FPGA-Enabled Machine Learning Applications in Earth Observation: A Systematic Review

CEDRIC LEONARD, CAPS, Technical University of Munich, Munich, Germany and MF-DAS, Remote Sensing Institute, German Aerospace Center (DLR), Weβling, Germany
DIRK STOBER, CAPS, Technical University of Munich, Munich, Germany
MARTIN SCHULZ, CAPS, Technical University of Munich, Munich, Germany

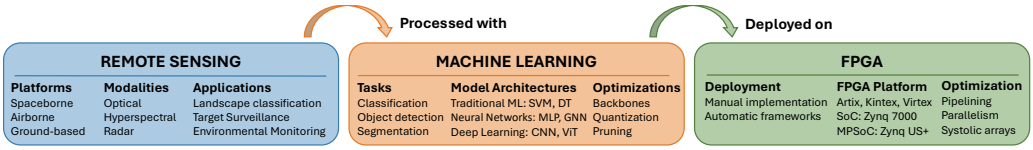

Fig. 1. Scope and content of the survey.

New UAV technologies and the NewSpace era are transforming Earth Observation missions and data acquisition. Numerous small platforms generate large data volume, straining bandwidth and requiring onboard decision-making to transmit high-quality information in time. While Machine Learning allows real-time autonomous processing, FPGAs balance performance with adaptability to mission-specific requirements, enabling onboard deployment. This review systematically analyzes 68 experiments deploying ML models on FPGAs for Remote Sensing applications. We introduce two distinct taxonomies to capture both efficient model architectures and FPGA implementation strategies. For transparency and reproducibility, we follow PRISMA 2020 guidelines and share all data and code at https://github.com/CedricLeon/Survey_RS-ML-FPGA.

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Computing methodologies** → *Machine learning*; • **Hardware** → *Reconfigurable logic and FPGAs*; • **Applied computing** → *Earth and atmospheric sciences;*

Additional Key Words and Phrases: Earth observation, remote sensing, neural networks, approximate computing

This work is supported by the Helmholtz Association under the joint research school "Munich School for Data Science – MUDS" and the German Federal Ministry of Research, Technology and Space, under the project BB-KI Chips, contract no. 16DHBKI020.
Authors' Contact Information: Cedric Leonard (corresponding author), CAPS, Technical University of Munich, Munich, Bavaria, Germany and MF-DAS, Remote Sensing Institute, German Aerospace Center (DLR), Weβling, Bavaria, Germany; e-mail: cedric.leonard@dlr.de; Dirk Stober, CAPS, Technical University of Munich, Munich, Bavaria, Germany; e-mail: dirk.stober@tum.de; Martin Schulz, CAPS, Technical University of Munich, Munich, Bavaria, Germany; e-mail: martin.w.j.schulz@tum.de.

ACM 0360-0300/2026/04-ART283
https://doi.org/10.1145/3800686

ACM Reference Format:
Cedric Leonard, Dirk Stober, and Martin Schulz. 2026. FPGA-Enabled Machine Learning Applications in Earth Observation: A Systematic Review. *ACM Comput. Surv.* 58, 11, Article 283 (April 2026), 36 pages. https://doi.org/10.1145/3800686

1 Introduction

The last decade of research has seen major innovations in **Machine Learning** (**ML**), particularly in **Deep Learning** (**DL**). Every discipline has seen publications leveraging ML's abstract learning capabilities to alleviate specific problems of the field. Nowadays, many state-of-the-art solutions incorporate ML methods in their pipeline. For instance, automatic navigation, for cars and **Unmanned Aerial Vehicles** (**UAVs**) alike, heavily relies on scene segmentation and obstacle detection. Similarly, video and image encoding apply learnable heuristics to select optimal patterns and splits in the data. Beyond vision, language processing benefits from the wide context available through transformer tokenization, significantly enhancing performance. By and large, these AI technologies have significant computational costs, making them inadequate in resource-constrained environments. Therefore, many applications deploying their algorithms close to the sensor, e.g., cars and UAVs On-Board Computers (OBCs), GoPros, or intelligent microphones, benefit from lightweight optimizations of these computationally intensive models.

In this survey, we take a systematic look at one application of ML for computer vision: **Remote Sensing** (**RS**) for **Earth Observation** (**EO**).[1] Contemporary challenges, such as the climate crisis, urban planning, or defense concerns, transformed RS into an essential field. In particular, the rising NewSpace[2] industry has decreased the costs and the production time of SmallSats. Because deploying CubeSats has never been so easy, we can observe a steady increase in the number of missions—to date, over 2,600 nanosatellites have been launched [62]. However, payload restrictions on these compact satellites reduce their onboard compute capacity relative to their larger counterparts, forcing them to run application-specific routines in flight and underscoring the need for efficient onboard data processing. In the particular example of RS, while the quantity of data grows, the bandwidth resources for the data transmission (downlink) stagnate [95]. Because of such limitations, most current missions acquire data on-demand or over certain pre-configured conditions, inevitably missing out relevant information. The increasing sensing resolution and number of in-orbit satellites, paired with further restrictions on the radio-frequency spectrum, will worsen the problem [103]. Processing the data directly onboard, transmitting only compressed or application-tailored information, constitutes an elegant solution for future missions. Equipping future missions with edge computing platforms opens up possibilities to relieve data processing stress. Because SmallSats' standard payloads typically include FPGA-based hardware [33, 110], this survey focuses on analyzing the literature deploying ML-based solutions for RS applications on FPGAs (Figure 1).

The main contributions of this survey are:

- **Systematic literature collection and analysis** ensuring a rigorous and transparent review process following the PRISMA 2020 guidelines (Section 2).
- **Comprehensive evaluation** of existing works, structured around eight **Research Questions** (**RQs**) to analyze the research landscape (Section 3), the design strategies (Section 4), and method synergies (Section 5).

[1]Throughout this survey, we will prefer the term RS, as it is the core technology of EO, which also encompasses in-situ measurements.
[2]https://philab.esa.int/tag/new-space/

— **Two specialized taxonomies**, one structuring RS applications based on their ML problem formulation (Section 3), and another categorizing FPGA implementations by frameworks and design patterns (Section 4).
— **Identification of research gaps**, structured under five subtopics to guide future advancements in FPGA-enabled ML for Earth Observation (Section 6).
— **Guidelines and recommendations** to facilitate benchmarking, enhance reproducibility, and drive future research efforts (Section 6).

1.1 Background

Remote Sensing (**RS**) is the science of obtaining reliable information about objects or areas on the Earth's surface without direct physical contact [14]. It involves the acquisition, processing, and interpretation of data captured by sensors that detect electromagnetic radiation. The discipline is broadly divided into active and passive approaches. Active systems, such as radar and LiDAR, emit their own energy signals and record the returning echoes, while passive systems depend on external sources of energy, primarily sunlight. Furthermore, RS platforms are classified by their operational altitude into spaceborne and airborne systems. Spaceborne platforms, typically mounted on satellites, provide extensive temporal and spatial coverage. They enable continuous monitoring over vast areas and across diverse time scales—from near real-time updates to the analysis of long-term environmental trends [145]. Airborne platforms, on the other hand, offer higher spatial resolution over more localized areas, making them ideal for detailed studies that require fine-grained spatial information.

Machine Learning (**ML**) methods have existed for several decades, but the past ten years have seen an unprecedented surge in their capabilities, largely driven by advances in computing power. On one hand, **Deep Learning** (**DL**) abilities to explore and learn from high-dimensional spaces have precipitated the growth of the discipline. In particular, foundation models and their generalization powers have been an intense research area over the last two years [114]. On the other hand, traditional ML methods such as **Decision Trees** (**DT**) [96] or **Support Vector Machines** (**SVM**) [25] are small and explainable methods still preferred in certain scenarios. Nevertheless, the vast majority of current ML research focuses on DL, equally in RS, where **Convolutional Neural Network** (**CNN**) and **Vision Transformer** (**ViT**) abilities to learn from images lay a solid basis for many RS applications [77, 114].

CPUs and, to a lesser extent, GPUs support a wide range of algorithms and kernels, requiring general-purpose hardware units such as multi-level caches, prefetching, and branch prediction. Many specific algorithms do not fully utilize these features, leading to wasted energy and transistors. FPGAs, by contrast, are reconfigurable integrated circuits adaptable to a specific algorithm, instantiating only the necessary hardware units. FPGAs' reconfigurability comes at the cost of lower operating frequency, requiring comparatively higher parallelism to compete with CPUs and the already highly parallel GPUs. In addition, such a high degree of freedom requires careful development. Programming FPGAs is similar to designing an integrated circuit and is traditionally done at low-level using **Hardware Description Languages** (**HDL**). However, with the growing adoption of AI on FPGAs, higher-level end-to-end toolchains are now available to implement **Neural Networks** (**NNs**) with popular kernels, such as CNNs [5, 117].

1.2 Related Work

We conduct this survey to combine the overlapping disciplines of Remote Sensing, Machine Learning, and FPGA deployment. To the authors' knowledge, no work has yet explored the intersection of these three research domains. However, multiple reviews have covered the usage of ML in RS, explored FPGAs for onboard processing, or investigated FPGA-based accelerators for ML separately—see

Table 1. Summary of Related Surveys

Reference	Year	RS Focus	ML Focus	FPGAs	Onboard	Years analyzed
Lopez et al. [78]	2013	✓	×	✓	Space	–
Bouhali et al. [12]	2017	≈	×	✓	UAV	–
Zhu et al. [145]	2017	✓	DL	×	×	2014–2017
Abdelouahab et al. [2]	2018	×	CNNs	✓	×	2015–2018
George and Wilson [33]	2018	≈	×	✓	Space	2000–2016
Shawahna et al. [108]	2019	×	DL	✓	×	2009–2018
Osco et al. [91]	2021	✓	DL	×	UAV	2012–2020
Lange et al. [66]	2024	≈	DL	×	Space	2018–2023
Our survey	2025	✓	DL + ML	✓	✓	2014–2024

✓ covered, × not covered, ≈ partially covered.

Table 1 for some relevant literature reviews. Notably, two surveys investigate the overlap between DL and FPGA design: Abdelouahab et al. [2] and Shawahna et al. [108]. Both studies draft a comprehensive overview of design paradigms to implement DL methods on FPGAs. Taking a closer look at onboard environments, George and Wilson [33] and Bouhali et al. [12] depict the specificities of onboard flying platforms, space- and airborne, respectively. With a similar scope, Lange et al. [66] specifically focus on the impact and mitigation techniques of radiation effects on DL models.

2 Methodology

This study follows the **Preferred Reporting Items for Systematic Reviews and Meta-Analyses** (**PRISMA**) 2020 guidelines [92] applicable to Computer Science, ensuring a complete, transparent, and reproducible review process.

2.1 Research Questions

This survey aims to examine articles using ML methods on FPGA hardware and discussing RS in EO. We specifically focus on RS sensors mounted on aerial vehicles (drones or UAVs) or spaceborne satellites. We explore this research intersection to summarize the community's work towards implementing automated methods onboard flying platforms, space- or airborne. To provide a clear structure for this survey, we address the following eight **Research Questions** (**RQs**). First, in Section 3, RQ1-3 establish the research landscape and define the population of studies included in our analysis.

- RQ1: Which Remote Sensing applications are addressed, and how are they formulated as ML problems?
- RQ2: Which Machine Learning models are most prevalent, and what are their key architectural characteristics?
- RQ3: What are the motivations for choosing FPGAs over other computing platforms (CPUs, GPUs, ASICs)?

Then, in Section 4, RQ4-6 explore the technical details of the deployed solutions. Here, we catalog ML optimization methods, FPGA design paradigms, and various levels of parallelism available to designers.

- RQ4: What are the effective optimization strategies to reduce model footprint and computational complexity?
- RQ5: How do FPGA design frameworks and patterns affect development effort and final performance?
- RQ6: How do different FPGA design approaches, such as pipelining, parallelization, and memory optimization, balance performance and constraints?

Finally, in Section 5, R7-8 take a holistic view of the accomplishments in this field, discussing key achievements and, in line with the current NewSpace era, the potential of AI on the edge for RS missions.

— RQ7: Given limited FPGA resources, which ML algorithms are suitable, and which RS applications can this platform support?
— RQ8: How are AI-powered edge computing solutions transforming RS missions?

To answer these research questions, we conduct a systematic search procedure adapted from the method proposed by the PRISMA 2020 guidelines [92]. The procedure starts by defining a prompt to frame for the population of records of interest. After collecting the results from the chosen research databases, we perform an initial screening pass via metadata analysis to filter out undesired records. The second screening pass involves reviewing the complete records, excluding any not clearly identified as being outside the survey's scope based on the metadata alone. The final curated collection of studies is enriched with additional articles coming from private libraries.

2.2 Search Procedure

We systematically retrieve relevant studies by querying Web of Science[3] with the prompt[4] below:

TOPICS: [*(Machine Learning* ***OR*** *Deep Learning* ***OR*** *Neural Network* ***OR*** *Convolutional Neural Network* ***OR*** *Recurrent Neural Network* ***OR*** *ML* ***OR*** *DL* ***OR*** *NN* ***OR*** *CNN* ***OR*** *RNN)* ***AND*** *(Earth Observation* ***OR*** *Remote Sensing* ***OR*** *Earth Science* ***OR*** *LiDAR* ***OR*** *SAR* ***OR*** *UAV* ***OR*** *Sentinel)* ***AND*** *(FPGA* ***OR*** *Field-Programmable Gate Array* ***OR*** *Field Programmable Gate Array* ***OR*** *Versal)*]
YEAR PUBLISHED: 2014-2024, ***LANGUAGE***: *ENGLISH*

The central part of the prompt consists of three expressions framing for any ML model, the usage of RS data in EO, and the explicit mention of FPGA devices. We choose not to consider articles published more than 10 years ago. Indeed, the FPGA technology, the state-of-the-art DL models, and even the tackled problems have evolved rapidly. Therefore, work from more than 10 years ago is often too different from modern studies to be compared fairly. After retrieving the 119 results of the prompt, we proceed to screen the articles to filter mismatches. To this effect, we pick four **Exclusion Criteria** (**EC**):

— EC1 - Discussion and review papers, without experiments.
— EC2 - Publication unrelated to or not using ML models;
— EC3 - Publication unrelated to or not using RS data for EO;
— EC4 - Publication unrelated to or not using FPGA devices;

While EC1 and EC2 are self-explanatory, EC3 mostly rules out non-EO applications of remote sensing technology, e.g., LiDAR for autonomous driving or SAR calibration. EC4 excludes studies with no experiments or details about the FPGA implementation. For example, EC4 eliminates records mentioning FPGA hardware as a possible solution or using a different FPGA acronym, e.g., fast patch-free global learning. We design this filtering process to retain only the most relevant, experiment-based studies that reflect the current state of FPGA-based ML in Remote Sensing. After screening, 46 articles are kept for review for a total of 48 records when adding the records from private libraries. Each article included in the survey is thoroughly read and tagged in Zotero [24]. In particular, we use Zotero's tagging system and API to systematically report the experiments' specificities and performance metrics. The entire database extracted via Zotero's API, along with the code used for analysis, figures, and main taxonomy tables, is available on GitHub at

[3] https://www.webofscience.com/, last consulted on 06 March 2025.
[4] The "**TOPICS**" field searches across each item's title, abstract, author keywords, and keyword plus in the database. The prompt is case-insensitive.

https://github.com/CedricLeon/Survey_RS-ML_FPGA. Our repository also hosts supplemental material, such as extended methodological considerations and the PRISMA 2020 flow diagram.

3 Mapping the Research Landscape: Applications, Models, and Hardware

In this section, we analyze the motivations, tackled problems, and solutions presented by each surveyed study.

3.1 Remote Sensing/Machine Learning Taxonomy

To help readers explore studies addressing similar problem settings, we present in Table 2 a taxonomy grouping the collected studies based on how RS applications are formulated as ML tasks. When a study conducts several relevant experiments, we extract each corresponding model and metric. For example, Suh et al. [112] present a methodological contribution and evaluate several sizes of the same architecture, i.e., SSD 0.25x, SSD 0.5x, and SSD 1.0x. We argue that reporting each experiment separately, rather than considering a single experiment per study, portrays a more complete picture of the conducted work. In total, the taxonomy includes 68 experiments drawn from 48 studies.

Each row in Table 2 corresponds to an experiment. In the first column, we classify the ML problem into different CV tasks. To facilitate comparison with other ML works, we describe the ML *Task* using CV terminology over RS nomenclature.[5] In addition, we distinguish between *Segmentation* models that process a pixel at a time ("Segmentation - Pixel") and those that process the patch as a whole ("Segmentation - Patch"), as the throughput requirements of the latter are significantly higher. Compared with the first column, the rest of the taxonomy uses RS terminology, as it focuses on the goals of the experiments, i.e., the RS-relevant information they generate. In the following group of columns, dedicated to the *RS Problem*, experiments are hierarchically grouped by data modality (*Mod.*), i.e., the type of RS data used as input, *Application*, and *Dataset*. As a result, experiments from the same study may not be juxtaposed. *Application* refers to the downstream task tackled by the study; applications are closely tied with the *Dataset*[6] and are detailed later in Figure 2. The *Article* group provides the experiment's context: *Ref.* lists the source article and *Year* indicates its publication date.

The rest of the table is dedicated to the ML component of the experiment. *ML Model* characterizes the model's architecture through three fields: *Original Name*, the exact appellation used in the study; *Core*, a high-level architectural descriptor (e.g., CNN, Shallow NN such as a **Multi-Layer Perceptron** (**MLP**), or Traditional ML like an SVM), further detailed in Figure 3; and, for **Deep Neural Networks** (**DNNs**), their *Backbone*, i.e., the core of the network responsible for extracting features. Next, the *FPGA* column indicates the family of the FPGA device used for implementation, more details in Figure 4. The last group of columns *Performance* gathers a few metrics of the experiment as reported in the original study: *Score* corresponds to the accuracy in **[%]** of the model on its dataset and *MF**[MB]*** is the Memory Footprint of the model in MegaBytes. Finally, *C**[OP]*** stands for the computational complexity of the model in numbers of operations. These performance metrics are the most closely related to the ML model; further metrics are reported during the second literature taxonomy in Table 4.

3.2 Formulating RS Applications as ML Problems (RQ1)

From disaster assessment to target monitoring, RS images support a wide and growing range of EO applications [142].

[5]For example, we use the term *Classification* for per-image class assignment while it implies a per-pixel class assignment in RS.

[6]Datasets carrying the *(cust.)* tag are unpublished datasets, meaning that the authors built their dataset but did not name or publish it.

