Giovanni Maria Capuano
Edge AI for Autonomous Spacecraft: Enabling Onboard Intelligence with FPGA Acceleration

Tutor: Prof. Strollo co-Tutor: Prof. Petra

Cycle: XXXVIII Year: 3

1

Candidate's information

- MSc degree in **Electronic Engineering @DIETI – Federico II**
- DIETI Research group/laboratory: **VLSI Group**
- PhD start date: **01/11/2022** – end date: **31/10/2025**
- Scholarship type: **PNRR – DM 352**
- Partner company: **Techno System Development (TSD-Space)**
- Periods in company: 12 Months
- Abroad Research Institution: **The European Space Research and Technology Centre (ESTEC) of ESA (Noordwijk, Netherlands)**
- Period abroad: 6 Months

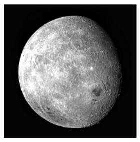

2

Summary of study activities

- AI-based Object Detection (OD) on optical satellite imagery
- AI-based Super-Resolution (SR) on optical satellite imagery
- AI-based Vision-based Satellite Pose Estimation
- AI-based Satellite Jitter Detection for EO imagery quality enhancement
- FPGA-based HW acceleration of Deep Learning algorithms
- FPGA-based Ethernet – Multigigabit COTS I/F

Ad hoc PhD Course

- ✓ Design methodologies for digital circuits and systems oriented to FPGA (Porf. Gennaro Di Meo
- ✓ Innovation and Entrepreneurship (Prof. Pierluigi Rippa)

Conference

- ✓ SciTech, American Institute of Aeronautics and Astronautics, Orlando (USA) (6-10 January 2025) (Oral Presentation)
- ✓ PRIME Conference, 20th International Conference on PhD Research in Microelectronics and Electronics, Taormina, Italy (21 – 24 September 2025) (Oral Presentation)
- ✓ IAC25, International Astronautical Congress, Earth Observation Symposium, Sydney, Australia (29 September – 3 October 2025) (Oral Presentation)

Seminar

Summary of study activities

PhD Year	Courses	Seminars	Research	Tutoring / Supplementary Teaching
1st	30	8,7	35	0
2nd	13	5,3	42	0
3rd	5,4	0,5	60	2

Research area: Edge AI in Space

Space missions are increasingly driving the need for spacecraft to operate with a high degree of **onboard autonomy**. The ability to process raw sensor data at the edge and execute time-critical decisions is essential to support mission scenarios requiring **low-latency responsiveness** and **reduced reliance on ground infrastructure**

AI technologies--including **machine learning (ML)** and **deep learning (DL)**--are increasingly integrated into space missions to strengthen onboard data processing and support autonomous operations

Earth Observation

- EO satellites generate TB of data daily
- Multiple Ground Contact
- Downlink bandwidth limitation
- Not useful data
- Delay between data acquisition and information delivery

Cloud Detection

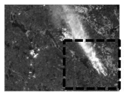

Fire Detection

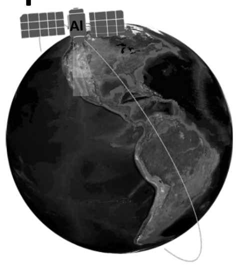

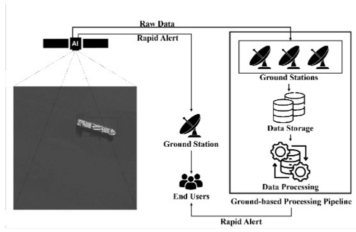

Data Filtering, Data Prioritization, Information Extraction

Giovanni Maria Capuano

5

5

Research area: Edge AI in Space

AI technologies—including **machine learning (ML)** and **deep learning (DL)**—are increasingly integrated into space missions to strengthen onboard data processing and support autonomous operations

DL for In-Space Operation

- ❑ Traditional spacecraft navigation methods involve pre-programmed instructions uploaded to the spacecraft before launch.
  - **Deep space exploration**
  - **Proximity operation**
  - **Rendezvous**
  - **Formation fly**

Require **constant monitoring** and **guidance** from ground control stations
Communication delay or **interruptions**

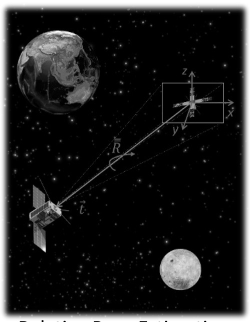

Relative Pose Estimation

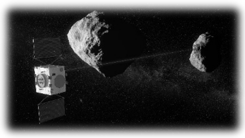

Asteroids Characterization

Autonomy: enable trajectory determination and control without reliance on ground-based systems.

AI for timely decision-making to avoid hazards

Giovanni Maria Capuano

6

6

Problem: Performance Gap

The computational demands of recent DL workloads often exceed the capabilities of traditional radiation-hardened space processors, typically designed using less dense semiconductor technologies to ensure reliability in harsh space environments.

