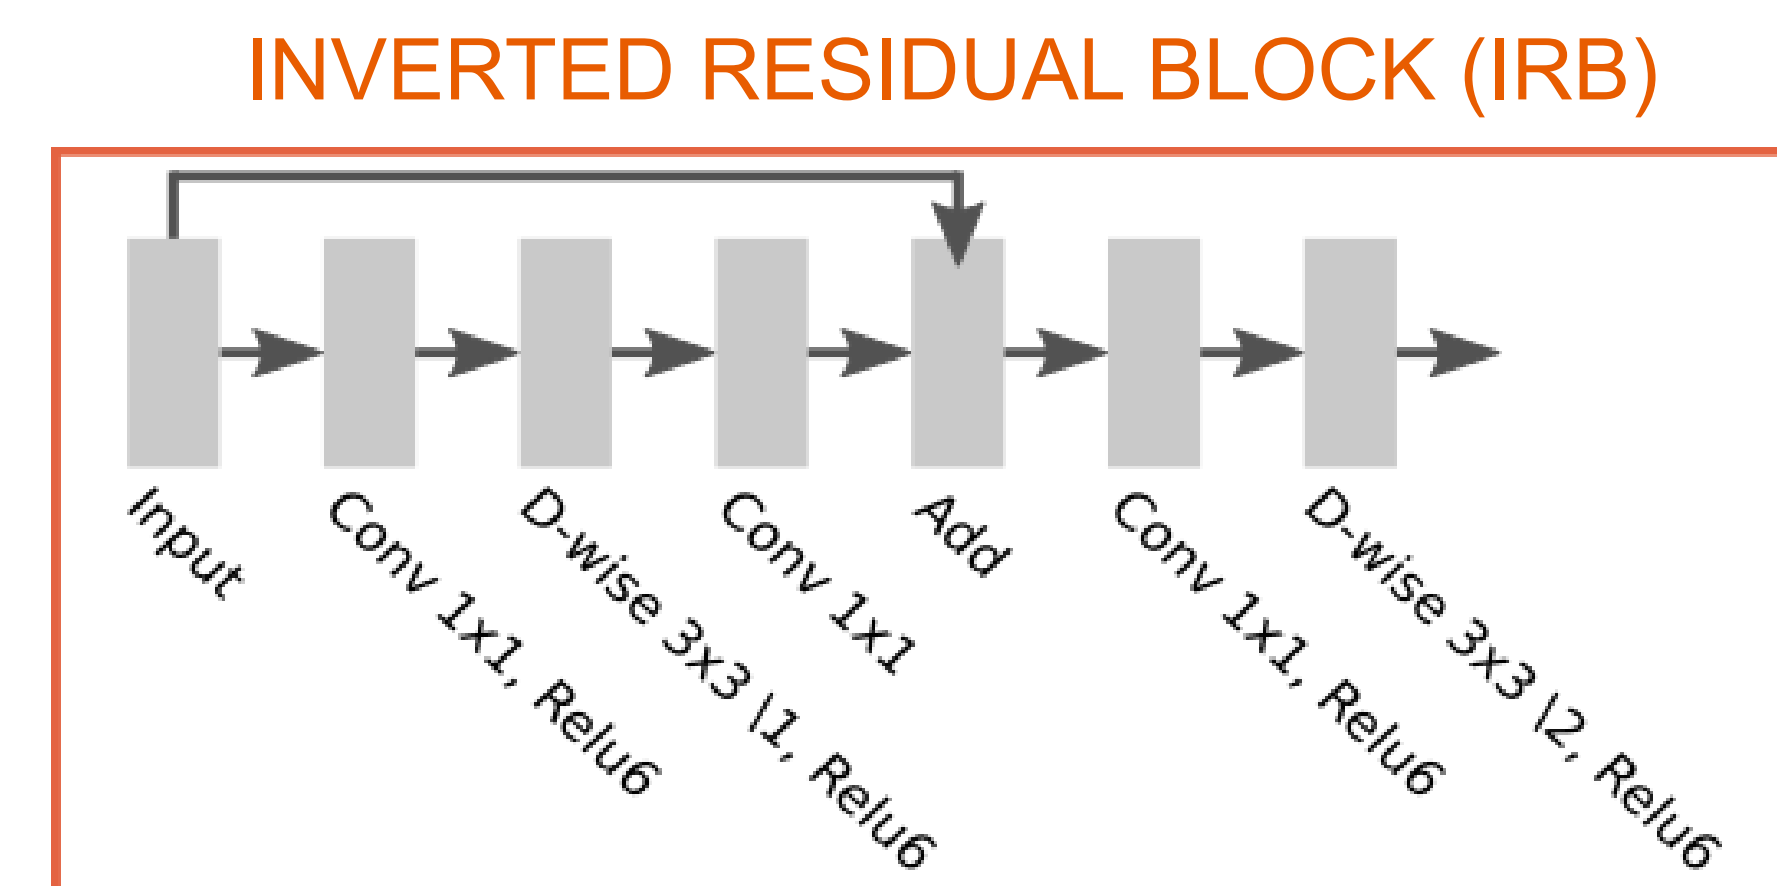
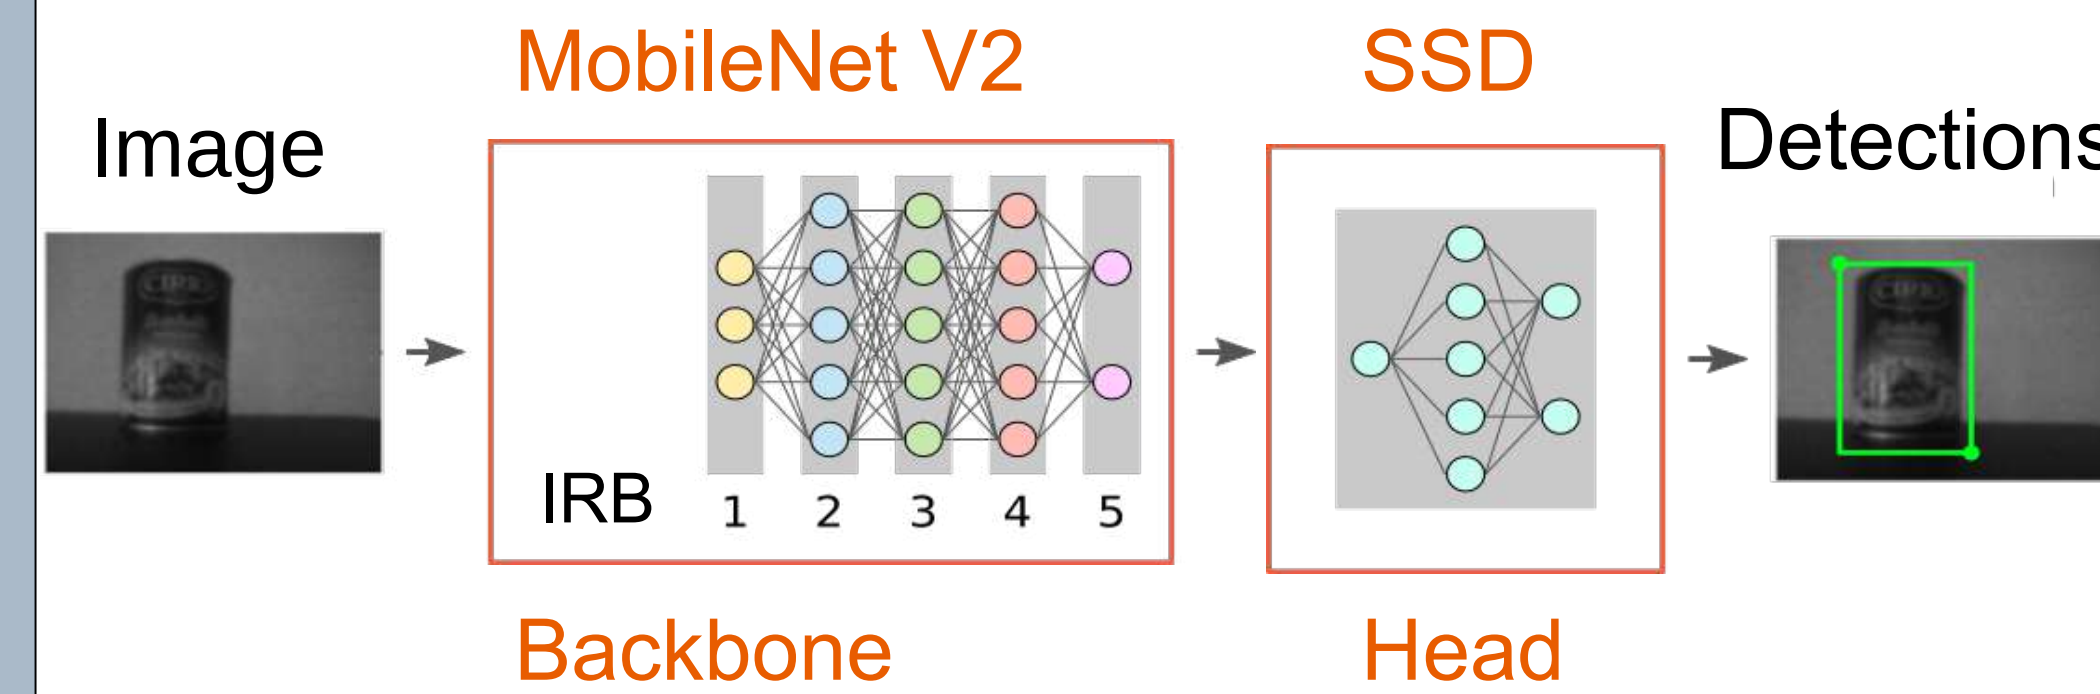
Core 0 Core 1 Core 2 Core 3
Core 4 Core 5 Core 6 Core 7