Table 2. RS/ML Taxonomy Table

ML	Remote Sensing			Article		ML Model				FPGA	Performance		
Task	Mod.	Application	Dataset	Ref.	Year	Original Name	Core	Backbone	Eq. Model	Family	Score [%]	MF[MB]	C[OP]
Classification	RGB	Cloud coverage class.	L8 Biome [27]	[95]	2022	CloudSatNet-1 Q2	Custom CNN	-	CNN	Zynq 7000	83.41 (OA)	1.43	-
				[95]	2022	CloudSatNet-1 Q4	Custom CNN	-	CNN	Zynq 7000	87.42 (OA)	3.06	-
			Sentinel-2 (cust.)	[98]	2021	CloudScout	Custom CNN	-	CNN	Zynq US+	92.0 (OA)	13.30	-
			RGB (cust.)	[118]	2024	ResNet-50	ResNet	ResNet-50	CNN	Zynq US+	-	-	-
				[118]	2024	SICNet	Custom CNN	-	CNN	Zynq US+	85.76 (OA)	-	-
			Landsat-8 (cust.)	[119]	2024	CNN	Custom CNN	-	CNN	Artix-7	89.64 (F1)	-	-
		Land Use Land Cover	UCM Land Use [134]	[139]	2023	A2NN	VGG	VGG11	CNN	Virtex-7	94.76 (OA)	-	14.99G
			NWPU-RESISC45 [20]	[141]	2020	Q-IORN	VGG	VGG16	CNN	Virtex-7	88.31 (OA)	121.51	-
				[131]	2022	Improved VGG16	VGG	VGG16	CNN	Artix-7	88.08 (OA)	14.80	40.96G
				[88]	2023	ResNet-34	ResNet	ResNet-34	CNN	Virtex-7	92.81 (OA)	21.29	7.33G
				[88]	2023	VGG16	VGG	VGG16	CNN	Virtex-7	91.90 (OA)	14.70	30.69G
			UAV RGB (cust.)	[81]	2019	MLP	Shallow NN	-	MLP	Spartan-3A	95.14 (OA)	-	-
		Ship id.	MASATI [30]	[85]	2022	AlexNetLite	AlexNet	AlexNet	AlexNet	Kintex US	89 (F1)	-	5.71G
				[85]	2022	MobileNetv1Lite	Custom CNN	MobileNetv1	CNN	Kintex US	94 (F1)	-	0.05G
			Kaggle SSI	[56]	2024	HO-ShipNet	Custom CNN	-	CNN	Zynq 7000	93.44 (F1)	-	-
	SAR	Military targets id.	MSTAR [107]	[125]	2019	LeNet-5 f32	LeNet	LeNet-5	LeNet-5	Kintex-7	98.76 (OA)	6.64	-
				[125]	2019	LeNet-5 i8	LeNet	LeNet-5	LeNet-5	Kintex-7	97.77 (OA)	1.66	-
				[18]	2020	LeNet-5	LeNet	LeNet-5	LeNet-5	Kintex-7	98.18 (OA)	-	-
				[137]	2022	GNN	GNN	-	GNN	Zynq US+	99.09 (OA)	0.96	2.13M
				[138]	2023	GNN	GNN	-	GNN	Alveo	-	-	-
				[126]	2024	VTR	ViT	-	ViT	Alveo	95.96 (OA)	-	18.28G
		Ship id.	*ALOS-2 (cust.)*	[47]	2019	CNN	Custom CNN	-	CNN	Zynq 7000	100 (OA)	-	-
				[83]	2020	BNN	Custom CNN	-	CNN	Zynq 7000	-	-	-
	HSI	Air quality class.	*UAV RGB (cust.)*	[54]	2024	ETAUS	ResNet	ResNet-50	CNN	Zynq US+	86.38 (F1)	-	-
Segmentation - Pixel	RGB	Deforestation det.	*UAV RGB (cust.)*	[116]	2020	Weightless Neural Systems	Shallow NN	-	WNS	Zynq 7000	90.0 (OA)	-	-
		UAV landing site id.	*UAV RGB (cust.)*	[28]	2018	Decision Tree	Trad. ML	-	ML	Zynq 7000	92.1 (OA)	-	-
	RGB†	Land Use Land Cover	*ALSAT-2A (cust.)*	[130]	2017	BRAM_DSP	Trad. ML	-	Fuzzy ARTMAP	Virtex-6	94.8 (OA)	-	-
				[130]	2017	LUT_MUL	Trad. ML	-	Fuzzy ARTMAP	Virtex-6	94.8 (OA)	-	-
	HSI	Land Use Land Cover	University of Pavia	[109]	2021	LPDBL	Trad. ML	-	ML	Virtex-7	-	-	-
				[49]	2023	2D CNN	Custom CNN	-	CNN	Zynq US+	98.24 (OA)	1.20	0.684M
				[49]	2023	3D CNN	Custom CNN	-	CNN	Zynq US+	94.09 (OA)	0.12	4.231M
				[49]	2023	HybridSN	Custom CNN	-	CNN	Zynq US+	100 (OA)	20.50	101.642M
				[79]	2024	SVM	Trad. ML	-	ML	Zynq 7000	82.48 (AA)	-	-
				[16]	2024	SAM-GNN	GNN	-	GNN	Virtex-7	95.05 (OA)	-	-
			AVIRIS-NG	[42]	2022	CAG-SC2S	Trad. ML	-	ML	Artix-7	98.30 (OA)	-	-
				[42]	2022	CAL-SC2S	Trad. ML	-	ML	Artix-7	99.69 (OA)	-	-
				[42]	2022	SVM	Trad. ML	-	ML	Artix-7	20.71 (OA)	-	-
		Land cover change det.	ABU [61]	[13]	2023	Deep Belief Network	Shallow NN	-	RBM	Zynq US+	-	-	-
Segmentation - Patch	RGB	Urban areas	Potsdam [57]	[104]	2021	ENet	Custom CNN	-	CNN	Zynq US+	63.3 (mIoU)	0.36	4.06G
				[104]	2021	ESPNet	Custom CNN	-	CNN	Zynq US+	55.5 (mIoU)	0.33	3.71G
				[104]	2021	FPN	Custom CNN	-	CNN	Zynq US+	65.3 (mIoU)	5.84	17.30G
				[104]	2021	U-Net	U-Net	-	CNN	Zynq US+	61.4 (mIoU)	7.40	96.68G
		Cloud extr.	38-Cloud [82]	[8]	2019	C-FCN++	Custom CNN	-	CNN	Cyclone V	79.30 (mIoU)	0.047	4.7K
			Landsat-8 (cust.)	[62]	2024	TriCloudNet + U-Net	U-Net	SqueezeNet	CNN	Zynq 7000	6.84 (CRE)	3.61	-
		Land Use Land Cover	*Landsat-8-OLI (cust.)*	[99]	2021	RDBC	Trad. ML	-	ML	Virtex-6	37.45 (SSIM)	-	O(n)
Object detection	RGB	Diverse det.	DOTAv1.0 [128]	[140]	2021	Improved YOLOv2	YOLO	DarkNet19	YOLOv2	Zynq 7000	67.32 (mAP)	-	379G
				[131]	2022	Improved YOLOv2	YOLO	DarkNet19	YOLOv2	Artix-7	67.30 (mAP)	49.40	379.55G
				[133]	2023	Ghost-YOLOS	YOLO	GhostNet	YOLOv3	Zynq 7000	62.58 (mAP)	12.60	8.41G
				[88]	2023	Improved YOLOv2	YOLO	DarkNet19	YOLOv2	Virtex-7	67.30 (mAP)	49.40	379.55G
			DIOR [73]	[143]	2023	YOLOv4-MobileNetv3	YOLO	MobileNetv3	YOLOv4	Zynq US+	82.61 (mAP)	5.69	-
				[70]	2023	RFA-YOLO	YOLO	MobileNeXt	YOLOv4	Zynq US+	64.85 (mAP)	24.75	-
			NWPU VHR-10 [21]	[74]	2019	CBFF-SSD	SSD	MobileNetv1	SSD	Zynq 7000	91.42 (mAP)	-	-
			DAC 2018	[122]	2019	CNN	Custom CNN	-	CNN	Zynq 7000	21 (mIoU)	-	-
			UAV RGB (cust.)	[75]	2020	AP2D-Net	Custom CNN	-	CNN	Zynq US+	55.0 (mIoU)	-	-
		Flying-object det.	PennSyn2Real [87]	[112]	2021	SSD 0.25x	SSD	VGG16	SSD	Zynq US+	76.2 (mAP)	-	4.04G
				[112]	2021	SSD 0.5x	SSD	VGG16	SSD	Zynq US+	83.9 (mAP)	-	15.65G
				[112]	2021	SSD 1.0x	SSD	VGG16	SSD	Zynq US+	88.4 (mAP)	-	61.60G
			UAV RGB (cust.)	[86]	2024	YOLOv4	YOLO	DarkNet53	YOLOv4	Zynq US+	83.7 (mAP)	245	105.9G
				[86]	2024	YOLOv4-tiny 3L	YOLO	DarkNet53-tiny	YOLOv4	Zynq US+	72.24 (mAP)	24.10	15.32G
		Aircraft det.	Airbus	[31]	2024	ResNet-18+YOLOv2	YOLO	ResNet-18	YOLOv2	Zynq US+	91 (AP)	-	-
			GES RGB (sim.)	[127]	2021	Improved YOLOv3	YOLO	DarkNet53	YOLOv3	Zynq US+	92 (OA)	8.50	-
		Railway defect det.	*FastenerDataset (cust.)*	[136]	2024	Improved YOLOv4-tiny	YOLO	DarkNet53-tiny	YOLOv4-tiny	Zynq US+	95.1 (mAP)	-	-
	RGB‡	UAV obstacles det.	*UAV RGB+MMW (cust.)*	[123]	2024	Lightweight CNN	ResNet	ResNet-18	CNN	Zynq 7000	70.1 (mAP)	-	4.6G
	SAR	Ship det.	SSDD [72]	[132]	2022	CNN2@0.7	YOLO	MobileNetv1	YOLOv2	Virtex-7	94 (AP)	0.09	0.45G
				[132]	2022	CNN4@0.7	YOLO	MobileNetv2	YOLOv2	Virtex-7	93.3 (AP)	0.11	0.56G
				[132]	2022	CNN6@1.6	YOLO	SqueezeNet	YOLOv2	Virtex-7	92.8 (AP)	0.09	0.6G
Regr.	SAR	Oil spills feature extr.	*SAR (sim.)*	[43]	2022	MLP	Shallow NN	-	MLP	Zynq 7000	-	-	-
				[44]	2024	TinyML	Shallow NN	-	MLP	Zynq 7000	99.00 (OA)	0.000037	-

Task: **Regr.** Regression; *Dataset:* **GES** Google Earth Studio, **MMW** MilliMeter-Wave radar; *Model:* **RDBC** Roller Dung Bettle Clustering, **WNS** Weightless Neural System
Modality: † = RGB + infrared; ‡ = RGB + MMW.

Overview of RS Applications. Figure 2 separates the surveyed RS applications into four distinct thematics. *Target Surveillance*[7] gathers all applications focused on recognizing a specific family of objects. The output of such a task can be the location—bounding box—of the targets, in such a

[7] *Target detection* is a more frequent appellation for RS applications. However, we will prefer the term "surveillance" to avoid any confusion with the ML task "object detection," sometimes also called "target detection."

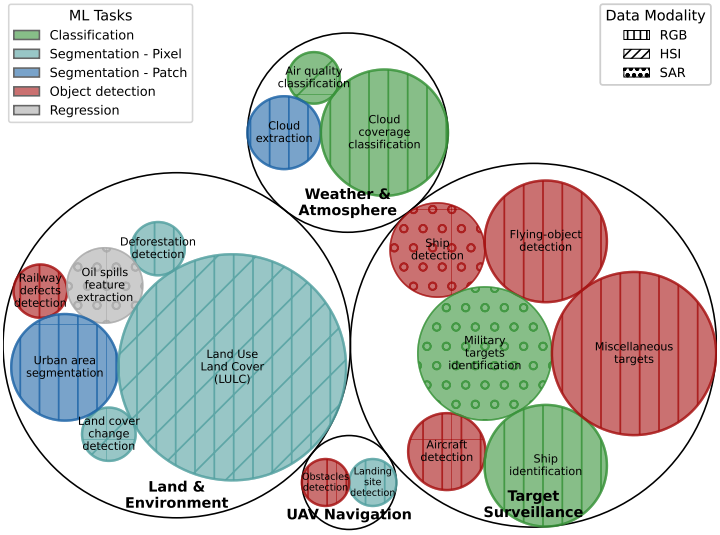

Fig. 2. Addressed RS applications grouped by thematic and colored as their corresponding ML task. The circles' size represents the number of experiments, while the hatching pattern represents the prevailing data modality.

case we denote it detection , or the binary presence of the target in the observed image, denoted identification . In this thematic, "Miscellaneous targets" refers to heterogeneous object families, such as ships and cars, but also tennis courts or roundabouts in the DOTAv1.0 [128] or DIOR [73] datasets. "Flying-object detection" and "Aircraft detection" differentiate objects in the air (planes, drones, birds, etc.) and standing aircraft on the ground. *Land & Environment* applications monitor the type of terrain present in the image, either on a general level, such as Land Use Land Cover (LULC), or for specific phenomena, such as oil spills. In particular, "Land cover change detection" performs semantic segmentation of the land in an unsupervised fashion to identify pixels with abnormal values. In *Weather & Atmosphere*, we distinguish experiments that extract cloud masks from images and experiments that classify complete scenes based on cloud coverage thresholds. Finally, *UAV Navigation* applications rapidly use the acquired images to support the UAV's itinerary. While the grouping of Figure 2 offers a clear structure for the surveyed studies, it does not capture the full range of RS applications. In particular, the records miss prevalent applications such as natural disaster monitoring or studies of the climate crisis.

Data Modalities in RS. The hash pattern in Figure 2 represents the modality of the images used for each application. In RS, the modality refers to the type of sensor used to acquire the data. The most common modality, as in this survey, is optical data,[8] with the three traditional **Red-Green-Blue** (**RGB**) channels. In RS, channels are commonly referred to as bands. Indeed, sensors capture data around specific wavelengths, each wavelength's spectrum resulting in a different band. **Hyperspectral Imaging** (**HSI**) emerges as the most extreme case, sometimes with over 200 bands, providing a very fine-grained spectral resolution that approximates a contiguous spectrum [78]. Lastly, **Synthetic Aperture Radar** (**SAR**) is an active sensor, in opposition to passive sensors

[8]Optical data is commonly named Multispectral in RS. However, as ~95% of the surveyed studies only use the RGB channels of the electromagnetic spectrum, we use the term optical or RGB for accuracy.

that collect light from the observed scene. SAR emits waves that scatter upon hitting the observed scene, then captures the reflected components that bounce back towards the sensor. Such a system enables data collection during nighttime and through clouds, making SAR a versatile modality.

Distinctions Between ML Tasks. In Figure 2, each application circle is colored depending on the ML task it is primarily formulated as. An ML task is the type of prediction made by the model. The most traditional tasks are *Classification* and *Regression*. Also referred to as *Scene Classification* and *Parameter Retrieval* in RS terminology, the former assigns a category to each input image, while the latter predicts a numerical value from its features. Beyond these, a Clustering task groups data instances into clusters based on shared characteristics. When applied to images, Clustering yields a segmentation map, leading us to prefer the term *Segmentation*. Semantic segmentation maps can be generated by processing the entire input image as a whole or by processing pixels one by one repetitively across the image. Finally, *Object Detection* not only assigns categories but also provides bounding boxes indicating the location of objects, enabling the identification and categorization of multiple objects within a single image. While some applications are closely tied to a specific ML task, others are more versatile and can be formulated in various ways. Figure 2 includes one example: "LULC", was formulated once as *Segmentation - Patch*, six times as *Classification*, and 12 times as a *Segmentation - Pixel* task. Ultimately, how RS applications are formulated as ML tasks reflects the decisions and priorities of each study.

Converting RS Problems into ML Solutions. Target Surveillance is a major focus in the literature. Though uncommon in satellite imagery due to the typical low resolution of the images, this trend might be explained by the proportion of airborne data. Indeed, UAVs offer higher-resolution imagery, which supports higher detection accuracy. To mitigate low accuracies, many experiments express *Target Surveillance* applications as scene-wise *Classification*, a solution relying on prudent patchification of the original full-size RS image.[9] Furthermore, we notice a significant focus on estimating cloud coverage. Such an application is motivated by the importance of clouds in optical images. Indeed, throughout the year, clouds cover over 66% of Earth's surface [119], significantly reducing the information obtainable from RGB images.[10] Regarding ML tasks, most surveyed studies frame their applications as supervised *Classification*. In contrast, *Regression* appears only twice, making it the most under-represented task, although retrieving parameters from a scene is a critical problem in RS. Similarly, most datasets used for Segmentation - Pixel are labeled and result in supervised problems, except for the LULC unsupervised classification of Gyaneshwar and Nidamanuri [42]. The semantic segmentation of urban areas follows the same trend, though Ratnakumar and Nanda [99] use unsupervised clustering. Finally, all *Object detection* applications use multi-class datasets and single-shot models to detect multiple instances in a unique pass.

RQ1 TAKEAWAYS

Target surveillance is the prominent focus of the surveyed literature (44%). Frequently simplified as scene-wise classification, 31% of the applications address defense concerns. Landscape and environment monitoring dominates the rest of the research (40%). Regarding data modalities, RGB data is dominant (65%), with HSI remaining niche and SAR primarily used for water environments.

[9]RS images are too large for most ML models to process directly. The process of patchification divides them into smaller, more manageable patches.

[10]Conversely, SAR wavelengths are specifically chosen to be outside the absorption spectrum of clouds, allowing radar signals to pass through without significant attenuation.

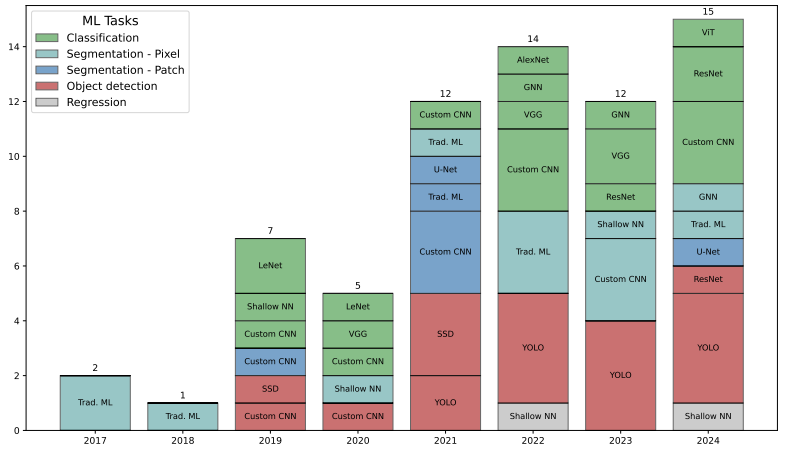

Fig. 3. ML model core architectures grouped per year and colored by their corresponding tasks.

3.3 Diversity and Trends in ML Models (RQ2)

With RQ1, we analyzed the wide spectrum of RS applications and discussed how each experiment translates into an ML task. In RQ2, we study the diversity of models used to interpret this information. Figure 3 shows the yearly distribution of models from the surveyed studies, categorized by ML task and labeled by each model's "core", i.e., its key architectural characteristics. Although this survey includes studies from 2014 onward, the earliest publication appears in 2017, with significant exploration beginning in 2019. From this point onward, the number of experiments and the diversity of architectures steadily increased, reflecting growing interest in the domain.