- ✓ Robust against radiation effects
  (Smaller transistor → Less charge stored → Easier to flip by radiation)
- ✓ **Key Limitations:**
  - ➢ Limited performance and lower clock speed
  - ➢ High Cost

Modern DL Model Requirements:
- ✓ Resnet-50: 4.1 billion FLOPs per frame (224x224)
- ✓ **Parallel processing:** Thousands of MAC operations simultaneously

RAD750 → ~ 20 s

Model	Input Size	Params. (M)	FLOPs (G)	Mem. (MB)
VGG-16	$224 \times 224 \times 3$	138.4	15.5	528
ResNet-50	$224 \times 224 \times 3$	25.6	4.1	98
MobileNet V1	$224 \times 224 \times 3$	4.2	0.57	16
YOLOv3	$416 \times 416 \times 3$	61.9	65.9	236
YOLOv5s	$416 \times 416 \times 3$	7.0	7.3	27

CNN Computational Requirements

Processor	Power (W)	Perf. (DMIPS)	Speed (MHz)	Peak GFLOPS	Mem. BW (GB/s)	Core Type	Node (nm)
RAD750	~10	~400	200	~0.4	~0.2	Single-core PowerPC 750	150–250
RAD5545	~20	5,592	up to 466	13.5	4.2	Quad-core QorIQ P5040	45
GR740	~2	~250	250	~0.5	~0.2	Quad-core LEON4	65

Space-grade Processor

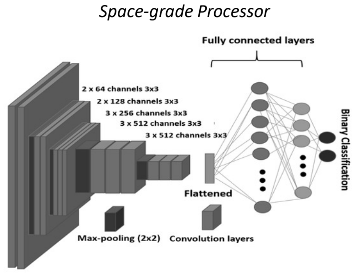

Problem: Space Radiation

The growing computational demands of onboard AI applications, combined with the increasing maturity of COTS Edge AI hardware, have driven the exploration of **heterogeneous processing architectures** that incorporate such **co-processors within the space domain**
- ➢ faster deployment cycles
- ➢ lowering entry barriers for space missions
- ➢ short-duration and LEO missions

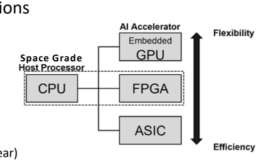

Low Earth Orbit (LEO)
- Altitude: ~200 – 1,200 km
- Radiation: Relatively low (~10-50 krad/year)
- **South Atlantic Anomaly** (200 Km) (low-inclination orbit / 6-7 pass a day)

Medium Earth Orbit (MEO)
- Altitude: ~2,000 – 20,000 km
- Directly in the outer Van Hallen belt
- Radiation: Extremely high (100-1000 krad/year.

Geostationary Orbit (GEO)
- Altitude: ~35,790 km
- Beyond both Van Allen belts, Exposed to solar wind and cosmic rays.
- No magnetic protection
- Radiation: high (~50-100 krad/year)

Device	TID (Krad)	SEL Immunity (LET)
Myriad 2 VPU	49	$8.8\ \mathrm{MeV{\cdot}cm^2/mg}$
RT PolarFire SoC	100	$68.2\ \mathrm{MeV{\cdot}cm^2/mg}$
RT PolarFire FPGA	100	$80\ \mathrm{MeV{\cdot}cm^2/mg}$
Myriad X VPU	20	$49.3\ \mathrm{MeV{\cdot}cm^2/mg}$
Google Coral TPU	30	$57.3\ \mathrm{MeV{\cdot}cm^2/mg}$
NVIDIA Jetson Xavier	20	$40\ \mathrm{MeV{\cdot}cm^2/mg}$

Edge AI Engine - Radiation Tolerance

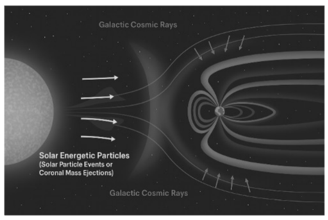

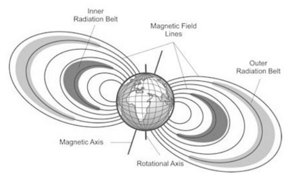

Problem: Space Radiation

Originally developed for terrestrial edge and mobile applications, COTS devices are demonstrating operational reliability in **short-duration missions** and **Low Earth Orbit (LEO)** environments, accumulating initial flight heritage.