GAP 8

- Cluster clock 175 MHz
- L3 memory 8 MB
- Fabric Controller 250 MHz
- 320X240 px grayscale camera
- Wi-Fi for data streaming

RISC-V
GREENWAVES TECHNOLOGIES

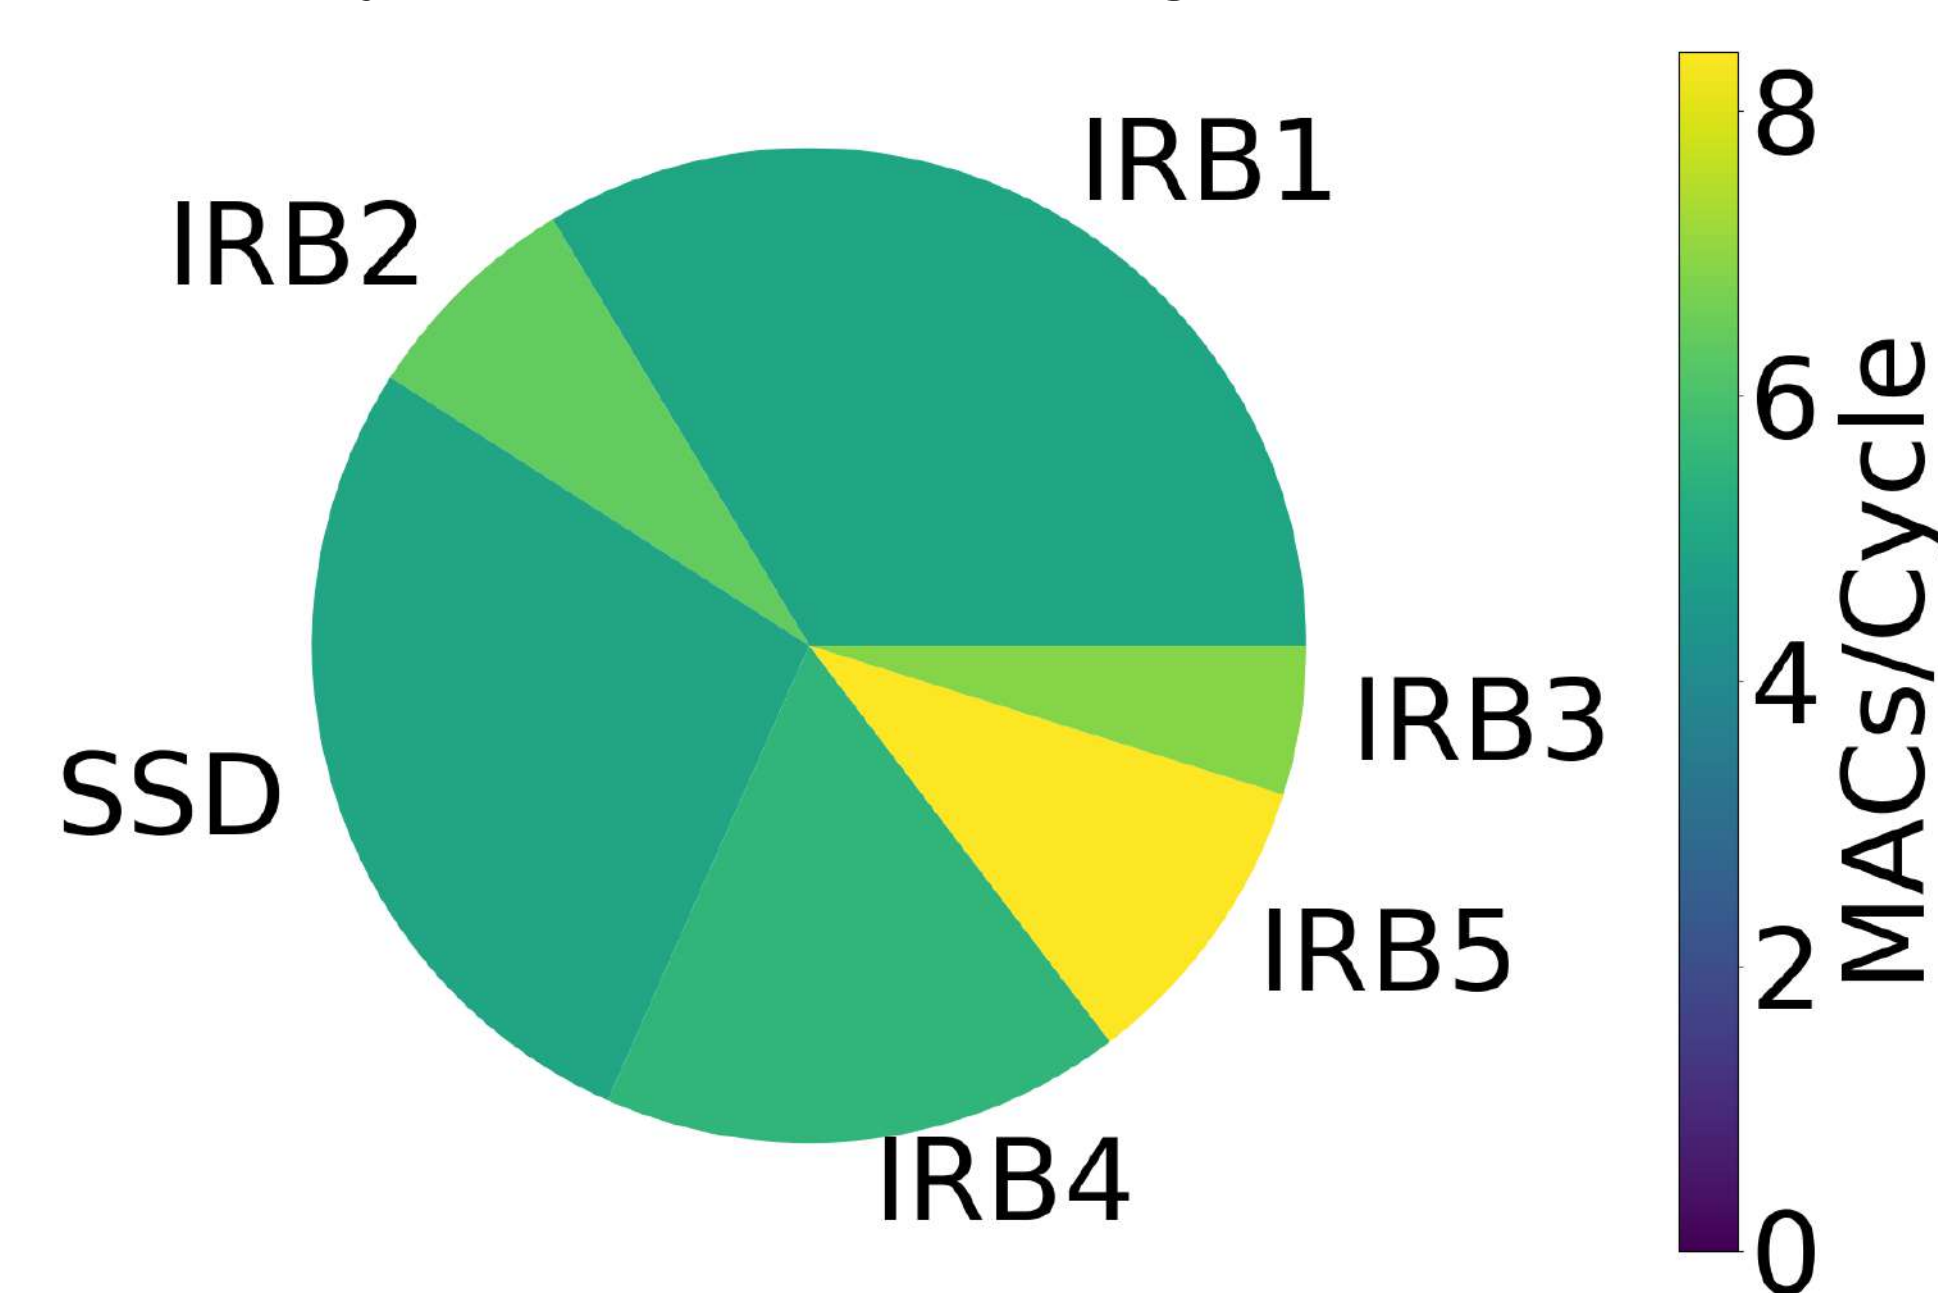
MobileNet V2 SSD



α depth multiplier: multiplies the number of channels of each layer

Depth multiplier	Parameters [M] (SSD)	MAC [M]
$\alpha=1$	4.67 (0.67)	534
$\alpha=0.75$	2.68 (0.55)	358
$\alpha=0.5$	1.34(0.43)	193

Cycles for each building block



Deploying SSD

- Models are trained on OpenImages dataset.
- RMSProp optimizer Learning rate of $8 \cdot 10^{-4}$.
- Images resized to 320x240
- Augmentations randomly applied with 50% probability: flipping, brightness adjustment, random cropping and grayscale conversion