Evolution of ML Approaches to RS Problems. Surprisingly, among semantic segmentation tasks, more models operate per-pixel rather than processing the input patch at once, despite the latter being the modern approach for pixel-level analysis [6]. Indeed, segmentation networks, like U-Net [102], have access to spatial information and generally outperform methods limiting spatial neighborhood information. However, resource constraints and computation overload also drive the model choice. Unlike *Segmentation - Patch* deep architectures, *Segmentation - Pixel* experiments often employ shallow networks or traditional ML, which are not only less demanding computationally but also easier to implement and parallelize. This combination of efficiency and simplicity likely explains their prevalence in this FPGA-focused survey, despite the broader RS trend favoring *Segmentation - Patch*. The remaining problems are formulated as *Classification* and *Object detection* tasks. Together, these two categories represent two-thirds of the experiments. Such a trend can be explained by the versatility of formulating problems as *Classification* tasks, as well as the importance of target monitoring in RS.

Examining Model Core Architectures. Model cores are chosen to highlight the fundamental differences between architectures. In particular, traditional ML methods, such as SVMs [25] or Decision Trees [96], are isolated from **Neural Networks** (**NNs**), just as Shallow Networks are distinguished from **Deep Neural Networks** (**DNNs**) [67]. We split CNNs—the vast majority of the models analyzed in this survey—between renowned architectures, such as LeNet [68], AlexNet [64], VGG [111], U-Net [102], or ResNet [48], and custom CNN architectures. Notably, "Custom CNN" includes both newly designed CNNs and named architectures used only once in this survey, such as SqueezeNet [55] or ENet [93]. Finally, we classify YOLO (v2, v3, v4) [11, 100, 101] and SSD [76] architectures separately, even though renowned CNN architectures are frequently used as the backbones of

detection networks. Indeed, the detection layers added to the feature extraction network (Fully Connected (FC) layers for YOLO and convolution layers for SSD) introduce fundamental differences to the architectures, such as the multiple heads of YOLO networks.

A CNN-Focused Landscape. Although shallow networks, such as MLPs, and traditional methods, like Decision Trees and SVMs, are featured in some studies [42, 79]—often due to their compatibility with resource-limited hardware—, the vast majority of contributions leverage DL algorithms. Within DL, convolutional architectures are strongly preferred. Models like AlexNet, ResNet, and VGG are occasionally used, often for benchmarking purposes, but custom-designed CNNs are the most prevalent (27% of all models), reflecting designers' needs for medium-sized architectures typically with 6 to 10 layers. All detection networks are single-shot architectures like YOLO and SSD. Two-stage detection alternatives like R-CNN [35] are not used in the analyzed experiments. Indeed, even if YOLO and SSD architectures are typically larger than R-CNN, the absence of a second pass through the networks highly reduces the computational complexity and processing time, making single-shot networks efficient solutions for resource-limited platforms like FPGAs [143].

RQ2 TAKEAWAYS

With a steadily growing number of publications per year, formulating RS applications as *Classification* problems is prevailing (35%). Of all the models, 74% are convolution-based, of which 26% are custom CNNs with an average of 9 layers. All detection models are single-shot networks and 89% of the traditional ML methods are used in *Segmentation - Pixel* tasks.

3.4 The Role of the Computing Platform (RQ3)

As explored in the previous sections, deploying AI-supported RS applications in resource-constrained environments requires formulating the application as a suitable ML task and selecting an appropriate model. These decisions are deeply intertwined with the selection of the hardware computing platform.

Hardware Constraints Across RS Platforms. RS platforms are vehicles that carry remote sensors; they can operate at various altitudes as spaceborne, airborne, or ground-based systems. Each platform operates under distinct hardware constraints. **Size, Weight, Power, and Cost (SWaP-C)** define the fundamental tradeoffs between computational capability and mission feasibility. Additionally, payloads may require radiation hardening or thermal management. For example, spaceborne systems face the most stringent constraints. They are exposed to harsh environmental conditions, including **Single-Event Effects (SEEs)**[11] in onboard electronics caused by radiation exposure [66], extreme temperature swings from −150°C to +150°C [40], and tight power budgets from 20W to 95W for SmallSats payloads [63]. These extreme conditions make spaceborne systems the most constrained environment for FPGA-based ML deployment and, accordingly, the primary focus of the surveyed studies. Airborne platforms, targeted by ∼ 25% of the articles, present moderate constraints that differ significantly between sub-categories [124]. Micro and nano UAVs operate at low altitudes and have short endurance, typically requiring light payloads and possibly high power if taking off vertically. On the other hand, high- and medium-altitude aircraft can typically accommodate substantial payloads and possess longer endurance. Finally, ground-based systems operate with minimal SWaP-C constraints and are absent from this review. Given these

[11] SEEs manifest as Single-Event Upsets (SEUs)—soft errors altering memory—or as Single-Event Latch-ups (SELs)—potentially destructive hard errors requiring a power reset.

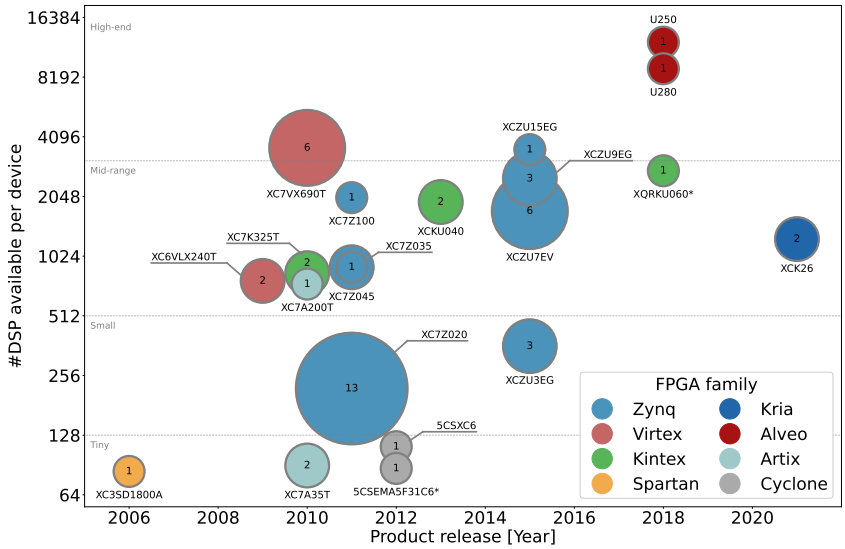

Fig. 4. Distribution of FPGAs used across articles. Unlike other figures/tables, it shows FPGA products used per article (48) instead of per experiment (68). * indicates additional boards used in the studies, absent from the experiments reported in Table 4. The horizontal dotted lines represent arbitrary distinctions between size categories used in RQ7

platform-specific limitations, this section examines how FPGAs address such challenges and how they compare with conventional computing hardware for edge AI deployment.

FPGA Usage in RS. FPGAs are reconfigurable integrated circuits, that offer design flexibility at the cost of programming complexity. Often available off-the-shelf, FPGAs can operate independently on the edge or connect to other computational platforms via network cables or PCIe. Figure 4 presents the different FPGAs used in each study. The devices are spread along the x-axis depending on their year of release, while the y-axis represents the number of DSP slices present on the board—relevant for arithmetic computations, like ML. Except for one experiment using the Cyclone V FPGA from Intel, we observe that all considered studies deploy hardware from AMD (formerly Xilinx). The most common family are Zynq FPGAs, used by 57% of the studies, which package FPGA programmable logic with CPUs on a single die to support networking and general-purpose tasks. Such products will be referred to as **System-on-Chip** (**SoC**) FPGAs. In contrast, datacenter FPGAs, such as the Alveo Series, are often offered as PCIe cards that can be inserted into a server. Furthermore, to address the need for ionizing-radiation tolerance to ensure mission reliability, FPGA vendors offer radiation-hardened and radiation-tolerant variants. AMD's KCU105 kit, for example, is a commercial equivalent of the space-grade Kintex Ultrascale XQRKU060 used in the VIDEO project [85]. In addition to FPGAs, designers can select from several families of **Processing Units** (**PUs**) when deploying AI models on edge devices. In the following subsections, we look at the alternatives: CPUs, GPUs, and ASICs (custom, TPUs, and VPUs). We analyze the authors' motivations to select FPGA over alternatives and mention, when available, the comparisons made by the studies. In the end, we summarize the unique advantages of each family of PUs in Table 3.

FPGAs vs. CPUs. **Central Processing Units** (**CPUs**) are ubiquitous general-purpose processing units, widely available as **Commercial Off-The-Shelf** (**COTS**) components. While well-suited for complex data transformations and system management, CPUs are limited in parallel processing

capabilities, often falling behind specialized hardware in raw computational power. Evaluating FPGA-based designs against CPUs is a common practice in the surveyed studies. These comparisons often involve desktop and server processors, including older Intel i7 models (up to 8th gen) [47, 70, 79, 133], newer i7/i9 models (12th and 13th gen) [54, 138], AMD Ryzen 3990x [137], and Broadwell Xeon server CPUs [88, 127, 141]. While these desktop and server CPUs offer high computational performance, they are not designed for energy-efficient applications. In contrast, mobile CPUs, such as the AMD 4800H [137, 143], provide a more power-conscious alternative. Similarly, Huang et al. [54] compare their ETAUS system against an embedded ARM Cortex A53 CPU and achieve significant performance improvements—×25 GOP/s/W and ×37 FPS. Superior power efficiency, compared with both CPUs and GPUs, is the most frequently cited advantage of FPGAs [8, 18, 42, 49, 70, 75, 81, 131, 133, 140]. For instance, Yang et al. [133] reports 29× higher efficiency compared with an Intel Core i7-8700 K, and Grover and Soni [41] demonstrate different design techniques to optimize power consumption. Although some CPUs and microcontrollers, such as ARM Cortex-R or LEON families, are radiant-tolerant, FPGAs' inherent reconfigurability facilitates the implementation of mitigation techniques, like **Triple Modular Redundancy** (**TMR**), where critical logic is triplicated and majority voting used to mask errors. On the downside, rad-hard FPGAs are significantly more expensive than rad-hard CPUs and microcontrollers. Their SRAM-based memory is also vulnerable to SEUs, requiring mitigation techniques [104, 110]. Siegle et al. [110] highlight that, although CPUs and microcontrollers remain essential onboard, especially for the OBC, modern payload processing applications increasingly rely on FPGAs. Over time, new solutions to deploy ML algorithms on embedded devices have been developed, shifting the focus from CPUs to domain-specific accelerators [97, 108].

FPGAs vs. GPUs. **Graphics Processing Units** (**GPUs**) possess many small cores, enabling high parallelism and performance for tasks involving massive amounts of data and large matrix computations. They are widely used for training NNs. Consequently, many studies compare FPGA-based designs against high-performance NVIDIA GPUs. These include Maxwell [47, 127] and Pascal architectures [74, 88, 131, 139–141], which introduce native half-precision support. Later architectures, such as Volta [70, 83] and Turing [86, 132, 133, 143], contain Tensor Cores, processing units dedicated for Deep Learning operations. Ampere GPUs [136–138] introduce native support for low precision integer operations, thus heavily benefiting from quantization. The latest Blackwell GPUs [90] support precision as low as 4-bit floating-point, decreasing FPGAs' benefits gained from using custom datatypes. FPGAs can exhibit better power efficiency [125] compared with desktop GPUs. For example, Li et al. [70] report 134% of the FPS of a Tesla V100 while consuming 10× less power. In contrast, Li et al. [74] find that an NVIDIA GTX 1070 Ti outperforms a Zynq-7000 in GOP/s/W efficiency by a factor of three showing, that power efficiency is highly implementation-dependent. Regardless of efficiency, GPUs' considerable processing capabilities come at the cost of significant memory and power consumption—often exceeding 300W [22]—making them impractical for most edge deployments. To address these constraints, computing boards, like the NVIDIA Jetson Nano [89], combine CPUs with energy-efficient edge GPUs. Even though edge GPUs usually deliver higher peak performance [74], FPGAs can still achieve better power efficiency [125]. For example, Huang et al. [54] report that their FPGA-based system ETAUS achieves 2.7× GOP/s/W compared with Jetson Nano. Similarly, the AP2D-Net presented by Li et al. [75] achieves 2.8× FPS/W compared with the best model of the DAC System Design Contest 2018 [129] deployed on a Jetson TX2. The complex architecture and high transistor density of GPUs pose significant challenges for circuit-level radiation hardening. While the NVIDIA Tegra K1 has withstood radiation levels present in short-duration LEO missions [7], no edge GPU, to the authors' knowledge, has been explicitly designed for space-grade reliability. In summary, while GPUs—especially edge

Table 3. Summary of Different Computing Platforms Used in the Different Studies

Chip	Dev. Time	Price/Unit	Rad-Hard (Avail.)	Updatable	Pow. Eff.	Latency	Peak Perf. (*fp32*)
CPU	Short	Low	Yes	Yes	Low	Medium	Medium
GPU	Medium	Medium/High	No	Yes	Medium	High	Very high
Custom ASIC	Very long	Very low	Yes	No	Very High	Very Low	Custom
FPGA	Long	High	Yes	Yes	High	Low	Medium

GPUs—offer strong computational performance and power efficiency trade-offs for ML tasks, FPGAs remain the preferred choice in power- and space-constrained environments due to their energy efficiency and radiation tolerance.

FPGAs vs. ASICs (TPUs/VPUs). **Application-Specific Integrated Circuits** (**ASICs**) are fully custom-designed chips optimized for specific applications, potentially offering unmatched performance and energy efficiency. Yet, their design is a complex, lengthy, and expensive process—no study has compared an FPGA-based design against a suitable custom ASIC. Many studies list reconfigurability as a key advantage of FPGAs over ASICs [75, 99, 116, 130, 133, 140]; in defense and space for example, re-programmability enhances mission versatility by enabling the "deploy and program" paradigm[12] [49, 70]. Consequently, FPGAs are increasingly replacing traditional ASICs due to their faster time-to-market, cost-efficiency in small-batch applications, and competitive performance [98]. Despite these advantages, ASICs remain the optimal choice for ultra-low-power applications, as the generic nature of FPGAs cannot match the efficiency of fully customized hardware [4]. **Tensor Processing Units** (**TPUs**) are commercially available ASICs optimized for Deep Learning, primarily used in data centers [59]. While the Google Coral TPUs [39] target edge devices, no study in this survey directly compared FPGA designs to a TPU. Li et al. [74] only compare their Zynq 7000 design against Google cloud TPU's theoretical peak performance. **Vision Processing Units** (**VPUs**), like Intel's Movidius Myriad series, focus on low-power CNN inference. However, compared with the Myriad 2 VPU [132], FPGAs hold a key advantage in radiation tolerance. In a direct comparison, Rapuano et al. [98] build on the CloudScout initiative [36] and conclude that the Myriad 2 VPU is suitable only for short **Low Earth Orbit** (**LEO**) missions due to radiation susceptibility—unlike the space-grade Kintex Ultrascale XQRKU060 they also evaluate and is SEL-immune. Performance-wise, their design shows a 2.4× speed-up compared with the VPU, but with 1.8× higher power consumption and a longer development time.

A Large Catalog of Computing Platforms. Through the motivations of the surveyed studies, and in line with the conclusions of Wang et al. [121] and Lentaris et al. [69], a generic guideline for hardware platform selection, considering development effort and performance targets, is reported in Table 3. GPUs, TPUs, and VPUs offer quicker paths to achieving initial performance, especially for computationally complex DNNs [9]. FPGAs, however, promise superior performance and efficiency if more development effort and hardware expertise are invested in custom accelerator design and optimization. For applications demanding ultimate performance and willing to bear the highest development costs and longest design cycles, custom ASICs represent the pinnacle, albeit with reduced updatability [4]. In addition, both ASICs and FPGAs can directly connect their compute units to input and output ports, enabling lower latencies compared with GPUs and CPUs. In the following sections, we examine the technical aspects of the surveyed studies, exploring the methodologies and techniques employed to implement and optimize ML models on FPGAs for RS applications.

[12]Allowing post-launch updates and in-orbit reconfiguration.

RQ3 TAKEAWAYS

AMD FPGAs overwhelmingly dominate the surveyed hardware, comprising nearly all devices (96%). The preferred FPGAs are SoC boards, particularly the Zynq family (57%). Most boards are equipped to run DNNs (60% with > 500 DSPs), and almost all devices have accelerated at least one DL model in an experiment. Reconfigurability, power efficiency, and radiation concerns are the primary motivations for choosing FPGAs over other DNN-friendly COTS accelerators.

4 From AI Models to FPGA Deployment: Design Strategies and Challenges

4.1 Optimizing AI Models for Efficient Inference (RQ4)

Before deploying a model on an FPGA, it typically requires adaptation and compression. These optimization steps are closely linked to hardware design decisions—discussed later in RQ5 and RQ6. In RQ4, we introduce key network compression techniques and discuss their impact and requirements. Since traditional ML methods are already lightweight, we primarily discuss optimizations for computationally intensive models, like NNs.

Architectural Optimizations via Lightweight Backbones. Most NNs include a feature extractor, or backbone, responsible for refining and compressing the input's essential information. Extracting rich, high-quality features often relies on large and deep backbones, which come with high computational costs. To lighten these costs, designers can exploit lightweight backbones. For instance, Yang et al. [133] replace YOLOv3's DarkNet53 backbone with GhostNet [45], reducing the number of learnable parameters by a factor of 3. Similarly, Li et al. [70] swap YOLOv4's CSPDarknet53 backbone for MobileNeXt [144], lowering the parameter count from 60.08M to 36.44M, with only a 1.65% mAP drop—compensated via other customizations. As discussed in RQ2, the existence of many architectures make selection a challenging task. Sabogal and George [104] compare four DL segmentation models, showing that ENet [93] outperforms U-Net [102] while being 23.81× more computationally efficient and requiring 20.56× less memory. Similarly, Yang et al. [132] compare three popular lightweight backbones: MobileNetv1 [53], MobileNetv2 [106], and SqueezeNet [55], concluding that MobileNetv1 achieves the best trade-off for their SAR ship detection application. MobileNetv1 also introduces a width multiplier for layer-wise scaling, leveraged by Suh et al. [112] to explore different deployment scenarios in terms of accuracy and FPGA resource use. Beyond swapping architectures, Kim et al. [62] propose a multi-phase cloud coverage detection pipeline where images first undergo a uniformity check and classification via a small CNN (TriCloudnet) before full segmentation with U-Net. Such workflow avoids the computation of cloud-clear or overcast images, reducing processing time and power consumption by 48.7%. When sacrificing some feature representation capabilities is manageable, lightweight and multi-phase approaches highlight that questioning a model's depth and width can result in significant computational savings.

Reducing the Complexity of Convolutional Layers. Given the ubiquity of convolution operations in Computer Vision, lightweight backbones commonly focus on efficiently computing them. For instance, MobileNets architectures [52, 53, 106] replace standard convolution layers by depth-wise separable convolutions. This form of factorized convolutions separate the standard convolution operation in two stages: filtering, with single 2D filters for each input channel, and combination, by linearly combining the output of the filtering with a 1x1 convolution. When considering an input tensor $h_i \cdot w_i \cdot d_i$, a kernel $k \cdot k \cdot d_i \cdot d_j$, and an output tensor $h_i \cdot w_i \cdot d_j$, a standard convolution costs $h_i \cdot w_i \cdot d_i \cdot d_j \cdot k \cdot k$, while a depth-wise separable convolution costs $h_i \cdot w_i \cdot d_i(k^2 \cdot d_j)$, a computational

reduction of a factor $\approx k^2$. For example, Howard et al. [53] use depth-wise separable convolution to reduce computation by 9× for a loss of 1% on ImageNet. As an alternative path to convolution approximation, All Adder Networks [17] replace the multiplications in convolution kernels by additions. While it does not reduce computational complexity, additions are generally less expensive than multiplications. In particular, the All Adder NN of Zhang et al. [139] achieves 3 TOP/s, the highest computational throughput of the survey.