Space Mission	Technology	TOPS	Power	AI App	Mission Type
PhiSat-1 (ESA)	Intel Myriad-2 VPU	2	2 W	Cloud Detection	LEO
PhiSat-2	Intel Myriad-2 VPU	2	2 W	Cloud Detection Vessel Detection Image Compression	
CogniSat-2	Intel Myriad X VPU (CogniSat-XE2)	4	2 W	Flood Detection RF Signal Classification	
Forest-2	NVIDIA Jetson Xavier	21	10 W	Wildfire Detection	
NewSat-27	NVIDIA Jetson TX2i	1.33 (TFLOPs)	7 – 15 W	Ship Detection	
SC-LEARN (NASA)	Google TPU	4	2 W	TBD	

Research area: FPGAs for AI in Space

Field-Programmable Gate Arrays (FPGAs), in contrast, have been adopted in space applications for several decades. **Radiation-hardened /tolerant FPGAs** are now commercially available and can withstand the harsh space environment even in demanding missions, such as **long-duration deployments in GEO and beyond.**

Device	Technology	TID (Krad)	SEL Immunity (LET)
Intel Myriad-2 VPU	VPU	49	8.8 $MeV \cdot cm^2/mg$
Intel Myriad X VPU	VPU	20	49.3 $MeV \cdot cm^2/mg$
Google Coral TPU	TPU	30	57.3 $MeV \cdot cm^2/mg$
NVIDIA Jetson Xavier	GPU	20	40 $MeV \cdot cm^2/mg$
AMD Xilinx Kintex UltraScale	**FPGA**	100	80 $MeV \cdot cm^2/mg$
AMD Xilinx Virtex-4QV	**FPGA**	300	125 $MeV \cdot cm^2/mg$
Microchip RT PolarFire	**FPGA**	100	68.2 $MeV \cdot cm^2/mg$
Microchip RT PolarFire	**SoC (FPGA + RISC-V)**	100	80 $MeV \cdot cm^2/mg$

ECSS Latch-up immunity: LET threshold > **68.2** MeV·cm2 /mg

FPGA Acceleration for Satellite Detection

ANHEO: A highly **integrated unit for autonomous absolute** and **relative navigation** of nano- and microsatellites, from **low Earth orbit** and **high Earth orbits**, up to **lunar altitude**.

- 6U-12U CubeSats
- Mass <0.8 Kg
- Power < 10W
- 0.8 U form factor

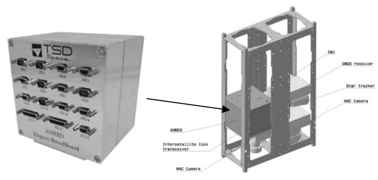

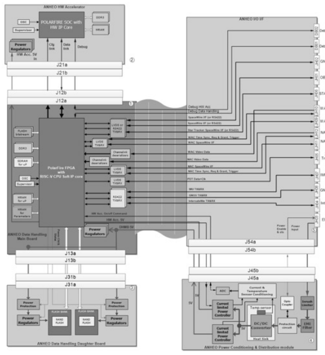

✓ **CoreVectorBlox in RT PolarFire SoC**

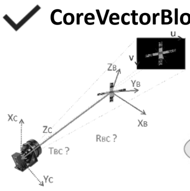

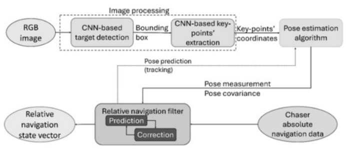

Agenzia Spaziale Italiana

iteePhD information technology electrical engineering

Giovanni Maria Capuano

11

11

FPGA Acceleration for Satellite Detection

Microchip CoreVBX: The HA leverages the advanced capabilities of Microchip's PolarFire SoC that combines a powerful **64-bit 5x core RISC-V** Microprocessor Sub-System (MSS), with the **PolarFire FPGA fabric** in a single device. The board enables the real-time execution of algorithms for autonomous absolute and relative navigation.

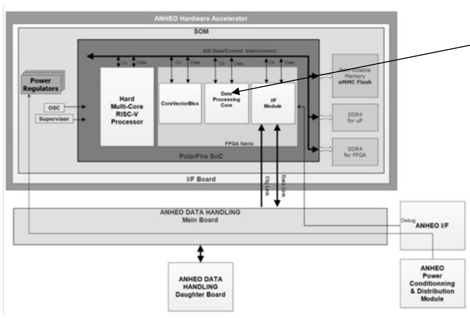

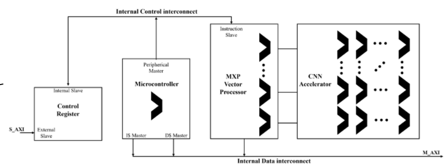

- Overlay architecture
  - ✓ Weights and biases are stored in off-chip non-volatile memory
  - ✓ A main controller coordinates the activities of the PEs, managing data flow across the buffers and executing layer-specific instructions to optimize inference efficiency.
  - ✓ Various operations multiplexed across layers

Agenzia Spaziale Italiana

iteePhD information technology electrical engineering

Giovanni Maria Capuano

12

12

FPGA Acceleration for Satellite Detection

The YOLOv5su algorithm for **Satellite Detection** cannot be directly accelerated using the AI engine due to layers and activation functions that are not supported by the PolarFire Core VBX architecture.