Google OpenImages example image

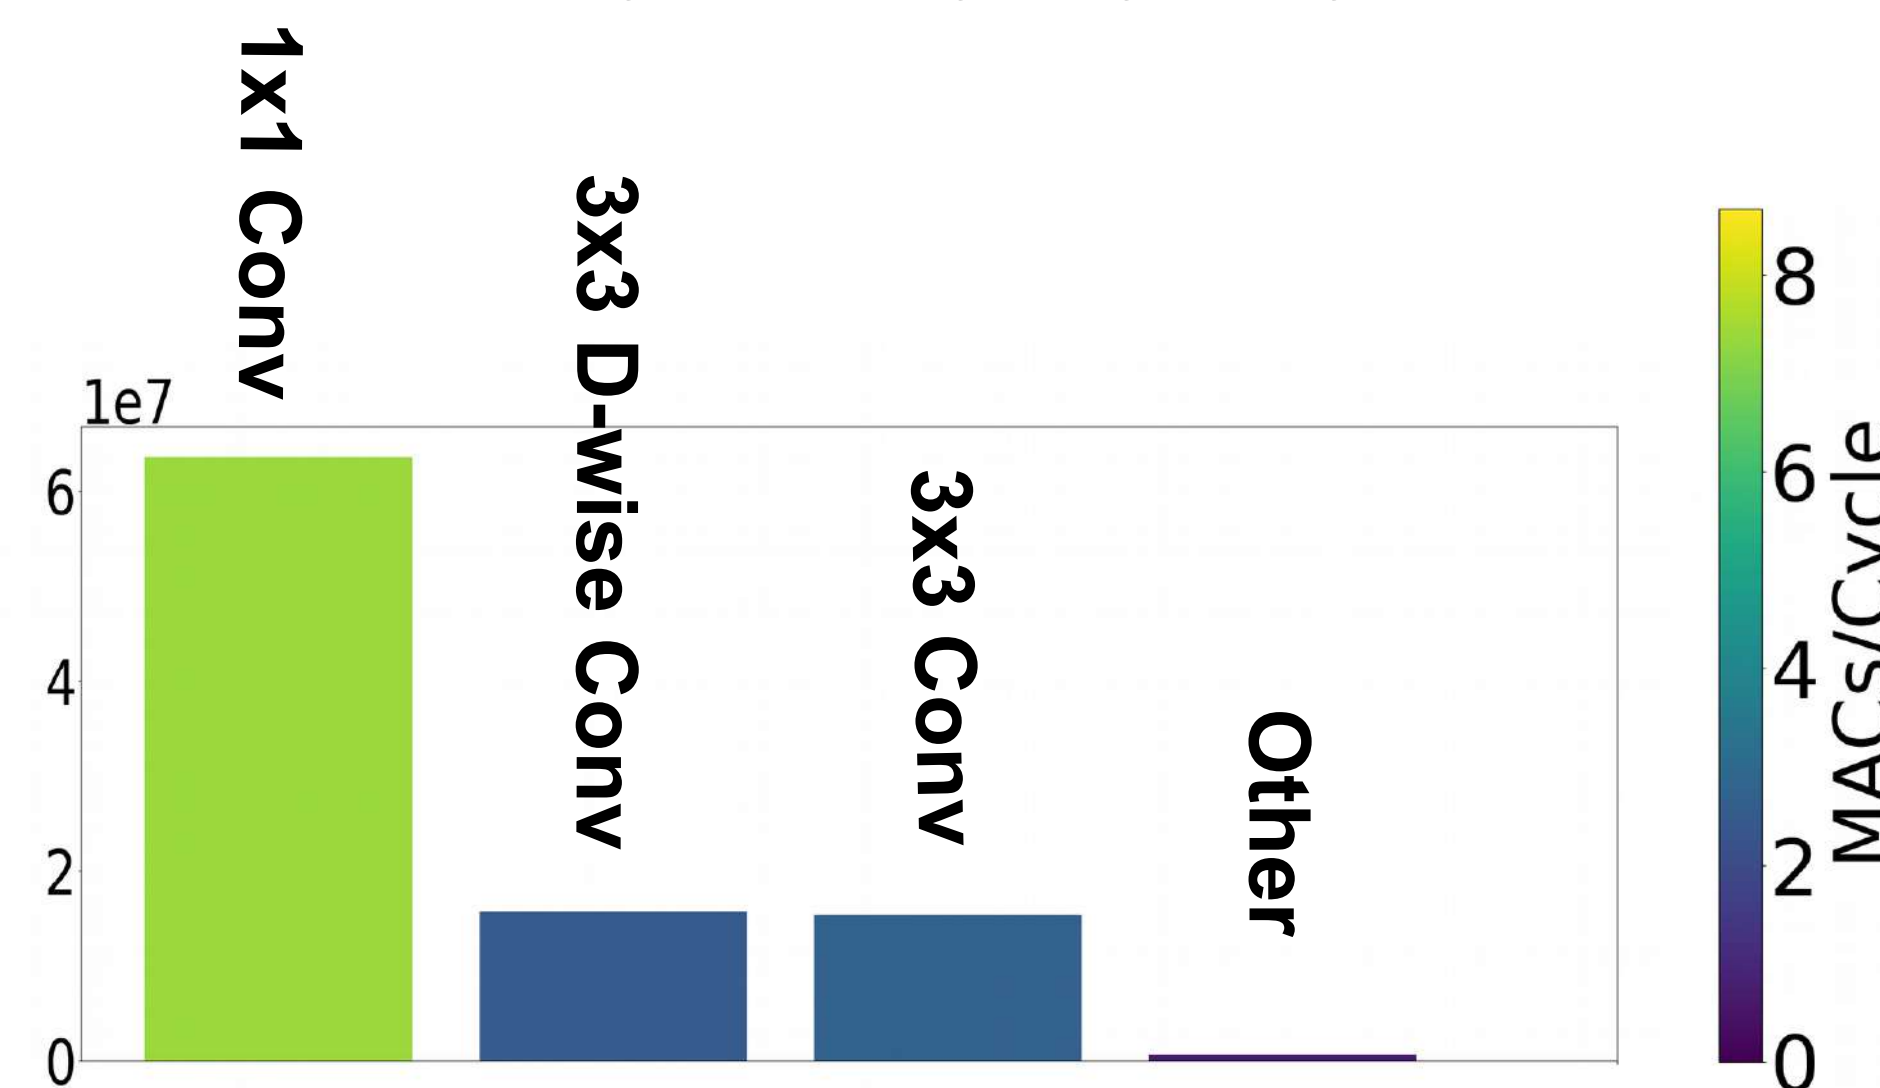
Himax camera image

Finetuning on Himax camera images