Pruning Techniques for Compact Networks. Pruning reduces computational complexity and memory costs by removing non-essential weights or entire channels from an NN. While the usual overparameterization of NNs aids training, it results in many redundant weights [113]. After training, pruning nullifies the least significant weights (in contribution to the output) to reduce this redundancy, but often requires fine-tuning to recover lost accuracy. The resulting tradeoff between parameter reduction and task accuracy can be adjusted with a pruning factor, which determines the percentage of weights to be nullified. To further enhance pruning efficiency, sparsity is often enforced during training via techniques like L1-regularization [137, 143]. For example, Zhang et al. [137] train their GNN with lasso regression—adding an L1-penalty term to the loss—to encourage weight shrinkage and facilitate pruning. Channel pruning removes entire feature channels, enabling more efficient hardware execution. Yang et al. [133] prune Ghost-YOLO channels, reducing the number of parameters from 23.52 M to 6.71 M at a cost of 2.93 % mAP.[13] Similarly, Nguyen et al. [86] prune 60 % of filter weights, incurring only a 0.85 % mAP loss. Beyond weights and channels, certain families of models also have the opportunity to prune input data. For example, Zhang et al. [137] reduce the computational complexity of their GNN by 92.8 % by pruning input graph vertices, causing only a 0.17 % OA drop. To leverage the efficiency improvements reached through pruning, Yang et al. [132] perform hardware-guided progressive pruning. The method integrates an initial coarse pruning, a resource-per-layer cost estimation, and a final resource-aware fine pruning. Pruning is frequently paired with quantization for further efficiency gains [46, 143].

Custom Datatypes Through Quantization. One of FPGA's key computational advantages over alternative hardware platforms, like GPUs, lies in custom datatype operations. Quantization reduces numerical precision by converting full-precision 32-bit floating-point (*fp32*) parameters into lower-precision formats such as 16-bit fixed-point (*i16*), 8-bit integer (*i8*), or even binary (*b*). Low-bandwidth datatypes enable highly efficient arithmetic operations, particularly on FPGA logic, where binary multiplications can be performed using simple XOR operations. While quantization reduces computational cost and memory footprint, it inevitably perturbs model parameters,[14] shifting them from their convergence point and causing accuracy degradation compared with full-precision implementations [34]. For instance, Sabogal and George [104] report an average accuracy drop of 1% across four segmentation models. Similarly, in object detection, Zhao et al. [143] observe a 0.42% drop in mAP, while Ni et al. [88] report a minimal reduction in classification accuracy of 0.03% for VGG16 and 0.06% for ResNet-34. Although the quantization-induced loss rarely exceeds 1%, it can be unacceptable in scenarios requiring high model reliability. To mitigate accuracy degradation, designers may opt to re-train the network, effectively choosing one of two standard approaches: **Post-Training Quantization** (**PTQ**) or **Quantization-Aware Training** (**QAT**). PTQ, sometimes called direct quantization, is straightforward to implement as it applies quantization directly after training. However, PTQ's simplicity can cause significant accuracy degradation, especially in

[13]High accuracy costs are usually mitigated with other optimizations, in this case via Knowledge Distillation (KD) [50], discussed further in Section 6.

[14]Several factors influence the effectiveness of a quantization scheme, such as its uniformity, symmetry, and granularity. For a more detailed discussion on quantization techniques, we recommend the comprehensive survey of Gholami et al. [34].

models with a wide parameter distribution [58]. For instance, Wei et al. [125] observe over 50% accuracy loss across all tested bit widths when applying PTQ to a LeNet-5 for MSTAR classification. In contrast, QAT simulates low-precision arithmetic already during training, resulting in better robustness to quantization noise [95], albeit at a higher computational cost [136]. Beyond PTQ and QAT, mixed-precision quantization assigns different bit widths to different sections of the model. This method preserves accuracy by keeping sensitive layers at higher precision while applying aggressive quantization to less critical layers. For example, Yang et al. [132] identify inefficient layers by estimating layer-wise latency and resource consumption. Rapuano et al. [98] take this idea further by manually assigning custom bit widths to each input, filter, and output, achieving a 48% memory footprint reduction for CloudScout. Alternatively, some approaches like Zhang et al. [141] mitigate accuracy loss by selectively retaining *fp32* precision for activations while quantizing only the weights. Similarly, Wei et al. [125] introduce a symmetric QAT method that quantizes weights but preserves normalization and activation in *fp32* format. The resulting computation overhead is motivated by the need to preserve the model's accuracy. Quantization is indispensable to achieve FPGA designs that are competitive with higher-frequency computing platforms. While PTQ is often sufficient to achieve 8-bit quantization with minimal accuracy degradation [84], QAT is preferable when maintaining accuracy is critical. Mixed-precision quantization further optimizes efficiency by balancing accuracy and computational cost.

Combining Optimization Techniques. All the methods discussed above effectively optimize NNs, but their combined use often yields the best results. Notably, combining lightweight backbones and fixed-point quantization is a recurring and effective strategy in the surveyed studies [62, 70, 74, 85, 86, 132, 133, 136, 143]. Additionally, lightweight backbone replacement or the direct use of efficient architectures, such as ENet [93] or *tiny* variants like YOLOv4-tiny [86, 136], appear in ∼20% of the surveyed experiments, offering an effective trade-off between accuracy and complexity. When the accuracy drop must be compensated, designers typically enhance the baseline architecture or training setup. For example, Nguyen et al. [86] add a third YOLOv4 head to detect targets at multiple spatial scales and optimize training with data augmentation to recover accuracy losses. Several surveyed studies [46, 86, 143] combine backbone replacement, pruning, quantization, and a recovery method in an incremental optimization strategy. Such gradual development approaches allow designers to evaluate the contribution of each step in an ablative manner, progressively refining both accuracy and resource efficiency before final FPGA deployment. Ultimately, the effectiveness of these optimization techniques depends not only on compression gains but also on their alignment with FPGA accelerator design, discussed in the next section.

RQ4 TAKEAWAYS

To reduce computational load, 24% of DNNs leverage lightweight backbones or efficient architectures, primarily MobileNets (46%). Out of all DL experiments, 89% use quantization, mostly to *int*8 format, and 26% leverage pruning to compress their model. Fine-granularity mixed-precision quantization presents significant opportunities for further network compression.

4.2 Finding the Best Framework to Efficiently Use FPGAs (RQ5)

Implementing NNs on CPUs and GPUs typically involves using stable and mature software frameworks and drivers. In contrast, automatic deployment frameworks for FPGAs have not seen the same level of adoption—most implementations depend on manual frameworks and varying design patterns. To support future work in selecting suitable implementation frameworks and design

patterns, we provide a dedicated second taxonomy in Table 4. It presents key metrics relevant to the respective FPGA implementation, as reported in the original study, organized into three column groups. First, *Implementation Choices* includes the implementation framework *Impl.*, the model family *Fam.*, the design pattern *P* (*S* for Specific, *F* for Flexible), the *FPGA* product, and the *Model Name*. Second, *Design Metrics* covers the precision post quantization *Prec.*, the model complexity *C**[OP]***, and the memory footprint ***[MB]*** (†: model parameters are located on-chip, and not loaded from external memory). Additionally, we report the DSP *D**[%]*** and BRAM *B**[%]*** resource utilization. Finally, *Performance Metrics* attempts to compare different designs with each other. In particular, we report the operating frequency in ***[MHz]***, computational throughput *T**[GOP/s]***, power consumption *P**[W]***, computational efficiency *T/P*, Latency *L**[s]***, and temporal throughput in ***[FPS]***.

Implementation Choices. The designs are split between manual implementations—**Hardware Description Language** (**HDL**) and **High-Level Synthesis** (**HLS**)—and automatic frameworks—FINN [117], Vitis AI [5], MATLAB, and VGT. HDL is the most commonly used language (≈ 45%), offering precise control of the generated circuit. However, most designs do not exploit this level of control to use all available resources nor push the theoretical limit of the clock frequency. HLS tools offer a more abstract approach to FPGA programming as they enable the use of high-level languages to create larger designs with less lines of code. Even with HLS, a solid grasp of FPGA architecture and hardware concepts remains necessary, especially when compared with GPU programming environments, like CUDA. In the reported designs, we detect no overarching trend in worse resource utilization or clock frequency of HLS designs compared with HDL designs. The average resource utilization of implementations are ≈30% DSPs and 50% BRAM for HDL and 46% DSPs and 43% for HLS with a high variance between individual designs. Most manual designs focus on CNN-based networks (≈ 70%), a well-studied network architecture commonly implemented using an array of **Processing Elements** (**PEs**) that pass intermediate results between each other.

To reduce the overhead of porting CNNs to FPGAs automatic frameworks have been developed. The two most commonly used in this survey are AMDs Vitis AI and the FINN compiler. The FINN compiler [117] automatically generates a model *Specific* accelerator for CNNs. Once implemented, the FINN design is solely useful for the specific network. It heavily depends on quantization to reduce the network complexity and memory footprint. Myojin et al. [83] binarize a 4 layers CNN, and Pitonak et al. [95] compare 2,3 and 4-bit quantization using FINN. The Vitis AI compiler [5] supports the *Flexible* design approach, using a **Deep Learning Processor Unit** (**DPU**) flashed to the FPGA able to run multiple different AI workloads. The peak **Operations Per Cycle** (**OPC**) of the DPU can be customized and matched to the algorithm [5]. Sabogal and George [104] test multiple networks with different DPU configurations, demonstrating that some networks have utilization as low as 14.2% (ESPNet) and up to 85.6% (U-Net) using the B1024 (1024 OPC) architecture. Vitis AI can leverage multiple DPU cores to increase performance as demonstrated by Nguyen et al. [86] and Yu et al. [136]. On the downside, Vitis AI requires PTQ or QAT, as it only supports the *i8* datatype [5]. Other automated frameworks include MATLAB (2023a) [31] and the **Xilinx System Generator** (**XSG**) [42], but both works do not go into detail concerning their implementation. Bahl et al. [8] use VGT, a VHDL compiler generating circuits for abstract CNNs similar to FINN; it has, however, no continuing support. Frameworks that are not classified do not provide enough information on their implementation.

Design Pattern. We split each workload into two distinct design patterns. *Flexible (F)* designs are bitstreams that support multiple different kernels and network topologies without requiring reconfiguration. Supporting different networks and kernels usually requires weights to be loaded from off-chip memory. Zhang et al. [137] use a *Flexible* design with on-chip memory, pre-loading

Table 4. FPGA Optimization Taxonomy Table

Implementation Choices						Design Metrics					Peformance Metrics					
Impl.	Fam.	P	Ref.	FPGA	Model Name	Prec.	C[OP]	MB	D[%]	B[%]	MHz	T[GOP/s]	P[W]	T/P	L[s]	FPS
Manual	HDL	CNN	S	[56]	7Z045	HO-ShipNet	i16	–	–	44	44	–	–	1.90	–	1.01 ms	–
				[98]	ZU7EV	CloudScout	i?	–	13.30	67	46	115	–	3.40	–	141.7 ms	–
				[119]	7A35T	CNN	f32*	–	-†	–	–	36	1.1	0.116	9.8	–	–
				[133]	7Z020	Ghost-YOLOS	i16	8.41G	12.60	58	60	150	29.5	2.98	9.9	320 ms	–
				[141]	7VX690T	Q-IORN	i8	–	121.51	21	27	200	209.6	6.32	33.2	147.62 ms	6.77
				[139]	7VX690T	A2NN	i8,f32	14.9G	-†	9	86	200	3047	8.27	368.4	4.92 ms	203.25
			F	[49]	ZU15EG	2D CNN	i8	684 K	1.20	5	71	100	7.1	8.40	0.8	0.097 ms*	–
					ZU15EG	3D CNN	i8	4.23M	0.12	9	76	100	3.8	8.40	0.5	1.11 ms*	–
					ZU15EG	HybridSN	i8	101 M	20.50	37	80	100	13.2	8.40	1.6	7.71 ms*	–
				[74]	Z100	CBFF-SSD	i16	–	–	57	60	200	452.8	19.52	23.2	42.59 ms	–
				[75]	ZU3EG	AP2D-Net	i(mix)	–	–	80	75	300	130.2	5.59	23.3	32.8 ms	30.53
				[88]	7VX690T	Improved YOLOv2	i8,i32	379 G	49.40	23	54	200	387	14.97	25.9	981.4 ms	–
					7VX690T	ResNet-34	i8,i32	7.33G	21.29	23	54	200	182	14.97	12.2	40.2 ms	–
					7VX690T	VGG16	i8,i32	30.6G	14.70	23	54	200	344	14.97	23.0	89.1 ms	–
				[125]	7K325T	LeNet-5 f32	f32	–	6.64	34	33	100	–	–	–	2.29 ms	–
					7K325T	LeNet-5 i8	i8	–	1.66	14	16	100	–	–	–	2.29 ms	–
				[131]	7A200T	Improved VGG16	i8	40.9G	14.80	12	29	200	23.1	3.41	6.8	1780 ms	–
					7A200T	Improved YOLOv2	i8	379 G	49.40	12	29	200	22.2	3.41	6.5	17120 ms	–
				[140]	7Z035	Improved YOLOv2	i8,f32	379 G	–	21	74	200	111.5	5.96	18.7	3400 ms	–
		ML	S	[28]	7Z020	Decision Tree	i8	–	-†	10	86	–	–	3.50	–	300 ms*	–
				[79]	7Z020	SVM	f32*	–	-†	38	–	200	–	–	–	0.11 ms*	–
				[81]	3SD1800A	MLP	i8,i32	–	-†	26	–	–	–	–	–	–	–
				[99]	6VLX240T	RDBC	–	–	-†	–	–	230	–	0.12	–	28.9 ms	–
				[109]	7VX690T	LPDBL	f32*	–	-†	14	15	200	–	–	–	–	–
				[130]	6VLX240T	BRAM_DSP	i8	–	-†	91	52	192	–	4.8	–	21.28 us	–
					6VLX240T	LUT_MUL	i8	–	-†	0	0	189	–	4.5	–	21.28 us	–
		GNN	S	[16]	7VX690T	SAM-GNN	f32*	–	–	13	14	200	–	–	–	–	–
	HLS	CNN	S	[62]	7Z020	TriCloudNet + U-Net	–	–	3.61	49	60	–	–	0.196	–	1028 ms	0.97
				[122]	7Z020	CNN	b	–	-†	11	100	–	69.8	2.38	29.3	–	7
				[132]	7VX690T	CNN2@0.7	i?	450 M	0.09†	64	21	250	298	5.80	51.4	0.7 ms	652
					7VX690T	CNN4@0.7	i?	560 M	0.11†	69	21	250	366	6.60	55.5	0.7 ms	636
					7VX690T	CNN6@1.6	i?	600 M	0.09†	76	18	250	235	5.10	46.1	1.6 ms	384
			F	[127]	ZU9EG	Improved YOLOv3	i8	–	8.50	–	–	–	–	–	–	7 ms	–
			–	[85]	KU040	AlexNetLite	i16	5.71G	-?	31	39	200	–	–	–	–	–
					KU040	MobileNetv1Lite	i16	50.0M	-?	31	39	200	–	–	–	–	–
		ML	S	[13]	ZU7EV	Deep Belief Network	f32	–	-†	10	0	250	–	–	–	5.4 ms	–
		GNN	F	[138]	U280	GNN	f32*	–	–	43	37	260	400	38.00	10.5	27 ms	759
				[137]	ZU7EV	GNN	f32	2.13M	0.96†	74	91	125	–	6.30	–	0.105 ms	9500
		ViT	S	[126]	U250	VTR	–	18.2G	-?	–	–	300	–	–	–	–	–
Automatic	FINN	CNN	S	[47]	7Z020	CNN	b	–	-†	–	–	–	–	–	–	330 ms	–
				[83]	7Z045	BNN	b	–	-†	–	–	–	–	–	–	21.14 ms	–
				[95]	7Z020	CloudSatNet-1 Q2	b	–	1.43†	14	–	100	–	2.52	–	–	15.46
					7Z020	CloudSatNet-1 Q4	i4	–	3.06†	14	–	100	–	2.59	–	–	15.47
	Vitis AI	CNN	F	[54]	ZU19EG	ETAUS	i8	–	–	–	–	106	329.9	1.65	199.9	107.8 ms	9.3
				[70]	ZU7EV	RFA-YOLO	i8	–	24.75	–	–	200	–	15.82	–	–	27.97
				[86]	ZU7EV	YOLOv4	i8	105 G	245	–	–	100	–	30.1	–	107.94 ms	17.9
					ZU7EV	YOLOv4-tiny 3L	i8	15.3G	24.10	–	–	100	–	26.40	–	19.18 ms	125
				[104]	ZU3EG	ENet	i8	4.06G	0.36	70	75	300	–	3.36	–	–	25.2
					ZU3EG	ESPNet	i8	3.71G	0.33	70	75	300	–	3.08	–	–	11.7
					ZU3EG	FPN	i8	17.3G	5.84	70	75	300	–	4.03	–	–	14.1
					ZU3EG	U-Net	i8	96.6G	7.40	70	75	300	–	3.38	–	–	2.7
				[118]	ZU9EG	ResNet-50	i8	–	–	–	–	–	–	–	–	10.7 ms	93.02
					ZU9EG	SICNet	i8	–	–	–	–	–	–	–	–	1.8 ms	530.66
				[136]	ZU7EV	Improved YOLOv4-tiny	i8	–	–	–	–	300	–	20.00	–	–	295.9
				[143]	K26	YOLOv4-MobileNetv3	i8	–	5.69	59	–	200	–	7.20	–	–	48.14
	MATLAB	CNN	S	[31]	ZU9EG	ResNet-18+YOLOv2	i8	–	–	–	–	250	–	–	–	–	46
		ML	S	[42]	7A35T	CAG-SC2S	i32	–	-?	93	21	70	–	0.31	–	5.27 us*	–
					7A35T	CAL-SC2S	i32	–	-?	96	25	65	–	0.36	–	5.70 us*	–
					7A35T	SVM	i32	–	-?	49	6	72	–	0.22	–	5.12 us*	–
	VGT	CNN	S	[8]	6	C-FCN++	i16	4.70K	0.047†	–	–	100	–	–	–	150 ms	7
Unclassified	–	CNN	S	[18]	7K325T	LeNet-5	f32	–	–	–	33	100	–	–	–	2.29 ms	–
			F	[112]	ZU3EG	SSD 0.25x	i8	4.04G	–	73	47	200	138	2.40	57.5	–	34.18
					ZU3EG	SSD 0.5x	i8	15.6G	–	73	55	200	150	2.60	57.7	–	9.6
					ZU3EG	SSD 1.0x	i8	61.6G	–	73	85	200	158	2.00	79.0	–	2.56
			–	[123]	7Z020	Lightweight CNN	i8	4.60G	–	70	25	100	–	3.30	–	–	60.0
		ML	S	[43]	7Z020	MLP	f32*	–	-†	26	1	–	–	0.13	–	–	–
				[44]	7Z020	TinyML	i8,i16	–	0.000037†	49	–	100	–	0.173	–	–	–
			–	[116]	7Z020	WNS	f32*	–	-†	–	–	–	–	–	–	–	–

FPGA: AMD/XILINX FPGA names are stripped of leading to increase readability
Mem: † Uses On-Chip Memory, ? Memory Location unknown
Latency: * Latency of a single PIXEL
Model Name: **RDBC** Roller Dung Beetle Clustering, **WNS** Weightless Neural System.

memory to the weight buffer, and amortizing the access over multiple classifications. The support for different layers comes at a small cost in computational performance, but allows implementing larger models that would otherwise not fit on the FPGA using a *Specific* approach. The average footprint of the 20 *Flexible* designs is $25MB$ compared with $14MB$ for the 11 *Specific* designs. Unlike *Flexible* designs, *Specific* designs only support a single network architecture. Deploying a different model requires re-implementation of the bitstream or is not supported. This allows the design to be hyper-specialized to the model, making full use of the customizability of the FPGA. *Specific* designs can use different datapaths for different pruning and quantization of individual layers. Differentiating model families (*Fam.*) is relevant as different ML models have vastly distinct requirements. For example, CNNs spend most of their time performing convolution, a task with high computational intensity [98], whereas GNNs require a high memory bandwidth [138]. The FPGA device also influences the performance, similar to CPUs, the number of transistors increases over the generations, enabling designers to use more resources such as DSPs, BRAM, and LUTs (see 3.4).