- SiLU Activation Function
- Concatenation Layer (Axis 2)

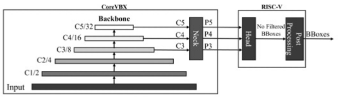

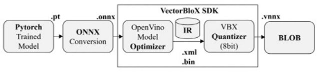

Model	Input Size (pixels)	mAP^{val} 50-95	mAP^{val} 50	Inference Speed (ms) CPU	Inference Speed (ms) A100	Params (M)	FLOPs (B)
YOLOv5s	640	37.4	56.8	98	$0.9^{\dagger}$	7.2	16.5
YOLOv5su	640	43.0	–	$120.7^{\ddagger}$	$1.27^{\S}$	9.1	24.0

Device	Precision	mAP50	Inference Speed	Avg. Power Consumption	Avg. Energy Consumption
Intel Core i9 CPU	FP32	0.995	~56 ms	~125 W	~7 $J/frame$
NVIDIA A100 GPU	FP32	0.995	~10 ms	~53 W	~530 $mJ/frame$
NVIDIA Jetson Orin Nano Embedded GPU	FP32	0.995	~35 ms	~9.6 W	~336 $mJ/frame$
PolarFire SoC (ANHEO HA)	INT-8	0.965	~170 ms	~2.3 W	~311 $mJ/frame$

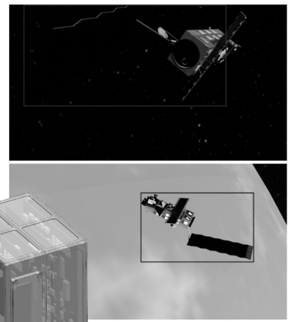

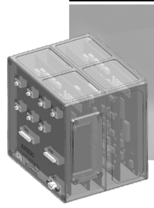

itee PhD information technology electrical engineering

Giovanni Maria Capuano

13

13

Main Benefits of FPGAs for AI in Space

Concurrent deployment of multiple VHDL modules,

- data management,
- Storage
- high-throughput sensor interfaces (SpaceWire, SpaceFibre)

Firmware Update (RT PolarFire FPGA supports up to 500 in-orbit reprogramming cycles)

ANHEO FPGA firmware is update to synthetize the newest release of CoreVectorBlox

Config.	Vector Proc. Width	Vector Scratch.	CNN Accel. Array Size
V250	64-bit	64 kB	16×8
V500	128-bit	128 kB	16×16
V1000	256-bit	256 kB	32×16

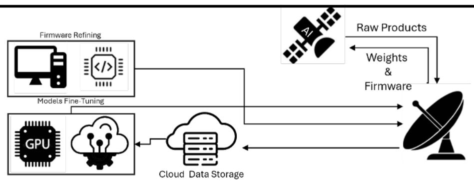

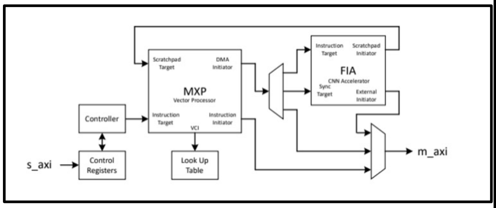

CoreVBX v2

Model	Accel. Core	CNN Speed	Power Consumption
YOLOv5s – ReLU	CoreVBX 1	~59.56 ms	1.16 W
YOLOv5s – ReLU	CoreVBX 2	~55.00 ms	1.32 W

Model	Accel. Core	CNN Speed	Post-Proc. Type	RISC-V Latency
Yolov5n – ReLU	CoreVBX 1	~20.9 ms	Anchor-based	802.7 ms
Yolov5n – SiLU	CoreVBX 2	~51.6 ms	Anchor-free	1908.6 ms
Yolov5nu – SiLU	CoreVBX 2	~126.78 ms	Anchor-free	629.6 ms

Device	Inference Speed	Avg. Dynamic Power	Avg. Dynamic Energy
CoreVectorBlox v1 + RISC-V	~791.5 ms	~0.58 W	~459 mJ/fr.
CoreVectorBlox v2	~240.5 ms	~0.96 W	~230.8 mJ/fr.

itee PhD information technology electrical engineering

Giovanni Maria Capuano

14

14

Research area: EO Imagery Quality Enhancement

- Spacecraft jitter refers to **small-amplitude, high-frequency variations in a spacecraft's line-of-sight (LoS)**
- A small angular disturbance in spacecraft's pointing induces a displacements of the scene's projection onto focal plane
- The jitter-induced motion results in **misregistration** between successive nitration intervals, **generating blur** and **geometric distortion**