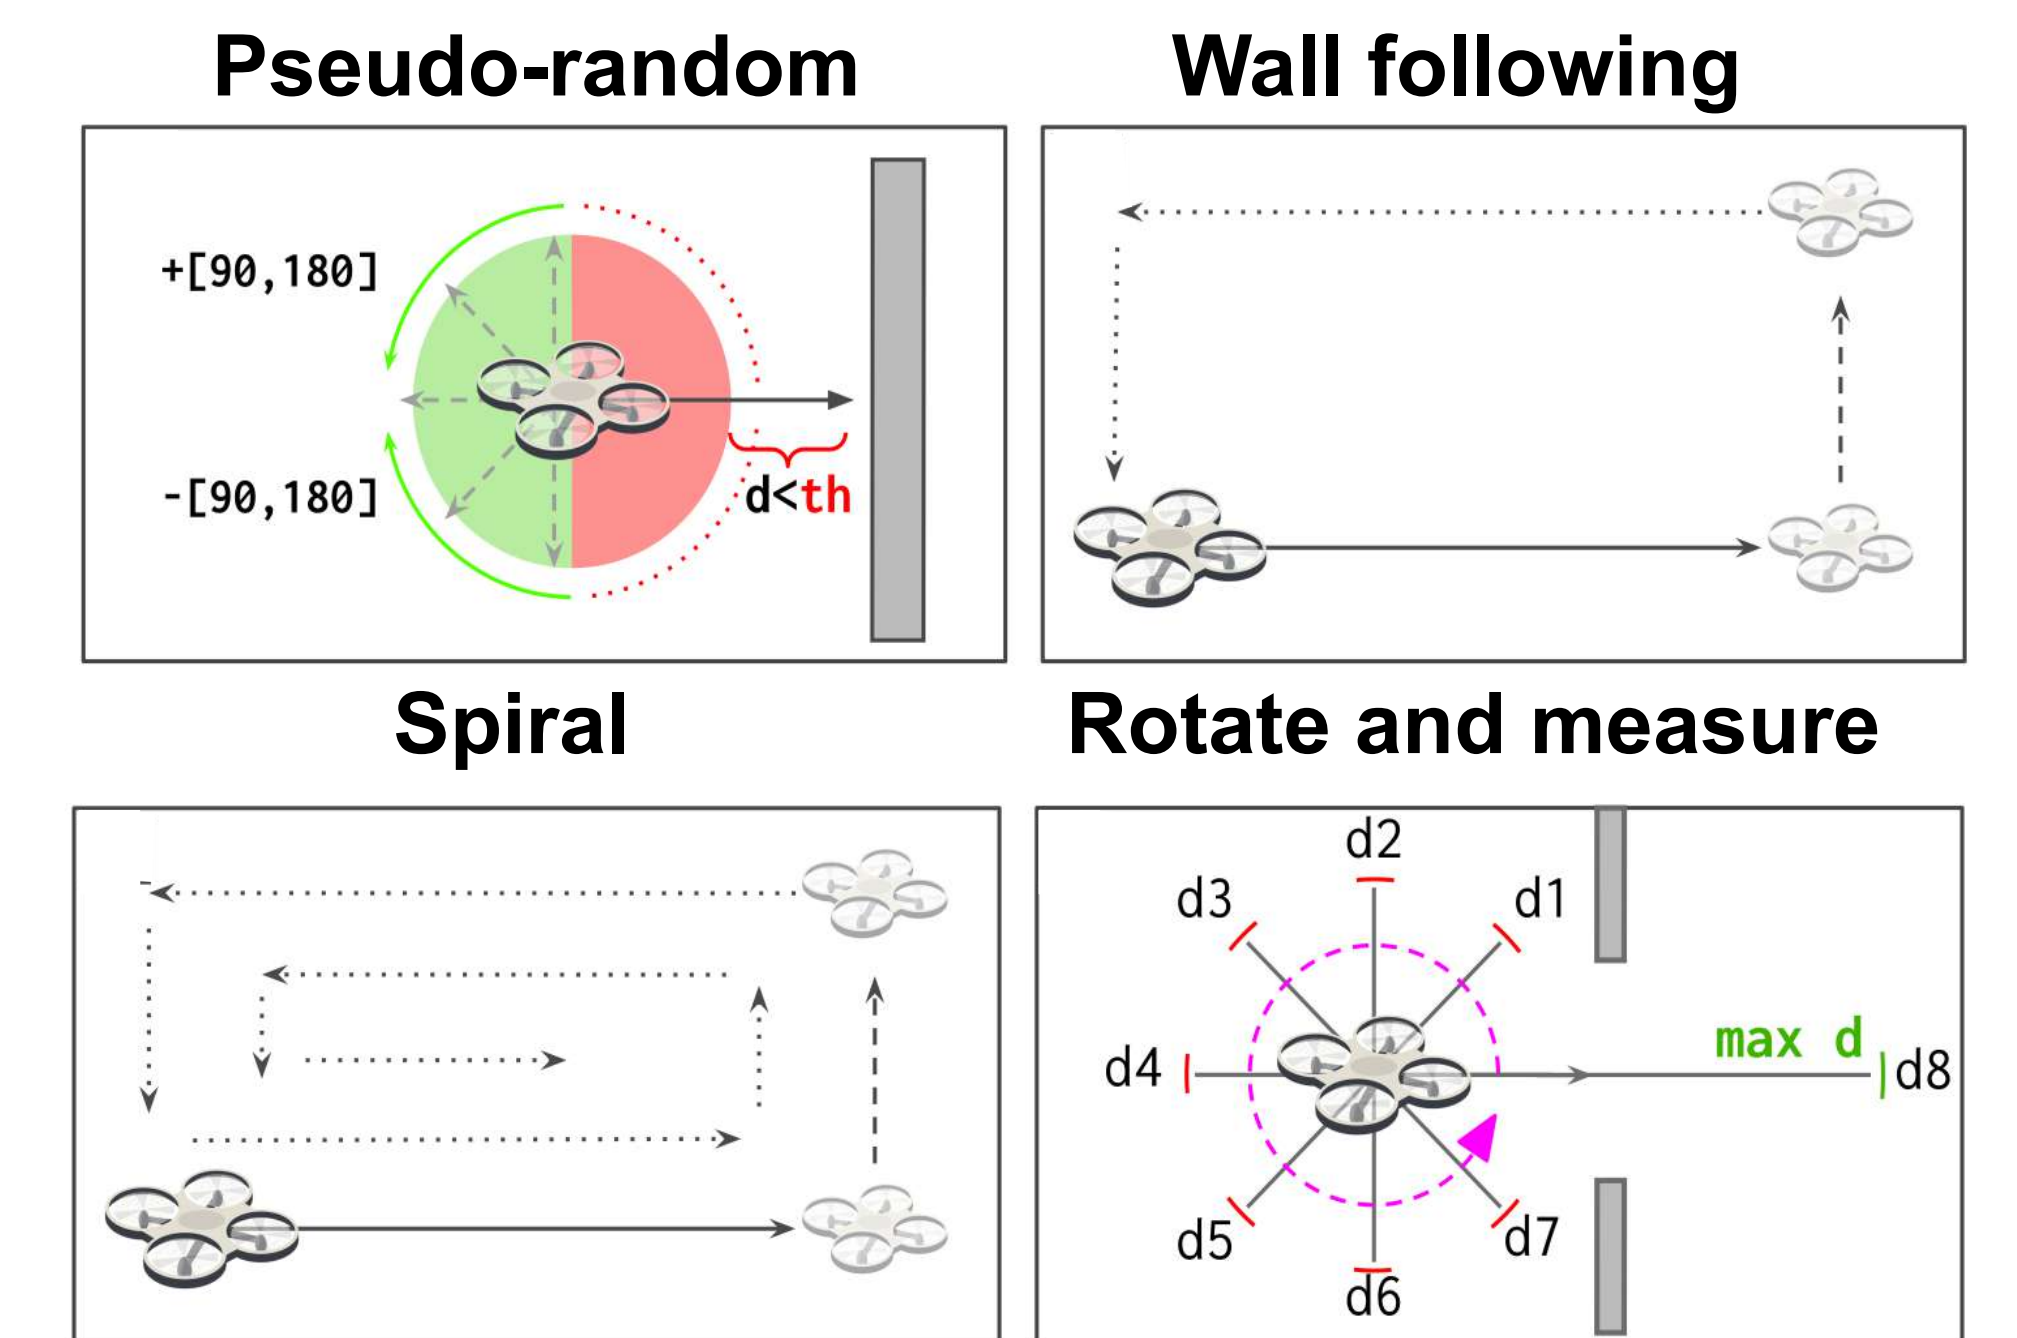
Testing dataset	Fine tuning	Format	SSD size		
			1	0.75	0.5
OpenImages	no	float32	59%	47%	43%
Himax	no	float32	50%	41%	29%
Himax	yes	float32	55%	46%	43%
Himax	yes	int8	50%	48%	32%