Design Metrics. The *Prec.* column reports the precision of operations post-quantization. The most popular quantization is *i8*, with 50% of all implementations and 43% of manual ones using either *i8* or a mixture of *i8* with *f32* or *i32*. Next, we report the model complexity ***C[OP]***, an essential metric as a higher complexity requires more computational resources to produce the same temporal throughput *FPS*. In addition, the relationship between model complexity and memory footprint can indicate whether a model is compute- or memory-bound. Most of the designs do not reach the compute boundary nor provide an off-chip memory bandwidth analysis making a further discussion impossible. The next metric is the model footprint and the location of parameters. FPGAs, similarly to CPUs, contain restricted on-chip memory, e.g., $<$ 5MB for Zynq MPSoC chips, and a main off-chip memory (512MB to 32GB). Using on-chip memory eliminates most memory bandwidth concerns as distributed BRAM provides very high bandwidth. On-chip memory weights can either be stored in BRAM—pre-loaded before execution—or stored inside LUTs—disabling any change in parameters. *Specific* designs mostly use on-chip memory with only 24% using off-chip memory, while 93% of *Flexible* designs use off-chip memory. In particular, the FINN compiler stores weights on-chip, using BRAM and LUTs, while Vitis AI loads weights from external memory. Storing the weights on-chip leads to the design being mostly limited by BRAM [28, 122, 139]. The average memory footprint of on-chip implementation is only $0.72MB$ compared with $28MB$ for off-chip memory implementation. The last two middle columns present the DSP & BRAM utilization. A high utilization can indicate a good selection of the model for the FPGA, while a low utilization can indicate bad scalability of the design. This excludes network architectures that cannot use DSPs and BRAMs efficiently, such as BNNs—using LUTs for most operations—or GNNs—with low data locality and computational intensity. Yahiaoui et al. [130] compare an implementation using only LUTs with a DSP- and BRAM-based implementation. They show that the BRAM_DSP implementation can double the number of pixel processed in parallel, while using the same FPGA.

Performance Metrics. Comparing different models on CPUs and GPUs is already difficult; for FPGAs, where the design can vary even more, finding a fair comparison is especially complex. One common metric relevant for comparison is the computational throughput, marking the peak number of operations per second. Nevertheless, operations with different precisions cannot be directly compared, as floating-point computations are more expensive than integer operations, and smaller bit widths yield higher **[GOP/s]**. For example, the highest throughput is reported as 3 TOP/s [139]. However, the solution uses an All Adder NN, which approximates the functionality of a convolution filter by using only additions. Moreover, the intermediate feature maps resulting from the convolution layers are sent back to the CPU. Thus, the **Fully-Connected** (**FC**) layers,

corresponding to 92% of the parameters, are executed on the CPU, drastically reducing the memory footprint. The remaining layers only represent $10M$ parameters, which can fit on the boards 6.6MB BRAM due to the quantization to 4-bit. Using BRAM eliminates any memory bottlenecks and the main computation, a 4-bit adder, can be implemented using solely LUTs, resulting in the extreme throughput. The second highest reported throughput is 452 GOP/s [74]. Compared with Zhang et al. [139] the design uses 16-bit values loaded from memory, and the main operations are **Multiply-Accumulate** (**MAC**), counted as two operations. The DSP block inside a 7-series Xilinx FPGA can perform one MAC operation per cycle and they manage to keep 1152 DSPs (60%) busy achieving $1.97 \times 1152 \times 200MHz \approx 452GOP/s$. Sabogal and George [104] and Suh et al. [112] use double data-rate DSPs to increase the performance—an optional optimization inside Vitis AI. Doubling the clock frequency of DSP components, which can be clocked higher than the rest of the design, roughly doubles the available computational capability. A majority of experiments ($\approx$ 72%) report the power consumed by the FPGA, a value pointless on its own as the total energy required to compute a sample is more important. To attempt a more informative comparison between different models and implementations, we calculate the computational efficiency T/P in $[\frac{GOP/s}{W}]$. Zhang et al. [139] achieve the highest computational efficiency $368\frac{GOP/s}{W}$ due to the high throughput detailed previously. On the other hand, He et al. [49] have an extremely low $< 2\ \frac{GOP/s}{W}$ efficiency, even though they use a low-precision datatype (*i8*). Another key design metric is the frequency *MHz*, which depends on the speed-grade and technology of the device, as well as the critical paths in the design. A low frequency, compared with designs on the same FPGA, can hint at a poorly designed circuit with long critical paths. *Latency* and *FPS* values are model- and implementation-specific, incomparable between different workloads and implementations, and only included to ensure completeness. It is important to separate *Latency* from *FPS*, as, generally, multiple inputs can be processed inside the accelerator pipeline, increasing the FPS significantly compared with a non-pipelined implementation that can process one element at a time.

RQ5 TAKEAWAYS

Extensive work in implementing CNNs on FPGAs has produced various automated frameworks supporting common layers. Within the surveyed studies we could not find a trend showing better performance of manual designs in any of the metrics. *Specific* design patterns are commonly used with limited on-chip memory ($\sim$75%), while *Flexible* design patterns tend to support larger models (+10MB in average).

4.3 Optimizing Performance for FPGA Implementations (RQ6)

FPGAs have a reconfigurable architecture, which enables designers to implement energy-efficient dataflow architectures that re-use data, and are adapted to the compute operations of the algorithm. However, FPGAs' reconfigurability comes at the cost of frequency—compared with CPUs and GPUs. Therefore, the use of different kinds of parallelism is necessary to achieve similar performance.

Instruction Level Parallelism (**ILP**) exploits the execution of consecutive instructions in parallel, for example overlapping load and MAC operations [112]. This level of parallelism is exploited in almost every design and can be performed manually in HDL or automatically using HLS control directives for loop statements. Consecutive iterations of a loop can overlap, increasing the throughput and amortizing the latency over multiple loop iterations. In addition, ILP designs commonly exploit **Data Level Parallelism** (**DLP**), for example by executing multiple MAC operations in parallel to perform convolution [131]. After optimizing individual tasks using ILP, designers can leverage **Task Level Parallelism** (**TLP**), for example, by overlapping different layers in a NN—a

mechanic heavily exploited by FINN [117]. The compiler connects individual layers with **First-In First-Out** (**FIFO**) streams and starts executing a layer as soon as its first inputs are produced by the previous layer. Increasing the batch size increases TLP, enabling different inputs to be processed at the same time. The disadvantage of larger batch sizes is a possible increase in latency, leading to most implementations using a batch size of 1 during inference.

Increasing computation puts pressure on the memory bandwidth, and most designs use multiple PEs in a Systolic Array [65] to reduce off-chip memory accesses. Furthermore, transformer-based networks [126], relying on dense matrix multiplication with a lower computational intensity than convolutional layers, require more off-chip bandwidth. Similarly, GNNs, exhibiting the lowest computational intensity, necessitate an even greater focus on creating an advanced network to reduce the number of accesses to off-chip memory as much as possible [137, 138]. Given enough bandwidth, multiple PEs enable simple scaling to use most of the FPGA resources; this approach can be found in many designs [16, 49, 74, 95, 98, 109, 112, 122, 125, 127, 130–133, 137–141]. Using multiple levels of parallelism, most designs can be scaled to reach the computational limit for floating-point operations on FPGAs, which require multiple DSPs per operation. To increase the **OP/s** available, designers can choose to reduce datatypes precision, hence reducing the number of resources required per operation. Most designs use quantization to a fixed-point datatype to increase the number of available operations. In addition to DSPs, which can perform basic arithmetic operations on integers, LUTs can perform any 6-bit to 1-bit mapping or a 5-bit to 2-bit operation. Thus, a BNN or low-precision network can be mapped to LUTs directly, allowing for high peak **OP/s** [139]. The tradeoff between LUT and DSP utilization is automatically explored by Yang et al. [132]. In addition to computational benefits, lower precision reduces off-chip memory requirements and can allow the network to fit entirely into BRAM. Another option to increase available operations is to run DSPs at double the frequency, effectively doubling their throughput [104, 112].

Finally, the large spatial scale of RS imagery also plays a decisive role in FPGA design, as it makes naive full-image processing impractical. Consequently, most implementations rely on patch-based strategies, which maintain throughput but introduce an additional pre-processing overhead—often negligible compared with models' complexity. In parallel, differences in sensing modalities also shape algorithmic design choices. For instance, HSI data, with hundreds of spectral bands, are frequently processed pixel-by-pixel [13, 16, 42, 79, 109], while others sample channels to decrease spectral redundancies and computational load [49]. In contrast, all studies handle SAR imagery as amplitude-only—effectively grayscale images—and therefore avoid any complex-valued operations or datatypes in the implementation.

RQ6 TAKEAWAYS

Most manual designs use a combination of different forms of parallelism and quantization to increase the throughput of the design. To reduce the off-chip memory bandwidth, designs can re-use data, commonly employing a Systolic Array of PEs. Comparing the quality of different designs proves difficult, and most surveyed studies do not compare their CNN implementations with previous work.

5 AI and FPGA Synergies: Towards Next-Gen Earth Observation Missions

In each experiment, designers must select an ML model, an FPGA device, and an accelerator architecture. In this vast search space, finding the best AI and FPGA technology combinations constitutes a tedious process. While the diversity of study setups makes any comparison challenging and generalization even more difficult, this section discusses promising associations between AI and FPGA technology.

5.1 Limits and Suitabilities of AI-Supported RS Applications on FPGAs (RQ7)

To demonstrate the feasibility of different ML models and RS applications for an increasing number of resources, we discuss trends among solutions based on the size of the FPGA. We use the amount of DSPs available, displayed in Figure 4, as a proxy for the size of the FPGA as other resources (LUTs, FFs, and BRAM) scale in a similar fashion.

Tiny and Small Devices. Among the four experiments using sub-150 DSP FPGAS, which we categorize as tiny, two designs deploy equally tiny models to perform pixel segmentation [42] and image classification [8]. For its part, Vitolo et al. [119] implement a larger network (7403 parameters) running at low frequency. All these designs achieve marginal power consumption (< 0.5W) and demonstrate the use case of tiny FPGA devices for traditional ML models and tiny CNNs for fast approximate classification. The next size of FPGAs are small FPGA devices, often featured in educational evaluation boards. Two FPGAs of this class are used: first, the XC7Z020 with 220 DSPs and DDR3, and second, the newer generation XCZU3EG with 360 DSPs and faster DDR4 memory. Both devices successfully deploy traditional ML or shallow NN designs, but also smaller CNN-based models are implemented to segment images [62, 104], classify ships [47] or clouds [95], or detect miscellaneous objects [75, 112, 122, 123, 133]. Deploying CNNs on small FPGAs requires the use of reduced models [95, 123, 133] or sacrifices in FPS [47, 62, 122, 133]. Interestingly, two studies [104, 112] examine the tradeoff between model size, throughput, and accuracy. Suh et al. [112] achieve particularly impressive results using their custom 8-bit UniPOT quantization scheme and leveraging DSP double rate, and perform inference $<$ 2.6W for all three SSD sizes.

Mid-Range FPGAs. FPGAs with more than 700 DSPs have sufficient resources for most classification or detection tasks using medium-sized models. However, to achieve high throughput, designers must still exploit significant model optimizations, such as heavy quantization, which leads to accuracy drops. For example, Rapuano et al. [98] deploy a $\sim$ 14MB custom CNN with layer-specific QAT on an XCZU7EV and achieve inference in 141.7 ms for 3.4W. Object detection tasks also induce a drop in accuracy, for example Zhang et al. [140] implement a lightweight version of YOLOv2, achieving 67% mAP on DOTAv1.0 [128], remaining significantly below the **State-Of-the-Art** (**SOTA**) (82%). Similarly, Zhao et al. [143] leverage PTQ, pruning, and depth-wise separable convolutions to detect objects in the DIOR [73] dataset at $\sim$ 50FPS and under 7.2W. Mid-range devices can still provide low-power implementations, such as the NN HO-ShipNet [56], consuming 1.90W on a Zynq 7045, or the ETAUS system [54], consuming 1.65W on a KV260.

High-End FPGAS. Only a few FPGAs incorporate 3,500 DSPs or more. Such available computing power enables deployment without significant alteration of the models and facilitates results competitive with the SOTA.[15] On the relatively old Virtex-7, Ni et al. [88] deploy large models for classification and object detection tasks. They achieve competitive accuracy on the NWPU-RESISC45 dataset [20] with both ResNet-34 (93%) and VGG16 (92%) deployed on their *Flexible* accelerator design compared with the SOTA (96%). As the DDR3 memory of the Virtex-7 can create a bottleneck on memory bandwidth, Yang et al. [132] stores model weights on-chip. Through pruning and mixed-precision QAT, they consistently compress YOLOv2 with different lightweight backbones to under 0.5MB while maintaining high accuracies. Their model optimizations, paired with optimized pipeline scheduling and the use of $\sim$70% of the available DSPs, lead to the lowest latencies of the surveyed DL models, achieving 0.7ms for their MobileNetv1 [53] and MobileNetv2 [106] designs. Two studies experiment with Alveo cards [126, 138] and demonstrate impressive

[15] Comparing SOTA-driven studies with HW-aware FPGA implementations is challenging without a cost/benefit analysis. Quantifying the computational cost of a 1% accuracy gain would clarify the trade-off between task performance and hardware efficiency, yet this is often missing in SOTA literature.

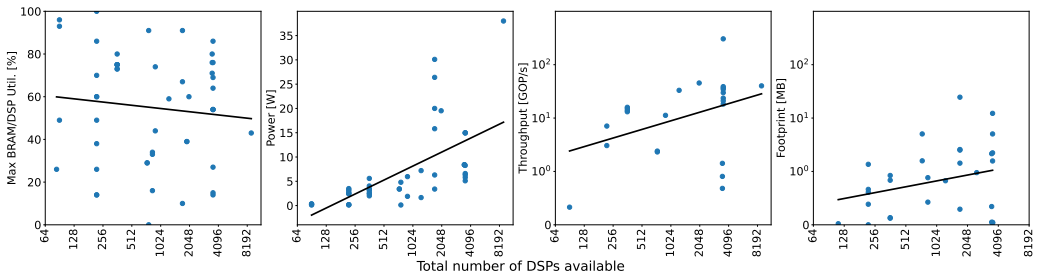

Fig. 5. Maximum utilization of DSP *D[%]* and BRAM *B[%]*, reported power *P[W]*, computational throughput ***[GOP/s]***, and footprint ***MB*** for different board sizes. We plot only reported metrics and fit the black line using linear regression.

throughput performance, including for the only transformer architecture of this survey. However, datacenter-grade boards, such as the Alveo family, have power consumptions ranging from 75 to 225W, making them unsuitable for edge operations.

Synergizing Model and FPGA. Several factors influence model selection, foremost the hardware constraints discussed above and in RQ3. Indeed, limited computational units (DSPs, LUTs) and on-chip storage (BRAMs) impose a tradeoff between accuracy and computational throughput. However, as seen in Figure 5, resource utilization varies significantly across all FPGA sizes. Optimizing utilization could make smaller FPGAs viable for a broader range of applications. Additionally, the characteristics of ML tasks and RS applications also restrict model selection. For example, object detection problems on the edge typically rely on single-shot models, such as YOLO or SSD, to save compute time. For segmentation, patch-based models (e.g., UNet) tend to yield better task accuracy [6] than methods that process pixels one at a time. Such models, however, have a larger memory footprint and higher computational costs. Classification, on the other hand, can be tackled with diverse models, ranging from SVMs that fit on all FPGA devices to larger models like ResNet50. As a final factor, because automatic toolchains only support a limited set of operations, designers who desire to accelerate development time and reduce the required design expertise must favor standard layers in their architectures. Ultimately, optimal model selection depends on each project's performance requirements, deployment constraints, and available development effort.

RQ7 TAKEAWAYS

Larger FPGAs naturally support higher model complexity and throughput, reducing the need for aggressive quantization or pruning. While real-time performance is frequently achieved, most models still require significant power (43% > 5W). High variance in resource utilization shows that efficient resource use, rather than raw device size, is a key determinant of achievable performance.

5.2 Advancing Onboard Processing: ML Approaches for Satellites and UAVs (RQ8)

One primary motivation for this survey was to understand how AI methods currently onboard spacecraft are deployed.

Onboard PhiSat-1: Pioneering Work in AI for RS. PhiSat-1[16] is a frequently cited mission [37] which deployed the first CNN model performing onboard cloud coverage classification: CloudScout [36].

[16] 6U CubeSat, launched 09/2020, https://www.esa.int/Applications/Observing_the_Earth/Ph-sat

ML-based methods accelerate data processing pipelines and recently became excellent at several essential tasks [114]. This is particularly true in RS, where the exponentially growing amount of data has outpaced the capacity for manual expert labeling. Deploying DL models directly onboard relieves downlink constraints and reduces the need for on-ground pre-processing. This saves time, conserves resources, and simplifies dataflows. For example, CloudScout classifies the cloud coverage of freshly acquired images, discarding overly cloudy images that contain little valuable information. Beyond hardware constraints, like limited memory and power, using a CNN for such a critical task implies a responsibility that comes with additional restrictions. As such, CloudScout has to ensure a False Positive rate under 1%, to avoid discarding appropriate images. In a follow-up study, Rapuano et al. [98] explore the possibility of using an FPGA instead of the original Myriad 2 VPU. They conclude that, although FPGA-based designs take longer to develop and consume more power, their superior customizability, lower latency, and radiation hardness make them suitable for long-term LEO missions. In a recent work, Cratere et al. [26] provide a comprehensive overview of the projects building on the CloudScout initiative, comparing all related cloud detection studies.