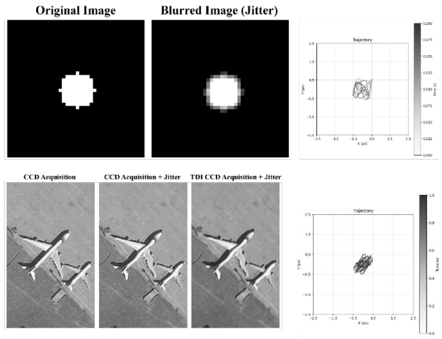

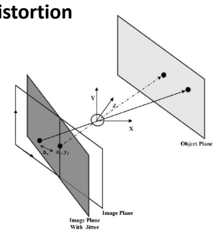

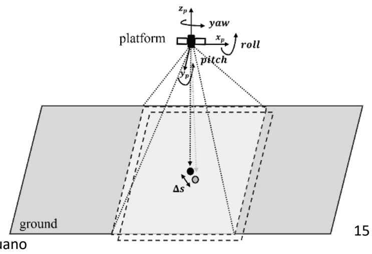

itee PhD information technology electrical engineering

Giovanni Maria Capuano

15

FOPAC System

Agenzia Spaziale Italiana

TSD Space

The objective of **FOPAC** is to design and validate an **active motion-control system** for the focal plane of **electro-optical payloads dedicated to Earth observation**, with the goal of enhancing image quality in terms of both spatial resolution and SNR.

FOPAC introduces fine **opto-mechatronic stabilization** directly at the **focal-plane level**.

- An actively controlled, piezo-actuated focal plane for the suppression of attitude-jitter components.
- By appropriately driving the piezoelectric nanopositioner, micro-displacements can be applied during image acquisition to compensate high-frequency sinusoidal attitude disturbances and stabilize the line of sight.

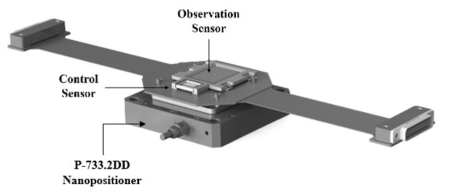

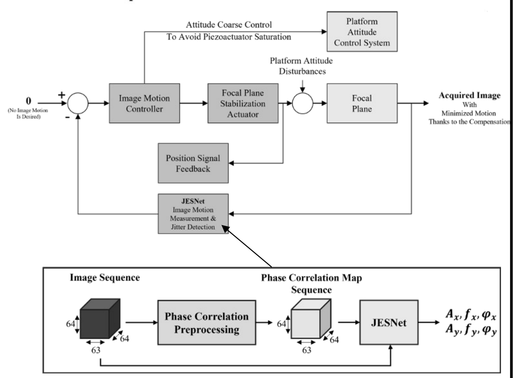

itee PhD information technology electrical engineering

Giovanni Maria Capuano

16

FOPAC System

JESNet is a **fully convolutional neural network** (CNN) specifically designed to process **temporal sequences** of images to simultaneously **infer six parameters** of **sinusoidal jitter.**

JESNet combines **spatial-domain** and **frequency-domain** information in a unified learning framework.

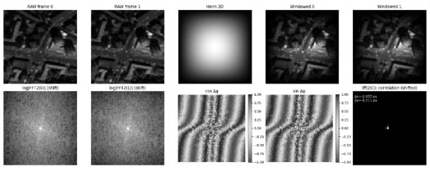

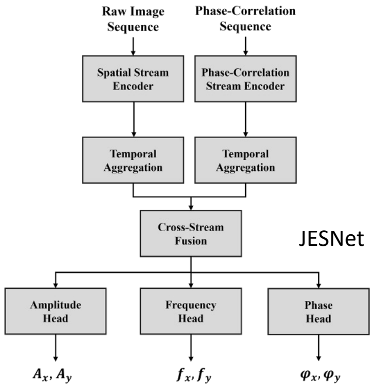

JESNet

Algorithm	MAE Amplitude	MAE Frequency	MAE Phase
JESNet	0.327 pixels	1.64 Hz	0.52 rad
Baseline	1.27 pixels	4.6 Hz	0.8 rad