SSD	mAP	Throughput	Efficiency
1x	50%	1.6 FPS	5.3 MAC/cycle
0.75x	48%	2.3 FPS	5.9 MAC/cycle
0.5x	32%	4.3 FPS	5.3 MAC/cycle

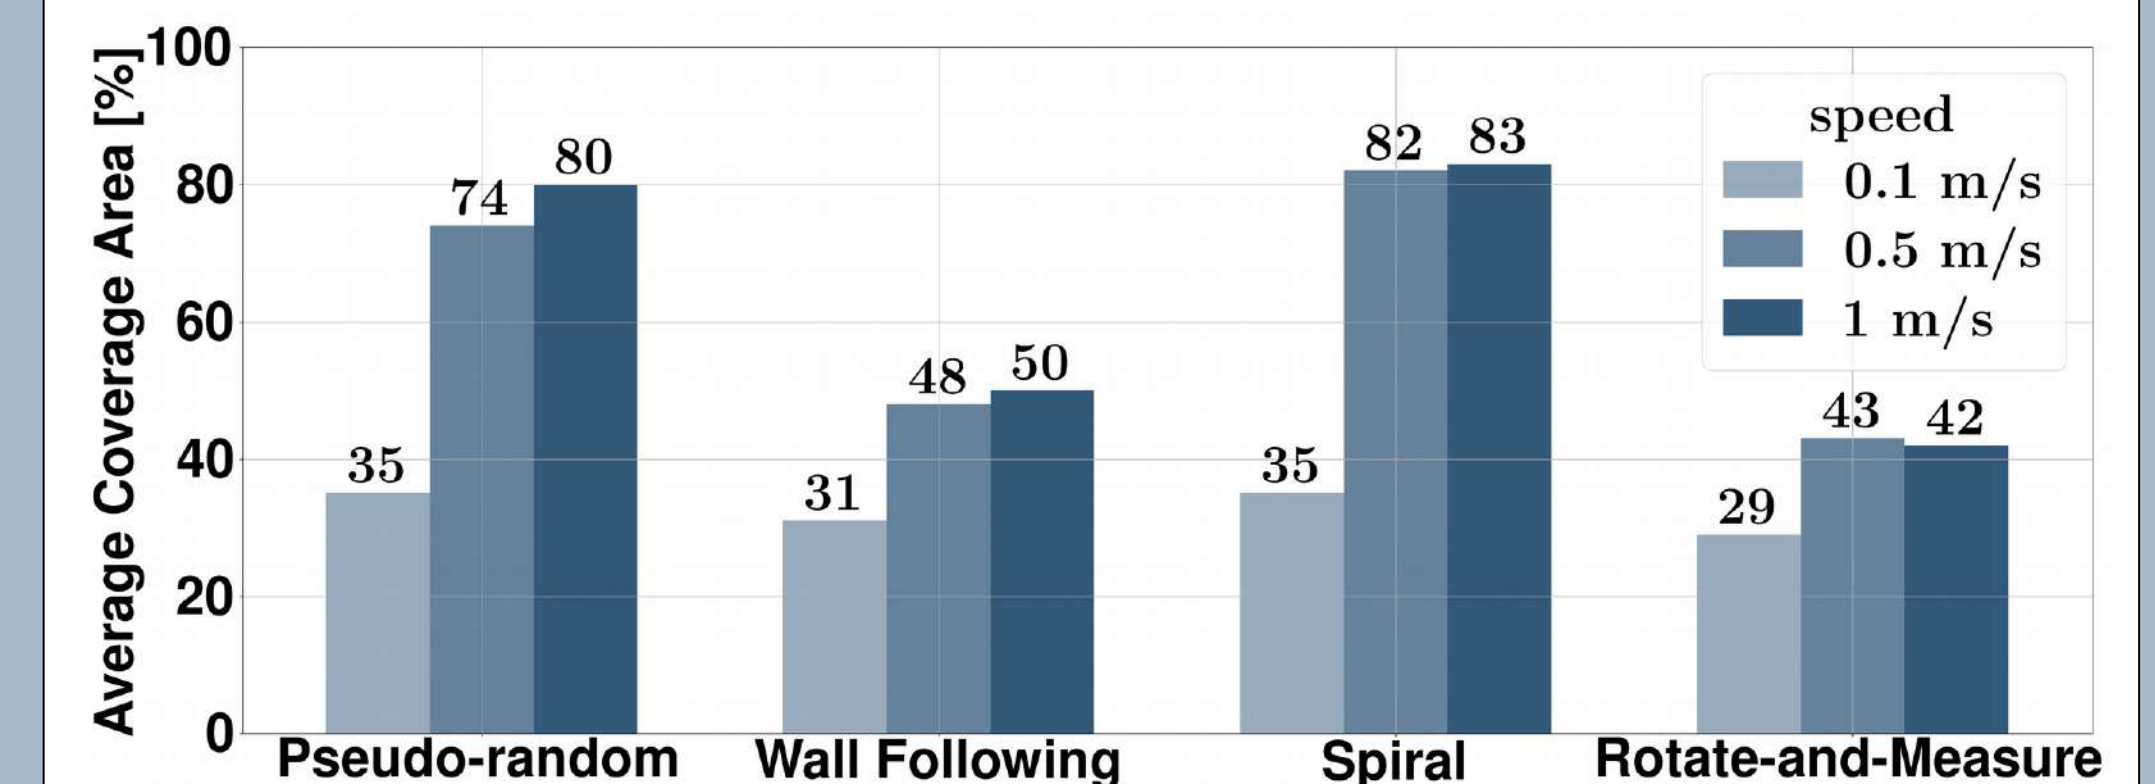
Cycles by layer type



Explorations policies




3 exploration speeds



End-to-end evaluation

- Targets:**
- 3 bottles
 - 3 tin cans

Metric:
Detection rate = $\frac{\#Detected\ Objects}{\#Present\ Objects}$



SSD	Flight speed [m/s]	Detection rate			
		Pseudo random	Wall following	Spiral	Rotate and measure
1.0x	0.1	27%	63%	67%	53%
	0.5	90%	50%	73%	53%
	1	83%	53%	70%	47%

Object detection is possible in real time (1.6FPS) with a detection rate of 90%