Onboard OPS-SAT: Experimental AI Deployment in Space. A second ESA mission is particularly relevant to onboard ML for RS: OPS-SAT.[17] This flying laboratory featured an Altera Cyclone V SoC and was designed for experimental ML deployments. The C-FCN++ model presented by Bahl et al. [8] was deployed onboard OPS-SAT. With 273 parameters (47KB), C-FCN++ is the only architecture explored in the study able to fit on-chip the Altera Cyclone V SoC. It uses dilated convolutions to expand early receptive fields and improve performance without increasing parameter count. Ultimately, C-FCN++ could perform cloud segmentation on a full-scale RS image in 150ms, less than required for the camera of OPS-SAT to acquire the next image, granting real-time performance. For further investigations of ML onboard OPS-SAT, see Kacker et al. [60], who tested several algorithms to prepare for the mission Beaver-Cube-2.[18]

Preparing Deep Learning for Spacecraft. Beyond deployed solutions, several studies contribute to preparing DL models for onboard spacecraft processing. In particular, Kim et al. [62] prepared for the A-HiREV mission[19] and present a cloud classification pipeline with goals similar to CloudScout. Their multi-phase approach uses three stages: a uniformity check, a lightweight ternary classifier called TriCloudNet, and a pruned U-Net. Each stage filters out extreme cases, reducing the workload for subsequent, more computationally complex models. This image prioritization method reduces downlink data by 40–50% and cuts processing time and power consumption by around 50%. Designed as a direct comparison to CloudScout, CloudSatNet-1 [95] classifies cloud coverage to relieve downlink, with a strong focus on low-power consumption and a minimized False Positive rate. More recently, Upadhyay et al. [118] develop SICNet, a cloud detection CNN optimized for real-time performance. The designers focus on deployment ease using Vitis AI and prepare further applications like fire and object detection. Zhang et al. [139] also target downlink reduction with A2NN, a model for direct onboard classification of RS images. As for solutions monitoring ships, Yang et al. [132] present OSCAR-RT, an end-to-end algorithm/hardware codesign framework dedicated to real-time onboard ship detection with SAR images. Their three experiments considerably focus on real-time performance using on-chip implementations. Neris et al. [85] use lightweight CNNs to identify ships and airplanes, deploying them on a Kintex US FPGA, which has a radiation-hardened counterpart, the XQRKU060—also used by Rapuano et al. [98]. Ieracitano et al. [56] also target ship monitoring with HO-ShipNet, a CNN enhanced with xAI techniques for increased

[17] 3U CubeSat, 12/2019 - 05/2024, https://www.esa.int/Enabling_Support/Operations/OPS-SAT
[18] 3U CubeSat educational mission from MIT, expected launch 2025, https://www.nanosats.eu/sat/beavercube-2
[19] 6U CubeSat platform from the Korea Aerospace Research Institute (KARI).

transparency, achieving 95% accuracy. As discussed in Section 3.4, radiation resilience remains a key requirement for space deployment [110]. Sabogal and George [104] study SEEs effects on DL applications, quantifying performance degradation in Xilinx DPU implementations [5]. Such effects are mitigated through circuit-level hardening or software solutions, like **Triple Modular Redundancy** (**TMR**). By contrast, UAV platforms operate in less radiation-exposed environments, enabling more flexible trade-offs in model design and hardware selection.

AI-Accelerated UAVs Payloads: Autonomous and Real-Time Insights. Spacecraft are not the only edge RS platforms benefiting from ML; numerous studies explore ML-supported pipelines for UAVs. Wang and Qiu [122] deploy a CNN on a Zynq-7000 to detect objects in images from the 2018 DAC System Design Contest, targeting real-time inference under low-power constraints. Suh et al. [112] design three sizes of SSD models to detect drones achieving high energy efficiency on a Zynq US+. Addressing a practical use case, Yu et al. [136] develop an Improved YOLOv4-tiny to detect abnormal railway track fasteners. Implemented with Vitis AI, their solution achieves ~300 FPS and 95.1% mAP, demonstrating performance levels suitable for onboard UAV deployment. Focusing on UAV navigation, Fraczek et al. [28] experiment with Decision Trees and Support Vector Machines to classify terrain and support automated landing of UAVs. Similarly, Wang et al. [123] merge optical and **MilliMeter-Wave** (**MMW**) radar data to quickly detect UAV obstacles. Their lightweight CNN reaches real-time performance with 60 FPS and low-power 3.3W on a Zynq-7020. By far the most complete surveyed UAV-focused study, Huang et al. [54] present the ETAUS system, a custom UAV based on the Pixhawk 4 drone to monitor Taiwan's **Air Quality Index** (**AQI**). ETAUS is powered by a CNN achieving high accuracy for AQI classification and cryptographic modules implemented on the FPGA logic. ETAUS also uses a YOLOv4 pre-trained from Vitis AI Model Zoo to detect—and later blur—private and sensitive information, such as license plates.[20] The study focuses on maintaining high accuracy for a reliable solution and achieving sufficient FPS for real-time performance.

Onboard Challenges: Real-Time, Memory, and Power Constraints. Real-time processing remains the central concern across the surveyed studies, as keeping pace with the continuous stream of newly acquired images is essential for operational success. To meet this demand, designers minimize model computations and maximize execution parallelism. In some contexts, such as the PhiSat-1 payload and CloudScout [36], strict memory footprint limitations introduce additional constraints. Low power consumption is also a recurring objective, though the diversity of mission requirements makes it difficult to extract consistent trends. For instance, Nguyen et al. [86] could afford to onboard a YOLOv4-tiny 3L with a 26.4W power budget on their UAV.

RQ8 TAKEAWAYS

Although only three studies explicitly target space missions, the widespread focus on onboard deployment suggests the broad applicability of the research to future onboard systems. Space missions are often geared towards surveillance, while UAVs address more localized applications.

6 The Missing Pieces: Gaps and Opportunities in FPGA-Enabled ML for RS

This section highlights key research gaps and opportunities of this emerging field. Following PRISMA 2020 guidelines [92], we first outline the limitations of our survey, then present promising study directions, and conclude with recommendations for future work.

[20] This section of the algorithm was not detailed in the study and is therefore not reported in this survey.

6.1 Survey Limitations

While Scopus[21] and **Web of Science (WoS)**[22] are the two most comprehensive public databases [15], we limited our search to WoS to maintain focus and feasibility; a preliminary analysis showed its journal coverage better aligns with the scope of this review. Besides, as with all systematic surveys, the search results inherently depend on the chosen query terms. This can lead to omissions, such as works deploying on FPGAs but using broader description like "resource-limited hardware" in their abstracts. Similarly, although we aimed for comprehensive ML-related queries, some lesser-known models may have been overlooked. While the heterogeneity in the surveyed literature contributes to a broad and extensive body of work, it complicates direct comparisons of experiments. Comparing FPGA implementations is particularly challenging, as different designs use distinct networks and devices. Furthermore, most CNN-based studies do not benchmark against prior FPGA-based CNNs implementations, limiting our ability to identify commonly compared architectures. Lastly, distinguishing between *Flexible* and *Specific* accelerators often relies on our interpretation of the described designs.

6.2 Research Gaps Analysis

In this section, we identify critical research gaps from the surveyed literature. To offer insights to guide future research directions, we structure the gaps into five subtopics.

Incomplete Spectrum of RS Use Cases. RQ1 shows that the surveyed applications do not cover the full spectrum of RS problems. While edge AI does not resonate well with unstressed applications that can be processed on the ground, like long-term climate monitoring, other onboard-relevant use cases remain unexplored, such as data compression and disaster response. Indeed, the NewSpace era enables real-time alerts for numerous emergency situations, such as wildfires. Similarly, local authorities greatly benefit from post-disaster damage assessments, e.g., after floods [80]. Exploring lightweight implementations of SOTA methods on FPGA could significantly advance both research and industry adoption. Concurrently, DL-based data compression pipelines have gained relevance over the past decade, including for RS data [38]. With SmallSats imaging payloads generating data at $\approx$ 640 Mbps [118], efficient onboard compression methods present a valuable research opportunity. In addition to these underexplored use cases, retrieval problems remain severely underrepresented, with only two studies focusing on feature estimation, despite these tasks being frequently essential for practical remote sensing outcomes.

Overlooked Prevalent DL Architectures. Although RQ2 highlights the vast diversity of encountered models, the surveyed literature overlooks several widespread model families. Indeed, almost all DL-based studies rely on CNN or GNN architectures, with a single example of transformer models [126]. This gap may partly reflect the recent emergence of **Vision Transformers (ViTs)**, which have not yet seen mature FPGA implementations. However, it may also highlight the relative implementation simplicity of graph flows and convolution pipelines on FPGA platforms. Nevertheless, ViTs have significantly advanced DL in RS research [114] and, despite their complexity, their superior performance makes them a promising target for edge deployment [105]. Similarly, no surveyed work explores recurrent models, despite their natural fit for temporal tasks—common in RS, such as monitoring change across time or along flight paths. Investigating the application of RNNs onboard flying platforms represents a valuable research direction. At the same time, the absence of support for transformer and recurrent architectures in FPGA toolchains (e.g., FINN, Vitis AI) poses a barrier to broader adoption. To unlock the potential of modern DL architectures in edge

[21] https://www.scopus.com/
[22] https://www.webofscience.com/

environments, automated frameworks must evolve to support these emerging layers, including their quantized and pruned variants.

Beyond FPGAs: Coarse-Grained Reconfigurable Arrays (CGRAs). While we discovered a wide range of FPGAs used for RS applications, we found limited use of CGRAs. Unlike FPGAs, CGRAs harden the commonly used arithmetic units, improving energy efficiency and allowing higher density. This reduces configurability, but offers enhanced performance and power efficiency. As discussed in Section 4.2, CNNs are often implemented with **Processing Elements** (**PEs**) organized in a systolic array. The AMD Versal[23] platform features a systolic array of small PEs (AI Engines) capable of basic vector operation. Perryman et al. [94] uses two Versal devices (AI Edge and AI Core) to implement semantic segmentation using distinct networks. Compared with Sabogal and George [104], which uses the same U-Net model and dataset [57], it increases the $\frac{FPS}{W}$ by 2.2x and 4.1x for the Edge and Core devices, respectively, demonstrating the potential of CGRAs. Given the availability of space-grade Versal devices, a detailed comparison between CGRA-based and FPGA-only solutions—particularly for highly quantized models—emerges as a concrete and promising research direction.

Further Refining NNs for Efficient Inference. RQ4 reveals that some widely used compression methods in Computer Vision remain underexplored in the surveyed records. In particular, **Knowledge Distillation** (**KD**) appears in only one study [133]. KD is a training strategy where a large "teacher" model with high learning capabilities transfers its learned representations to a smaller "student" model [50]. In the on-the-edge context, the student network can be tailored to the target FPGA, maximizing resource utilization and performance. KD can also guide weight pruning to improve compression while limiting accuracy loss [3]. This synergy holds strong potential for deployment in constrained environments. Although a few studies adopt hardware-aware design methods, hardware insights could benefit many optimization techniques. Notably, no surveyed work explores mixed-precision quantization driven by layer-wise hardware feedback, as in Yao et al. [135], who improve energy efficiency and latency by assigning optimal bitwidths per layer. Similarly, we see potential in using AutoML, particularly **Neural Architecture Search** (**NAS**), to identify DNN architectures tailored to hardware constraints. Approaches like EfficientNet [115] or MobiletNetv3 [52] are designed using platform-aware NAS, but hardware-aware NAS (HW-NAS) is still a young field with significant room for growth [10]. In this survey, only Hammoud et al. [44] used an approach similar to HW-NAS to determine the best model for their application. As GPUs increasingly integrate dedicated CNN units, such techniques become even more relevant for FPGAs, which must exploit their configurability to remain competitive. FPGA-aware NAS offers a promising path to design performant models well-matched to FPGA deployment, a research direction supported by datasets like HW-NAS-Bench [71] that estimate FPGA performance.

Trust in and Interpretability of Onboard AI. Another underexplored area in this survey is Uncertainty Quantification (UQ). UQ methods estimate the confidence of model predictions to improve interpretability and support more reliable decision-making onboard. Only Myojin et al. [83] apply UQ techniques, using **Monte Carlo** (**MC**) Dropout [29] to estimate model uncertainty through multiple stochastic forward passes. Other established methods remain unexploited in this context and deserve further investigation; for a comprehensive overview of UQ in Deep Learning, we refer readers to Gawlikowski et al. [32]. Practical benefits of UQ include uncertainty rejection, where predictions exceeding a confidence threshold are discarded, potentially easing downlink stress and increasing reliability. Similarly, **explainable AI** (**xAI**) is crucial for critical decision-making onboard, such as alert triggering or data prioritization. Yet, only Ieracitano et al. [56] employ xAI

[23] https://www.amd.com/en/products/adaptive-socs-and-fpgas/versal.html

tools to interpret model decisions. We believe that lightweight xAI and UQ methods can significantly increase the trustworthiness and adoption of AI solutions for onboard processing [51].

6.3 Recommendations for Future Work

We encourage researchers to use open-source data or to publish their datasets. When a full release is not possible due to confidentiality agreements, studies should detail, as rigorously as possible, sensor specifications, selected scenes, ground truth, data splits, and so on. Such transparency is crucial for future readers to fully comprehend the problem and replicate results. In addition, we strongly encourage future authors to share their code and scripts as thoroughly as possible, as this provides valuable information for the readers and increases the long-term impact of their work. Furthermore, a significant limitation across many reviewed studies is the lack of thorough metric reporting, which prevents comprehensive evaluation and comparison. Fundamental performance indicators are often incomplete. For example, some object detection works may report **mAP** but omit **mIoU**, preventing comparisons across studies. Similarly, hardware results may mention throughput and power, but use inconsistent bit-widths (e.g., *i4* vs. *fp32*), skewing energy efficiency comparisons. Even more critical, memory usage is rarely analyzed in depth. While some papers mention model footprint, very few examine off-chip memory transfers—a common bottleneck for FPGA designs [2, 23].

Future research should also build upon existing work in FPGA implementations. Despite the ubiquity of convolution operations in the surveyed applications, most HDL and HLS studies re-implement convolution kernels from scratch, often without an in-depth study or comparison to established designs. This is surprising given the rich body of prior work on optimized CNN accelerators for FPGAs [1, 2, 117, 120], as well as mature ASIC tape-outs from both academia [19] and industry [59]. Rather than reinventing the convolutional microarchitecture, new studies could build on these foundations and shift their focus toward application-specific challenges. Furthermore, we have observed that many studies highlight low FPGA resource utilization as a positive outcome; however, it may also indicate poor scalability. Ideally, FPGA resources should be fully utilized, except when limited by external factors like power budgets or memory bandwidth.

7 Conclusion

The integration of ML into RS pipelines is reshaping Earth Observation, enabling real-time, autonomous data processing onboard satellites and UAVs. Facing the growing demands of NewSpace and the need for efficient ML accelerators, FPGAs stand out as the standard payload for SmallSats due to their adaptability, potential for computational speedup, and energy efficiency. This article followed PRISMA 2020 guidelines to systematically review the literature on ML models implemented on FPGA-based platforms for RS applications. We proposed two taxonomies to classify existing research and answered eight research questions to map the literature landscape, unveil design paradigms, and discuss technological synergies. Our analysis highlighted the dominance of CNN-based architectures, the prevalence of AMD FPGAs, and the necessary synergy of quantization with FPGA programmable logic. Despite progress in the field, our findings reveal gaps in research, including limited diversity in RS tasks, underexplored model compression techniques, like knowledge distillation, and the scarcity of uncertainty quantification and explainability methods in FPGA-based ML. Addressing these gaps will be crucial for advancing reliable and efficient onboard AI systems. We hope this survey serves as a valuable resource for researchers and stimulates future advancements in FPGA-enabled Machine Learning for Remote Sensing.

Acknowledgements

This work benefited from the use of ChatGPT-4o, which provided alternative formulations of some parts of the manuscript, for example, section titles and the abstract.

References

[1] Mohamed S. Abdelfattah, David Han, Andrew Bitar, Roberto DiCecco, Shane O'Connell, Nitika Shanker, Joseph Chu, Ian Prins, Joshua Fender, Andrew C. Ling, et al. 2018. DLA: Compiler and FPGA overlay for neural network inference acceleration. In *2018 28th International Conference on Field Programmable Logic and Applications (FPL)*. IEEE, Dublin, Ireland, 411–4117. DOI: 10.1109/FPL.2018.00077

[2] Kamel Abdelouahab, Maxime Pelcat, Jocelyn Serot, and François Berry. 2018. Accelerating CNN Inference on FPGAs: A Survey. arXiv:1806.01683 [cs] DOI: 10.48550/arXiv.1806.01683

[3] Nima Aghli and Eraldo Ribeiro. 2021. Combining weight pruning and knowledge distillation for CNN compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3191–3198.

[4] Amara Amara, Frédéric Amiel, and Thomas Ea. 2006. FPGA vs. ASIC for low power applications. *Microelectronics Journal* 37, 8 (Aug. 2006), 669–677. DOI: 10.1016/j.mejo.2005.11.003

[5] AMD. 2019. AMD Vitis™AI Software. Retrieved from https://www.amd.com/en/products/software/vitis-ai.html

[6] Saeid Asgari Taghanaki, Kumar Abhishek, Joseph Paul Cohen, Julien Cohen-Adad, and Ghassan Hamarneh. 2021. Deep semantic segmentation of natural and medical images: A review. *Artificial Intelligence Review* 54, 1 (Jan. 2021), 137–178. DOI: 10.1007/s10462-020-09854-1

[7] Jose M. Badia, German Leon, Jose A. Belloch, Almudena Lindoso, Mario Garcia-Valderas, Yolanda Morilla, and Luis Entrena. 2022. Reliability evaluation of LU decomposition on GPU-accelerated system-on-chip under proton irradiation. *IEEE Transactions on Nuclear Science* 69, 7 (July 2022), 1467–1474. DOI: 10.1109/TNS.2022.3155820

[8] Gaetan Bahl, Lionel Daniel, Matthieu Moretti, and Florent Lafarge. 2019. Low-power neural networks for semantic segmentation of satellite images. In *2019 IEEE/CVF International Conference on Computer Vision Workshops (Iccvw) (IEEE International Conference on Computer Vision Workshops)*. 2469–2476. DOI: 10.1109/ICCVW.2019.00302

[9] Robert Bayer, Julian Priest, and Pinar Tözün. 2024. Reaching the edge of the edge: Image analysis in space. In *Proceedings of the Eighth Workshop on Data Management for End-to-End Machine Learning*. ACM, Santiago AA Chile, 29–38. DOI: 10.1145/3650203.3663330

[10] Hadjer Benmeziane, Kaoutar El Maghraoui, Hamza Ouarnoughi, Smail Niar, Martin Wistuba, and Naigang Wang. 2021. A Comprehensive Survey on Hardware-Aware Neural Architecture Search. arXiv:2101.09336 [cs]

[11] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. 2020. YOLOv4: Optimal Speed and Accuracy of Object Detection. arXiv:2004.10934 [cs] DOI: 10.48550/arXiv.2004.10934

[12] Mustapha Bouhali, Farid Shamani, Zine Elabadine Dahmane, Abdelkader Belaidi, and Jari Nurmi. 2017. FPGA applications in unmanned aerial vehicles - A review. In *Applied Reconfigurable Computing*. Springer International Publishing, Cham, 217–228. DOI: 10.1007/978-3-319-56258-2_19

[13] Samuel Boyle, Aksel Gunderson, and Milica Orlandic. 2023. High-Level FPGA design of deep learning hyperspectral anomaly detection. In *2023 Ieee Nordic Circuits and Systems Conference, Norcas*. DOI: 10.1109/NorCAS58970.2023.10305467

[14] James B. Campbell and Randolph H. Wynne. 2011. *Introduction to Remote Sensing, Fifth Edition*. Guilford Press.