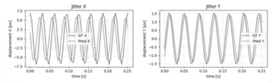

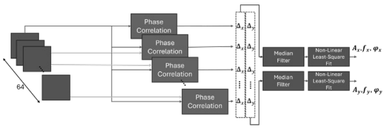

Model-based Baseline

Giovanni Maria Capuano

Research products

[P1]	G.M. Capuano, S. Capuozzo, A.G.M. Strollo, N. Petra, Super-Resolution YOLO Object Detection for Maritime Surveillance with Real-Time FPGA Processing Onboard Spacecraft, **International Journal of Acta Astronautica, Elsevier** vol. 238, Part A, pp. 1291-1310, 2026, DOI: 10.1016/j.actaastro.2025.09.018
[P2]	T.A. La Marca, M.D. Graziano, M. Grassi, L. Iannascoli, L. Cavicchiolo, A. Capanni, M. Duzzi, G.M. Capuano, V. Fortunato, C. Cardenio, G. Leccese, M. Rinaldi, M. Di Clemente, R. Luciani, R. Natalucci, V. Pulcino, Mission and system design for the EarthNext Cubesat VLEO mission, **CEAS Space Journal, Springer** vol. 238, pp. 1-17, 2025, DOI: 10.1007/s12567-025-00663-2
[P3]	G. M. Capuano, F. Saggiomo, A. G. M. Strollo, N. Petra, B. Confuorto, E. Zaccagnino & M. Di Clemente, Enhancing Satellite Imagery through Advanced Focal Plane In-flight Active Motion Control: The FOPAC System Implementation and Results **International Astronautical Congress, Earth Observation Symposium,** Sydney, Australia, October 2025, In Proceedings of the International Astronautical Congress, IAC-25-B1,IPB,17,x96978
[P4]	G.M. Capuano, S. Capuozzo, A.G.M. Strollo, N. Petra, Benefits of FPGAs for Deep Learning Acceleration, ***PRIME Conference, 20th International Conference on PhD Research in Microelectronics and Electronics,*** Taormina, Italy, September 2025, In Proceedings of the International Astronautical Congress, IAC-24-B1.IP.92.X89480
[P6]	G. M. Capuano, V. Capuano, G. Napolano, R. Opromolla, M. Severi, A. G. M. Strollo, N. Petra, G. Cuciniello, E. Zaccagnino & G. Varacalli FPGA Hardware Acceleration for Deep Learning-Based Satellite Relative Pose Estimation, **American Institute of Aeronautics and Astronautics SciTech Forum on aerospace research, development, and technology**, Orlando, USA, January 2025, In Proceedings of the AIAA SCITECH 2025 Forum (p. 2715)

Name Surname

Research products

[P7]	G.M. Capuano, S. Capuozzo, A.G.M. Strollo, N. Petra, FPGA-Based Hardware Acceleration for Real-Time Maritime Surveillance and Monitoring Onboard Spacecraft, **Spaice 2024: Conference on AI in and for Space** European Centre for Space Applications and Telecommunications (ECSAT), Didcot, United Kingdom, October 2024, In Proceedings of SPAICE2024: The First Joint European Space Agency/IAA Conference on AI in and for Space (pp. 156-161)
[P8]	G.M. Capuano, S. Capuozzo, A.G.M. Strollo, N. Petra, Super-Resolution-Based Small Object Detection for Real-Time Surveillance and Monitoring: An Onboard Satellite FPGA Implementation, **International Astronautical Congress, Earth Observation Symposium,** Milan, Italy, October 2024, In Proceedings of the International Astronautical Congress, IAC-24-B1.IP.92.X89480
[P9]	G.M. Capuano, A.G.M. Strollo, N. Petra, Super Resolution CNN for a quincunx sampling-based panchromatic earth observation imager for nanosatellites, **International Astronautical Congress 2023** Baku, Azerbaijan, October 2023, In Proceedings of the International Astronautical Congress, IAC-23,E2,4,5,x79398

19

PhD thesis overview

- Problem: Reduce reliance on ground-based processing in space missions
  - EO satellites generate TB of data daily
  - Multiple Ground Contact
  - Limited Downlink Bandwidth
  - Not Useful Data
  - Latency between image acquisition and information extraction & delivery
- Objective: Enable AI Intelligence Onboard Spacecraft
  - Rapid Alert
  - Reliable Information Extraction
  - Data Prioritization

20

PhD thesis

- Integrate *deep-learning super-resolution* with *object detection* to enable **reliable, real-time target recognition onboard satellites**, reducing dependence on ground processing and communication bandwidth.
- First framework combining **SR-based enhancement and YOLO detection** optimized for **FPGA deployment**, with a **custom maritime dataset** including synthetic low-visibility and cloud-covered scenes.
- **Portability toward space-grade RT PolarFire devices**, enabling in-orbit deployment and technology transfer to flight-qualified payloads.
- Validate through a rigorous experimental pipeline:
  - Multi-Source Dataset
  - Five-Fold Cross-Validation
  - Object-Size-based Accuracy Metrics
  - **Hardware-aware model design** tailored for **CoreVectorBlox architecture**, ensuring **full accelerability** on embedded AI cores;
  - Hardware performance test on PolarFire SoC (benchmark with NVIDIA Jetson Orin Nano GPU)

21

Research area: Maritime Surveillance

Accurate vessel detection and **timely information extraction** are essential for a wide range of **maritime surveillance and monitoring** operations, both civilian and defense-related:

- Vessel Tracking
- Unauthorized Fishing
- Detecting Oil spills
- Illegal Migration
- Search and Rescue Missions

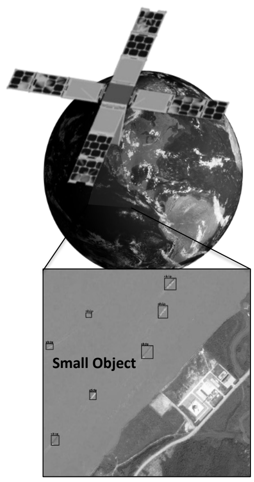

Object Detection enables the **identification** and **localization** of objects within an image by recognizing them from various classes and determining their boundaries with bounding boxes.