[15] Arezoo Aghaei Chadegani, Hadi Salehi, Melor Md Yunus, Hadi Farhadi, Masood Fooladi, Maryam Farhadi, and Nader Ale Ebrahim. 2013. A comparison between two main academic literature collections: Web of science and scopus databases. *Asian Social Science* 9, 5 (April 2013), p18. DOI: 10.5539/ass.v9n5p18

[16] C. Chellaswamy, M. M. Manjula, B. Ramasubramanian, and A. Sriram. 2024. FPGA-based remote target classification in hyperspectral imaging using multi-graph neural network. *Microprocessors and Microsystems* 105 (March 2024). DOI: 10.1016/j.micpro.2024.105008

[17] Hanting Chen, Yunhe Wang, Chunjing Xu, Boxin Shi, Chao Xu, Qi Tian, and Chang Xu. 2020. AdderNet: Do we really need multiplications in deep learning?. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1468–1477.

[18] Lei Chen, Xin Wei, Wenchao Liu, He Chen, and Liang Chen. 2020. Hardware implementation of convolutional neural network-based remote sensing image classification method. In *Communications, Signal Processing, and Systems*, Qilian Liang, Xin Liu, Zhenyu Na, Wei Wang, Jiasong Mu, and Baoju Zhang (Eds.). Vol. 516. Springer Singapore, Singapore. DOI: 10.1007/978-981-13-6504-1_19

[19] Yu-Hsin Chen, Tien-Ju Yang, Joel Emer, and Vivienne Sze. 2019. Eyeriss v2: A flexible accelerator for emerging deep neural networks on mobile devices. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 9, 2 (June 2019), 292–308. DOI: 10.1109/JETCAS.2019.2910232

[20] Gong Cheng, Junwei Han, and Xiaoqiang Lu. 2017. Remote sensing image scene classification: Benchmark and state of the art. *Proc. IEEE* 105, 10 (Oct. 2017), 1865–1883. DOI: 10.1109/JPROC.2017.2675998

[21] Gong Cheng, Junwei Han, Peicheng Zhou, and Lei Guo. 2014. Multi-class geospatial object detection and geographic image classification based on collection of part detectors. *ISPRS Journal of Photogrammetry and Remote Sensing* 98 (Dec. 2014), 119–132. DOI: 10.1016/j.isprsjprs.2014.10.002

[22] C. Collange, David Defour, and Arnaud Tisserand. 2009. Power consumption of GPUs from a software perspective. In *Computational Science – ICCS 2009*. Springer, Berlin, 914–923. DOI: 10.1007/978-3-642-01970-8_92

[23] Jason Cong, Zhenman Fang, Michael Lo, Hanrui Wang, Jingxian Xu, and Shaochong Zhang. 2018. Understanding performance differences of FPGAs and GPUs. In *2018 IEEE 26th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE, Boulder, CO, USA, 93–96. DOI: 10.1109/FCCM.2018.00023

[24] Corporation for Digital Scholarship. 2026. Zotero (7.0.11) [Software]. https://www.zotero.org/. (Original work published 2006).

[25] Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine Learning* 20, 3 (Sept. 1995), 273–297. DOI: 10.1007/BF00994018

[26] Angela Cratere, M. Salim Farissi, Andrea Carbone, Marcello Asciolla, Maria Rizzi, Francesco Dell'Olio, Augusto Nascetti, and Dario Spiller. 2025. Efficient FPGA-accelerated convolutional neural networks for cloud detection on CubeSats. *IEEE Journal on Miniaturization for Air and Space Systems* (2025), 1–1. DOI: 10.1109/JMASS.2025.3533018

[27] Steve Foga, Pat L. Scaramuzza, Song Guo, Zhe Zhu, Ronald D. Dilley, Tim Beckmann, Gail L. Schmidt, John L. Dwyer, M. Joseph Hughes, and Brady Laue. 2017. Cloud detection algorithm comparison and validation for operational landsat data products. *Remote Sensing of Environment* 194 (June 2017), 379–390. DOI: 10.1016/j.rse.2017.03.026

[28] P. Fraczek, A. Mora, and T. Kryjak. 2018. Embedded vision system for automated drone landing site detection. In *AGH University of Krakow*, L. J. Chmielewski, R. Kozera, A. Orlowski, K. Wojciechowski, A. M. Bruckstein, and N. Petkov (Eds.). Vol. 11114. 397–409. DOI: 10.1007/978-3-030-00692-1_35

[29] Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of The 33rd International Conference on Machine Learning*. PMLR, 1050–1059.

[30] Antonio-Javier Gallego, Antonio Pertusa, and Pablo Gil. 2018. Automatic ship classification from optical aerial images with convolutional neural networks. *Remote Sensing* 10, 4 (April 2018), 511. DOI: 10.3390/rs10040511

[31] D. Garg, D. Shah, V. Chintapalli, S. Paul, and A. B. Mishra. 2024. Aircraft detection in satellite imagery based on deep learning/AI techniques and FPGA /SoC based hardware implementation. In *2024 IEEE SPACE, Aerospace and Defence Conference, SPACE 2024*. 112–115. DOI: 10.1109/SPACE63117.2024.10667719

[32] Jakob Gawlikowski, Cedrique Rovile Njieutcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. 2023. A Survey of Uncertainty in Deep Neural Networks. *Artificial Intelligence Review* 56, 1 (Oct. 2023), 1513–1589. DOI: 10.1007/s10462-023-10562-9

[33] Alan D. George and Christopher M. Wilson. 2018. Onboard processing with hybrid and reconfigurable computing on small satellites. *Proc. IEEE* 106, 3 (March 2018), 458–470. DOI: 10.1109/JPROC.2018.2802438

[34] Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W. Mahoney, and Kurt Keutzer. 2021. A Survey of Quantization Methods for Efficient Neural Network Inference. arXiv:2103.13630 [cs] DOI: 10.48550/arXiv.2103.13630

[35] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, Columbus, OH, USA, 580–587. DOI: 10.1109/CVPR.2014.81

[36] Gianluca Giuffrida, Lorenzo Diana, Francesco de Gioia, Gionata Benelli, Gabriele Meoni, Massimiliano Donati, and Luca Fanucci. 2020. CloudScout: A deep neural network for on-board cloud detection on hyperspectral images. *Remote Sensing* 12, 14 (Jan. 2020), 2205. DOI: 10.3390/rs12142205

[37] Gianluca Giuffrida, Luca Fanucci, Gabriele Meoni, Matej Batič, Léonie Buckley, Aubrey Dunne, Chris van Dijk, Marco Esposito, John Hefele, Nathan Vercruyssen, et al. 2022. The Φ-Sat-1 mission: The first on-board deep neural network demonstrator for satellite earth observation. *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022), 1–14. DOI: 10.1109/TGRS.2021.3125567

[38] Carlos Gomes, Isabelle Wittmann, Damien Robert, Johannes Jakubik, Tim Reichelt, Stefano Maurogiovanni, Rikard Vinge, Jonas Hurst, Erik Scheurer, Rocco Sedona, et al. 2025. Lossy neural compression for geospatial analytics: A review. *IEEE Geoscience and Remote Sensing Magazine* 13, 3 (Sept. 2025), 97–135. DOI: 10.1109/MGRS.2025.3546527

[39] Google. 2019. Coral TPU Dev Board. Retrieved from https://coral.ai/docs/dev-board/datasheet/#system-components

[40] P. Gordo, T. Frederico, R. Melicio, S. Duzellier, and A. Amorim. 2020. System for space materials evaluation in LEO environment. *Advances in Space Research* 66, 2 (July 2020), 307–320. DOI: 10.1016/j.asr.2020.03.024

[41] Naresh Grover and M. K. Soni. 2012. Reduction of power consumption in FPGAs - An overview. *International Journal of Information Engineering and Electronic Business* 4, 5 (Oct. 2012), 50–69. DOI: 10.5815/ijieeb.2012.05.07

[42] Dubacharla Gyaneshwar and Rama Rao Nidamanuri. 2022. A real-time SC¡SUP¿2¡/SUP¿S-based open-set recognition in remote sensing imagery. *Journal of Real-Time Image Processing* 19, 5 (Oct. 2022), 867–880. DOI: 10.1007/s11554-022-01226-y

[43] Bilal Hammoud, Charbel Bou Maroun, Jonas Ney, and Norbert Wehn. 2022. Artificial neural networks-based radar remote sensing to estimate geographical information during oil-spills. In *2022 30th European Signal Processing Conference (EUSIPCO)*. 1801–1805. DOI: 10.23919/EUSIPCO55093.2022.9909727

[44] B. Hammoud, J. Ney, C. B. Maroun, and N. Wehn. 2024. FPGA onboard processing of tiny-ML models using radar-sensing for oil spill monitoring. In *Lebanese American University*. 6339–6342. DOI: 10.1109/IGARSS53475.2024.10640763

[45] Kai Han, Yunhe Wang, Qi Tian, Jianyuan Guo, Chunjing Xu, and Chang Xu. 2020. GhostNet: More features from cheap operations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1580–1589.

[46] Song Han, Huizi Mao, and William J. Dally. 2016. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.).

[47] Shintaro Hashimoto, Yohei Sugimoto, Ko Hamamoto, and Naoki Ishihama. 2019. Ship classification from SAR images based on deep learning. In *Intelligent Systems and Applications, Vol 1 (Advances in Intelligent Systems and Computing, Vol. 868)*, K. Arai, S. Kapoor, and R. Bhatia (Eds.). 18–34. DOI: 10.1007/978-3-030-01054-6_2

[48] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.

[49] Wenjing He, Yuesong Yang, Shaohui Mei, Jian Hu, Wanqiu Xu, and Shiqi Hao. 2023. Configurable 2D–3D CNNs Accelerator for FPGA-Based Hyperspectral Imagery Classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 16 (2023), 9406–9421. DOI: 10.1109/JSTARS.2023.3321965

[50] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the Knowledge in a Neural Network. arXiv:1503.02531 [stat] DOI: 10.48550/arXiv.1503.02531

[51] Adrian Höhl, Ivica Obadic, Miguel Ángel Fernández Torres, Hiba Najjar, Dario Oliveira, Zeynep Akata, Andreas Dengel, and Xiao Xiang Zhu. 2024. Opening the Black-Box: A Systematic Review on Explainable AI in Remote Sensing. arXiv:2402.13791 [cs] DOI: 10.48550/arXiv.2402.13791

[52] Andrew Howard, Mark Sandler, Bo Chen, Weijun Wang, Liang-Chieh Chen, Mingxing Tan, Grace Chu, Vijay Vasudevan, Yukun Zhu, Ruoming Pang, et al. 2019. Searching for MobileNetV3. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Seoul, Korea (South), 1314–1324. DOI: 10.1109/ICCV.2019.00140

[53] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. 2017. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. arXiv:1704.04861 [cs] DOI: 10.48550/arXiv.1704.04861

[54] CH Huang, WT Chen, YC Chang, and KT Wu. 2024. An edge and trustworthy AI UAV system with self-adaptivity and hyperspectral imaging for air quality monitoring. *IEEE Internet of Things Journal* 11, 20 (Oct. 2024), 32572–32584. DOI: 10.1109/JIOT.2024.3422470

[55] Forrest N. Iandola, Song Han, Matthew W. Moskewicz, Khalid Ashraf, William J. Dally, and Kurt Keutzer. 2016. SqueezeNet: AlexNet-level Accuracy with 50x Fewer Parameters and <0.5MB Model Size. arXiv:1602.07360 DOI: 10.48550/arXiv.1602.07360

[56] Cosimo Ieracitano, Nadia Mammone, Fanny Spagnolo, Fabio Frustaci, Stefania Perri, Pasquale Corsonello, and Francesco C. Morabito. 2024. An explainable embedded neural system for on-board ship detection from optical satellite imagery. *Engineering Applications of Artificial Intelligence* 133, 108517 (July 2024). DOI: 10.1016/.engappai.2024.108517

[57] ISPRS. 2022. 2D Semantic Labeling Contest - Potsdam. Retrieved from https://www.isprs.org/education/benchmarks/UrbanSemLab/2d-sem-label-potsdam.aspx

[58] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. 2018. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Salt Lake City, UT, 2704–2713. DOI: 10.1109/CVPR.2018.00286

[59] Norm Jouppi, George Kurian, Sheng Li, Peter Ma, Rahul Nagarajan, Lifeng Nai, Nishant Patil, Suvinay Subramanian, Andy Swing, Brian Towles, et al. 2023. TPU v4: An optically reconfigurable supercomputer for machine learning with hardware support for embeddings. In *Proceedings of the 50th Annual International Symposium on Computer Architecture* (Orlando, FL, USA) *(ISCA'23)*. Association for Computing Machinery, New York, NY, USA, Article 82, 14 pages. DOI: 10.1145/3579371.3589350

[60] Shreeyam Kacker, Alex Meredith, Kerri Cahoy, and Georges Labrèche. 2022. Machine learning image processing algorithms onboard OPS-SAT. *Small Satellite Conference* (Aug. 2022).

[61] Xudong Kang. 2017. Airport-Beach-Urban Website. Retrieved from http://xudongkang.weebly.com/data-sets.html

[62] JH Kim, Y Kim, DH Cho, and SM Kim. 2024. On-orbit AI: Cloud detection technique for resource-limited nanosatellite. *International Journal of Aeronautical and Space Sciences* (Dec. 2024). DOI: 10.1007/s42405-024-00865-8

[63] Vivek Kothari, Edgar Liberis, and Nicholas D. Lane. 2020. The final frontier: Deep learning in space. In *Proceedings of the 21st International Workshop on Mobile Computing Systems and Applications*. ACM, Austin TX USA, 45–49. DOI: 10.1145/3376897.3377864

[64] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, Vol. 25. Curran Associates, Inc.

[65] Kung. 1982. Why systolic architectures? *Computer* 15, 1 (Jan. 1982), 37–46. DOI: 10.1109/MC.1982.1653825

[66] Kevin Lange, Federico Fontana, Francesco Rossi, Mattia Varile, and Giovanni Apruzzese. 2024. Machine learning in space: Surveying the robustness of on-board ML models to radiation. In *2024 IEEE Space Computing Conference (SCC)*. IEEE, Mountain View, CA, USA, 51–64. DOI: 10.1109/SCC61854.2024.00012

[67] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (May 2015), 436–444. DOI: 10.1038/nature14539

[68] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (Nov. 1998), 2278–2324. DOI: 10.1109/5.726791

[69] George Lentaris, Konstantinos Maragos, Ioannis Stratakos, Lazaros Papadopoulos, Odysseas Papanikolaou, Dimitrios Soudris, Manolis Lourakis, Xenophon Zabulis, David Gonzalez-Arjona, and Gianluca Furano. 2018. High-performance embedded computing in space: Evaluation of platforms for vision-based navigation. *Journal of Aerospace Information Systems* 15, 4 (2018), 178–192. DOI: 10.2514/1.I010555

[70] Chao Li, Rui Xu, Yong Lv, Yonghui Zhao, and Weipeng Jing. 2023. Edge real-time object detection and DPU-based hardware implementation for optical remote sensing images. *Remote Sensing* 15, 3975 (Aug. 2023). DOI: 10.3390/rs15163975

[71] Chaojian Li, Zhongzhi Yu, Yonggan Fu, Yongan Zhang, Yang Zhao, Haoran You, Qixuan Yu, Yue Wang, Cong Hao, and Yingyan Lin. 2021. HW-NAS-bench: Hardware-aware neural architecture search benchmark. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

[72] Jianwei Li, Changwen Qu, and Jiaqi Shao. 2017. Ship detection in SAR images based on an improved faster R-CNN. In *2017 SAR in Big Data Era: Models, Methods and Applications (BIGSARDATA)*. 1–6. DOI: 10.1109/BIGSARDATA.2017.8124934

[73] Ke Li, Gang Wan, Gong Cheng, Liqiu Meng, and Junwei Han. 2020. Object detection in optical remote sensing images: A survey and a new benchmark. *ISPRS Journal of Photogrammetry and Remote Sensing* 159 (Jan. 2020), 296–307. DOI: 10.1016/j.isprsjprs.2019.11.023

[74] Lin Li, Shengbing Zhang, and Juan Wu. 2019. Efficient object detection framework and hardware architecture for remote sensing images. *Remote Sensing* 11, 2376 (Oct. 2019). DOI: 10.3390/rs11202376

[75] S. Li, K. Y. Sun, Y. K. Luo, N. Yadav, and K. Choi. 2020. Novel CNN-based AP2D-net accelerator: An area and power efficient solution for real-time applications on mobile FPGA. *Electronics* 9, 5 (May 2020). DOI: 10.3390/electronics9050832

[76] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. 2016. SSD: Single shot multibox detector. In *Computer Vision – ECCV 2016*, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.). Springer International Publishing, Cham, 21–37. DOI: 10.1007/978-3-319-46448-0_2

[77] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. 2022. A ConvNet for the 2020s. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, New Orleans, LA, USA, 11966–11976. DOI: 10.1109/CVPR52688.2022.01167

[78] S. Lopez, T. Vladimirova, C. Gonzalez, J. Resano, D. Mozos, and A. Plaza. 2013. The promise of reconfigurable computing for hyperspectral imaging onboard systems: A review and trends. *Proc. IEEE* 101, 3 (2013), 698–722. DOI: 10.1109/JPROC.2012.2231391

[79] Lucas Amilton Martins, Felipe Viel, Laio Oriel Seman, Eduardo Augusto Bezerra, and Cesar Albenes Zeferino. 2024. A real-time SVM-based hardware accelerator for hyperspectral images classification in FPGA. *Microprocessors and Microsystems* 104, 104998 (Feb. 2024). DOI: 10.1016/j.micpro.2023.104998

[80] Gonzalo Mateo-Garcia, Joshua Veitch-Michaelis, Lewis Smith, Silviu Vlad Oprea, Guy Schumann, Yarin Gal, Atılım Güneş Baydin, and Dietmar Backes. 2021. Towards global flood mapping onboard low cost satellites with machine learning. *Scientific Reports* 11, 1 (March 2021), 7249. DOI: 10.1038/s41598-021-86650-z