Small Vessel Detection Problem:

- **Small Size**
- **Limited Spatial Resolution**
- **Minimal Object Pixelation**
- **Low SNR**
- **No Adequate Structure Information**
- **Complex Background**

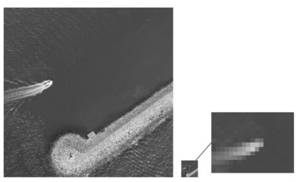

22

SR-Backbone for Object Detection

The **SR backbone** novel CNN architecture specifically developed for **real-time SISR applications.** The model is optimized as an **integral component of object detection networks**, featuring a **lightweight design** that **minimizes memory footprint** and **computational overhead**

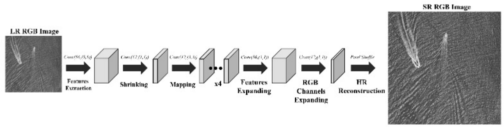

- **Feature Extraction:** Directly from the LR input using a series of convolutional layers
- **Shrinking:** Feature reduction through a 1x1 CONV layer to lower computational complexity
- **Non-linear Mapping:** To capture deeper image representations.
- **Expansion:** The shrunk features are expanded back to a higher dimension using another 1x1 convolution layer
- **RGB Channel Regression & Pixel-Shuffle:** Maps aggregation and SR image reconstruction

Model	Parameters	PSNR (dB)	SSIM
SRCNN	20,099	31.668	0.880
FSRCNN	24,683	31.978	0.889
EDSR	1,395,788	32.790	0.898
SR-Backbone	**16,960**	**30.640**	**0.884**

Result in ShipRSImage

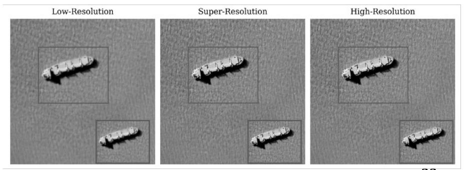

itee PhD information technology electrical engineering

Giovanni Maria Capuano

23

23

SR-YOLOv5 for Small Vessel Detection

SR-YOLOv5s is an enhanced version of the YOLOv5 object detector that incorporates the **SR-Backbone** for the performance improvement of small vessel detection

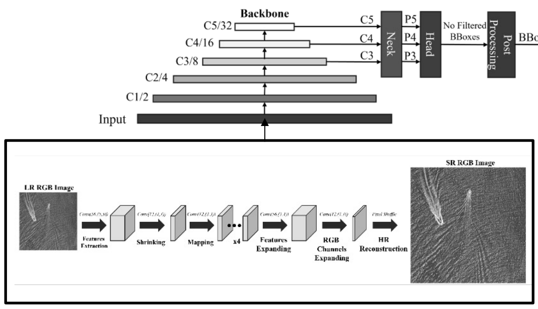

LR-ShipRSImage (3,435 RGB):
- 80% Training
- 20% Validation
- 10% Test

LR Dataset

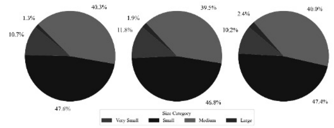

HR Dataset

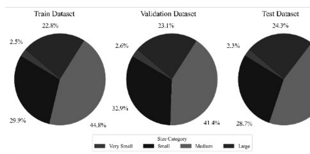

Model	mAP50 All	mAP50:95 All	mAP50 Small	mAP50 Very Small
YOLOv5s	0.8544 ± 0.0054	0.6610 ± 0.0046	0.7972 ± 0.0129	0.2832 ± 0.0615
SR-YOLOv5s	$\mathbf{0.9272 \pm 0.0029}$	$\mathbf{0.7604 \pm 0.0072}$	$\mathbf{0.9272 \pm 0.0079}$	$\mathbf{0.4658 \pm 0.0221}$

Result on LR-ShipRSImage

itee PhD information technology electrical engineering

Giovanni Maria Capuano

24

SR-YOLOv5:Experimental Validation

Multi-source maritime dataset:
- ShipRSImageNet
- MASATI
- synthetic cloudy scenes

5-fold cross-validation

Statistical Significance test on SR-YOLOvs Vs. YOLOv5s Baseline

Analysis of false positives under low-visibility conditions

SR-YOLOv5s significantly outperforms the baseline YOLOv5s,

mAP50 of **93.27% ± 0.80%,** compared to **84.86% ± 1.14%** (mean difference = +8.40% ± 1.32%, $p < 0.001$)

5-fold cross-validation

Model	mAP50 All	mAP50:95 All	mAP50 Small	mAP50 Very Small
YOLOv5s	0.8544 ± 0.0054	0.6610 ± 0.0046	0.7972 ± 0.0129	0.2832 ± 0.0615
SR-YOLOv5s	$\mathbf{0.9272 \pm 0.0029}$	$\mathbf{0.7604 \pm 0.0072}$	$\mathbf{0.9272 \pm 0.0079}$	$\mathbf{0.4658 \pm 0.0221}$

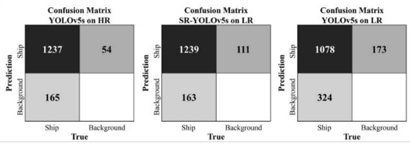

Vessel detection accuracy.

Model	Input size	Param.	mAP50	mAP50:95
YOLOv5s Original	208	7,012,822	0.854	0.661
YOLOv5s Original	416	7,012,822	0.935	0.792
YOLOv5s ReLU	416	7,012,822	0.920	0.761
SR-YOLOv5s	**208**	**7,029,782**	**0.927**	**0.764**

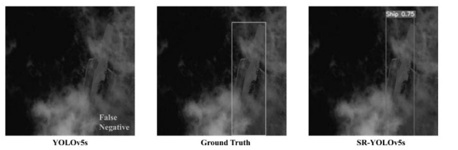

itee PhD information technology electrical engineering

Giovanni Maria Capuano

25

25

SR-YOLOv5:Experimental Results

Model mAP50 results on other datasets.

Model	LR-CloudyTest	MASATI-ship
YOLOv5s	0.610	0.529
SR-YOLOv5s	**0.660**	**0.751**

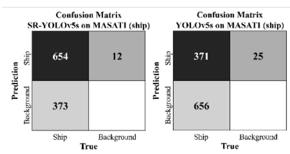

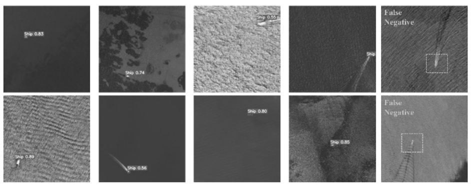

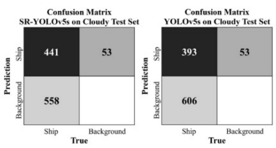

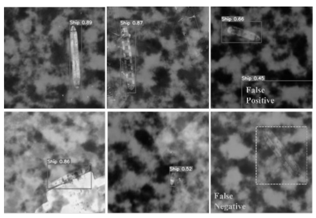

itee PhD information technology electrical engineering

Giovanni Maria Capuano

26

26

FPGA Acceleration for Rapid Alert

PolarFire SoC FPGA (Available in RT Version) + **CoreVectorBlox IP Core**

- Five Core Linux Capable Hard RISC-V
- Non-Volatile Flash-based FPGA
- SEU Immune

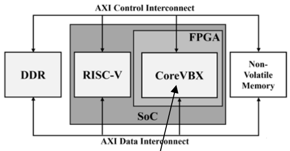

Device	Prec.	mAP50	Inference speed	Avg. power consumption	Avg. energy per frame	mAP50/W
A100	FP32	**0.927**	7.30 ms	53.04 W	384.2 mJ/frame	0.017
Jetson Orin Nano (15 W)	FP32	**0.927**	20.80 ms	10.50 W	218.4 mJ/frame	0.087
Jetson Orin Nano (15 W)	FP16	0.916	18.07 ms	10.50 W	189.7 mJ/frame	0.087
Jetson Orin Nano (15 W)	INT8	0.843	16.00 ms	6.51 W	**104.2** mJ/frame	0.127
Jetson Orin Nano (7 W)	INT8	0.843	23.54 ms	6.30 W	148.3 mJ/frame	0.132
PolarFire SoC	INT8	0.916	55.00 ms	**2.46** W	135.3 mJ/frame	**0.366**

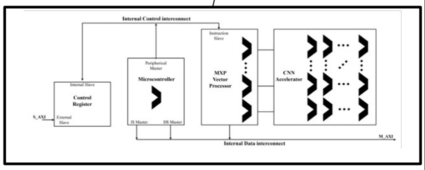

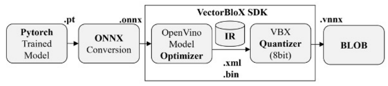

CoreVectorBlox SDK Flow

Thank you for your attention