[81] J. P. Matos-Carvalho, F. Moutinho, A. B. Salvado, T. Carrasqueira, R. Campos-Rebelo, D. Pedro, L. M. Campos, J. M. Fonseca, and A. Mora. 2019. Static and dynamic algorithms for terrain classification in UAV aerial imagery. *Remote Sensing* 11, 21 (Nov. 2019). DOI: 10.3390/rs11212501

[82] Sorour Mohajerani and Parvaneh Saeedi. 2019. Cloud-net: An end-to-end cloud detection algorithm for landsat 8 imagery. In *IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium*. 1029–1032. DOI: 10.1109/IGARSS.2019.8898776

[83] T. Myojin, S. Hashimoto, and N. Ishihama. 2020. Detecting uncertain BNN outputs on FPGA using monte carlo dropout sampling. In *Japan Aerospace Exploration Agency (JAXA)*, I. Farkas, P. Masulli, and S. Wermter (Eds.). Vol. 12397. 27–38. DOI: 10.1007/978-3-030-61616-8_3

[84] Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart van Baalen, and Tijmen Blankevoort. 2021. A White Paper on Neural Network Quantization. arXiv:2106.08295 [cs] DOI: 10.48550/arXiv.2106.08295

[85] Romén Neris, Adrián Rodríguez, Raúl Guerra, Sebastián López, and Roberto Sarmiento. 2022. FPGA-based implementation of a CNN architecture for the on-board processing of very high-resolution remote sensing images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 15 (2022), 3740–3750. DOI: 10.1109/JSTARS.2022.3169330

[86] DD Nguyen, DT Nguyen, MT Le, and QC Nguyen. 2024. FPGA-SoC implementation of YOLOv4 for flying-object detection. *Journal of Real-Time Image Processing* 21, 3 (May 2024). DOI: 10.1007/s11554-024-01440-w

[87] Ty Nguyen, Ian D. Miller, Avi Cohen, Dinesh Thakur, Arjun Guru, Shashank Prasad, Camillo J. Taylor, Pratik Chaudhari, and Vijay Kumar. 2021. PennSyn2Real: Training object recognition models without human labeling. *IEEE Robotics and Automation Letters* 6, 3 (July 2021), 5032–5039. DOI: 10.1109/LRA.2021.3070249

[88] Shuo Ni, Xin Wei, Ning Zhang, and He Chen. 2023. Algorithm–hardware co-optimization and deployment method for field-programmable gate-array-based convolutional neural network remote sensing image processing. *Remote Sensing* 15, 24 (Jan. 2023), 5784. DOI: 10.3390/rs15245784

[89] NVIDIA. 2019. NVIDIA Jetson Nano Module. Retrieved from https://developer.nvidia.com/embedded/jetson-nano

[90] NVIDIA. 2025. NVIDIA Blackwell Architecture. Retrieved from https://www.nvidia.com/en-us/data-center/technologies/blackwell-architecture/

[91] Lucas Prado Osco, José Marcato Junior, Ana Paula Marques Ramos, Lúcio André de Castro Jorge, Sarah Narges Fatholahi, Jonathan de Andrade Silva, Edson Takashi Matsubara, Hemerson Pistori, Wesley Nunes Gonçalves, and Jonathan Li. 2021. A review on deep learning in UAV remote sensing. *International Journal of Applied Earth Observation and Geoinformation* 102 (Oct. 2021), 102456. DOI: 10.1016/j.jag.2021.102456

[92] Matthew J. Page, David Moher, Patrick M. Bossuyt, Isabelle Boutron, Tammy C. Hoffmann, Cynthia D. Mulrow, Larissa Shamseer, Jennifer M. Tetzlaff, Elie A. Akl, Sue E. Brennan, et al. 2021. PRISMA 2020 explanation and elaboration: updated guidance and exemplars for reporting systematic reviews. *BMJ* 372 (March 2021), n160. DOI: 10.1136/bmj.n160

[93] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. 2016. ENet: A Deep Neural Network Architecture for Real-Time Semantic Segmentation. arXiv:1606.02147 DOI: 10.48550/arXiv.1606.02147

[94] Noah Perryman, Sebastian Sabogal, Christopher Wilson, and Alan George. 2025. Dependable DPU architectures on AMD-xilinx versal adaptive SoCs for space applications. *IEEE Trans. Aerospace Electron. Systems* 61, 3 (June 2025), 6629–6646. DOI: 10.1109/TAES.2025.3527938

[95] R. Pitonak, J. Mucha, L. Dobis, M. Javorka, and M. Marusin. 2022. CloudSatNet-1: FPGA-based hardware-accelerated quantized CNN for satellite on-board cloud coverage classification. *Remote Sensing* 14, 13 (July 2022). DOI: 10.3390/rs14133180

[96] J. R. Quinlan. 1986. Induction of decision trees. *Machine Learning* 1, 1 (March 1986), 81–106. DOI: 10.1007/BF00116251

[97] Amir Raoofy, Gabriel Dax, Max Ghiglione, Martin Langer, Carsten Trinitis, Martin Werner, and Martin Schulz. 2021. Benchmarking machine learning inference in FPGA-based accelerated space applications. *Proceedings of the Workshop on Benchmarking Machine Learning Workloads Co-Located with IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)* (2021).

[98] Emilio Rapuano, Gabriele Meoni, Tommaso Pacini, Gianmarco Dinelli, Gianluca Furano, Gianluca Giuffrida, and Luca Fanucci. 2021. An FPGA-based hardware accelerator for Cnns inference on board satellites: Benchmarking with myriad 2-based solution for the cloudscout case study. *Remote Sensing* 13, 1518 (April 2021). DOI: 10.3390/rs13081518

[99] R. Ratnakumar and S. J. Nanda. 2021. A high speed roller dung beetles clustering algorithm and its architecture for real-time image segmentation. *Applied Intelligence* 51, 7 (July 2021), 4682–4713. DOI: 10.1007/s10489-020-02067-7

[100] Joseph Redmon and Ali Farhadi. 2017. YOLO9000: Better, faster, stronger. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7263–7271.

[101] Joseph Redmon and Ali Farhadi. 2018. YOLOv3: An Incremental Improvement. arXiv:1804.02767 [cs] DOI: 10.48550/arXiv.1804.02767

[102] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi (Eds.). Springer International Publishing, Cham, 234–241. DOI: 10.1007/978-3-319-24574-4_28

[103] Vít Růžička, Anna Vaughan, Daniele De Martini, James Fulton, Valentina Salvatelli, Chris Bridges, Gonzalo Mateo-Garcia, and Valentina Zantedeschi. 2022. RaVÆn: Unsupervised change detection of extreme events using ML on-board satellites. *Scientific Reports* 12, 1 (Oct. 2022), 16939. DOI: 10.1038/s41598-022-19437-5

[104] Sebastian Sabogal and Alan George. 2021. A methodology for evaluating and analyzing FPGA-accelerated, deep-learning applications for onboard space processing. In *2021 IEEE Space Computing Conference (Scc)*. 143–154. DOI: 10.1109/SCC49971.2021.00022

[105] Shaibal Saha and Lanyu Xu. 2025. Vision transformers on the edge: A Comprehensive survey of model compression and acceleration strategies. *Neurocomputing* 643 (Aug. 2025), 130417. DOI: 10.1016/j.neucom.2025.130417

[106] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. MobileNetV2: Inverted residuals and linear bottlenecks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4510–4520. DOI: 10.1109/CVPR.2018.00474

[107] SDMS. 2010. MSTAR Dataset. Retrieved from https://www.sdms.afrl.af.mil/index.php?collection=mstar

[108] Ahmad Shawahna, Sadiq M. Sait, and Aiman El-Maleh. 2019. FPGA-based accelerators of deep learning networks for learning and classification: A review. *IEEE Access* 7 (2019), 7823–7859. DOI: 10.1109/ACCESS.2018.2890150

[109] Sherin C. Shibi and R. Gayathri. 2021. Onboard target detection in hyperspectral image based on deep learning with FPGA implementation. *Microprocessors and Microsystems* 85, 104313 (Sept. 2021). DOI: 10.1016/j.micpro.2021.104313

[110] Felix Siegle, Tanya Vladimirova, Jørgen Ilstad, and Omar Emam. 2015. Mitigation of radiation effects in SRAM-based FPGAs for space applications. *ACM Comput. Surv.* 47, 2 (Jan. 2015), 37:1–37:34. DOI: 10.1145/2671181

[111] Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

[112] H. S. Suh, J. Meng, T. Nguyen, S. K. Venkataramanaiah, V. Kumar, Y. Cao, and J. S. Seo. 2021. Algorithm-hardware co-optimization for energy-efficient drone detection on resource-constrained FPGA. In *Arizona State University*. 133–141. DOI: 10.1109/ICFPT52863.2021.9609840

[113] Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, and Joel S. Emer. 2017. Efficient processing of deep neural networks: A tutorial and survey. *Proc. IEEE* 105, 12 (Dec. 2017), 2295–2329. DOI: 10.1109/JPROC.2017.2761740

[114] Daniela Szwarcman, Sujit Roy, Paolo Fraccaro, Þorsteinn Elí Gíslason, Benedikt Blumenstiel, Rinki Ghosal, Pedro Henrique de Oliveira, Joao Lucas de Sousa Almeida, Rocco Sedona, Yanghui Kang, et al. 2025. Prithvi-EO-2.0: A Versatile Multi-Temporal Foundation Model for Earth Observation Applications. arXiv:2412.02732 [cs] DOI: 10.48550/arXiv.2412.02732

[115] Mingxing Tan and Quoc V. Le. 2020. EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proceedings of Machine Learning Research Volume 97*. ICML. arXiv:1905.11946 [cs].

[116] VAMF Torres, B. R. A. Jaimes, E. S. Ribeiro, M. T. Braga, E. H. Shiguemori, H. F. C. Velho, L. C. B. Torres, and A. P. Braga. 2020. Combined weightless neural network FPGA architecture for deforestation surveillance and visual navigation of UAVs. *Engineering Applications of Artificial Intelligence* 87 (Jan. 2020). DOI: 10.1016/j.engappai.2019.08.021

[117] Yaman Umuroglu, Nicholas J. Fraser, Giulio Gambardella, Michaela Blott, Philip Leong, Magnus Jahre, and Kees Vissers. 2017. FINN: A framework for fast, scalable binarized neural network inference. In *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. ACM, Monterey California USA, 65–74. DOI: 10.1145/3020078.3021744

[118] G. Upadhyay, S. Ghosal, S. Kart, K. Jain, SH Shantala, J. T. LalitKrushna, S. Srividhya, and S. A. Balwantrao. 2024. Design and implementation of CNN-based custom net architecture with improved inference time for realtime remote sensing application. In *2024 IEEE SPACE, AEROSPACE AND DEFENCE CONFERENCE, SPACE 2024*. 647–651. DOI: 10.1109/SPACE63117.2024.10667853

[119] P. Vitolo, A. Fasolino, R. Liguori, L. Di Benedetto, A. Rubino, and G. D. Licciardo. 2024. Real-time on-board satellite cloud cover detection hardware architecture using spaceborne remote sensing imagery. In *University of Salerno*, M. F. Carlsohn (Ed.). Vol. 13000. DOI: 10.1117/12.3017554

[120] Erwei Wang, James J. Davis, Peter Y. K. Cheung, and George A. Constantinides. 2019. LUTNet: Rethinking inference in FPGA soft logic. In *2019 IEEE 27th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE, San Diego, CA, USA, 26–34. DOI: 10.1109/FCCM.2019.00014

[121] Erwei Wang, James J. Davis, Ruizhe Zhao, Ho-Cheung Ng, Xinyu Niu, Wayne Luk, Peter Y. K. Cheung, and George A. Constantinides. 2020. Deep neural network approximation for custom hardware: Where we've been, where we're going. *Comput. Surveys* 52, 2 (March 2020), 1–39. DOI: 10.1145/3309551

[122] EY Wang and DH Qiu. 2019. Acceleration and implementation of convolutional neural network based on FPGA. In *Capital Normal University*. 321–325. DOI: 10.1109/iccsnt47585.2019.8962428

[123] Xiyue Wang, Xinsheng Wang, Zhiquan Zhou, and Yanhong Song. 2024 AUG 26 2024. Fast detection and obstacle avoidance on UAVs using lightweight convolutional neural network based on the fusion of radar and camera. *Applied Intelligence* (2024 AUG 26 2024). DOI: 10.1007/s10489-024-05768-5

[124] Adam C. Watts, Vincent G. Ambrosia, and Everett A. Hinkley. 2012. Unmanned aircraft systems in remote sensing and scientific research: Classification and considerations of use. *Remote Sensing* 4, 6 (June 2012), 1671–1692. DOI: 10.3390/rs4061671

[125] X. Wei, W. C. Liu, L. Chen, L. Ma, H. Chen, and Y. Zhuang. 2019. FPGA-based hybrid-type implementation of quantized neural networks for remote sensing applications. *SENSORS* 19, 4 (Feb. 2019). DOI: 10.3390/s19040924

[126] Sachini Wickramasinghe, Dhruv Parikh, Bingyi Zhang, Rajgopal Kannan, Viktor Prasanna, and Carl Busart. 2024. VTR: An optimized vision transformer for SAR ATR acceleration on FPGA. In *Image Sensing Technologies: Materials, Devices, Systems, and Applications XI*, Nibir K. Dhar, Achyut K. Dutta, and Sachidananda R. Babu (Eds.). SPIE, National Harbor, United States, 37. DOI: 10.1117/12.3013580

[127] W. B. Wu, C. H. Cao, Y. Li, R. X. Zhang, and S. L. Dong. 2021. Design and implementation of remote sensing object detection system based on MPSoC. In *AOPC 2021: Optical Sensing and Imaging Technology*, Y. Jiang, Q. Lv, D. Liu, D. Zhang, and B. Xue (Eds.). Vol. 12065. DOI: 10.1117/12.2606142

[128] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. 2018. DOTA: A large-scale dataset for object detection in aerial images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3974–3983.

[129] Xiaowei Xu, Xinyi Zhang, Bei Yu, Xiaobo Sharon Hu, Christopher Rowen, Jingtong Hu, and Yiyu Shi. 2021. DAC-SDC low power object detection challenge for UAV applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43, 2 (Feb. 2021), 392–403. DOI: 10.1109/TPAMI.2019.2932429

[130] Réda Yahiaoui, Farid Alilat, and Saliha Loumi. 2017. Parallelization of fuzzy ARTMAP architecture on FPGA: Multi-spectral classification of ALSAT-2A images. *IEEE Transactions on Industrial Electronics* 64, 12 (Dec. 2017), 9487–9495. DOI: 10.1109/TIE.2017.2708028

[131] Tianwei Yan, Ning Zhang, Jie Li, Wenchao Liu, and He Chen. 2022. Automatic deployment of convolutional neural networks on FPGA for spaceborne remote sensing application. *Remote Sensing* 14, 13 (Jan. 2022), 3130. DOI: 10.3390/rs14133130

[132] G. Yang, J. Lei, W. Y. Xie, Z. M. Fang, Y. S. Li, J. X. Wang, and X. Zhang. 2022. Algorithm/hardware codesign for real-time on-satellite CNN-based ship detection in SAR Imagery. *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022). DOI: 10.1109/TGRS.2022.3161499

[133] RH Yang, ZK Chen, BA Wang, YF Guo, and LT Hu. 2023. A lightweight detection method for remote sensing images and its energy-efficient accelerator on edge devices. *SENSORS* 23, 14 (July 2023). DOI: 10.3390/s23146497

[134] Yi Yang and Shawn Newsam. 2010. Bag-of-visual-words and spatial extensions for land-use classification. In *Proceedings of the 18th Sigspatial International Conference on Advances in Geographic Information Systems (GIS'10)*. Association for Computing Machinery, New York, NY, USA, 270–279. DOI: 10.1145/1869790.1869829

[135] Zhewei Yao, Zhen Dong, Zhangcheng Zheng, Amir Gholami, Jiali Yu, Eric Tan, Leyuan Wang, Qijing Huang, Yida Wang, Michael Mahoney, et al. 2021. HAWQ-V3: Dyadic neural network quantization. In *Proceedings of the 38th International Conference on Machine Learning*. PMLR, 11875–11886.

[136] Q. Yu, A. Liu, X. X. Yang, and W. M. Diao. 2024. An improved lightweight deep learning model and implementation for track fastener defect detection with unmanned aerial vehicles. *ELECTRONICS* 13, 9 (May 2024). DOI: 10.3390/electronics13091781

[137] Bingyi Zhang, Rajgopal Kannan, Viktor Prasanna, and Carl Busart. 2022. Accurate, low-latency, efficient SAR automatic target recognition on FPGA. In *2022 32nd International Conference on Field-Programmable Logic and Applications (FPL)*. IEEE Computer Society, 1–8. DOI: 10.1109/FPL57034.2022.00013

[138] B. Y. Zhang, R. Kannan, V. Prasanna, and C. Busart. 2023. Accelerating GNN-based SAR automatic target recognition on HBM-enabled FPGA. In *University of Southern California*. DOI: 10.1109/HPEC58863.2023.10363615

[139] Ning Zhang, Shuo Ni, Tingting Qiao, Wenchao Liu, and He Chen. 2023. An extremely pipelined FPGA-based accelerator of all adder neural networks for on-board remote sensing scene classification. In *2023 International Conference on Field Programmable Technology, Icfpt (IEEE International Conference on Field Programmable Technology)*. 258–261. DOI: 10.1109/ICFPT59805.2023.00036

[140] Ning Zhang, Xin Wei, He Chen, and Wenchao Liu. 2021. FPGA implementation for CNN-based optical remote sensing object detection. *ELECTRONICS* 10, 282 (Feb. 2021). DOI: 10.3390/electronics10030282

[141] X. L. Zhang, X. Wei, Q. B. Sang, H. Chen, and Y. Z. Xie. 2020. An efficient FPGA-based implementation for quantized remote sensing image scene classification network. *ELECTRONICS* 9, 9 (Sept. 2020). DOI: 10.3390/electronics9091344

[142] Qiang Zhao, Le Yu, Zhenrong Du, Dailiang Peng, Pengyu Hao, Yongguang Zhang, and Peng Gong. 2022. An overview of the applications of earth observation satellite data: Impacts and future trends. *Remote Sensing* 14, 8 (Jan. 2022), 1863. DOI: 10.3390/rs14081863

[143] Yonghui Zhao, Yong Lv, and Chao Li. 2023. Hardware acceleration of satellite remote sensing image object detection based on channel pruning. *Applied Sciences-Basel* 13, 10111 (Sept. 2023). DOI: 10.3390/app131810111

[144] Daquan Zhou, Qibin Hou, Yunpeng Chen, Jiashi Feng, and Shuicheng Yan. 2020. Rethinking bottleneck structure for efficient mobile network design. In *Computer Vision – ECCV 2020*. Springer International Publishing, Cham, 680–697. DOI: 10.1007/978-3-030-58580-8_40

[145] Xiao Xiang Zhu, Devis Tuia, Lichao Mou, Gui-Song Xia, Liangpei Zhang, Feng Xu, and Friedrich Fraundorfer. 2017. Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine* 5, 4 (Dec. 2017), 8–36. DOI: 10.1109/MGRS.2017.2762307

Received 3 June 2025; revised 27 November 2025; accepted 19 January 2026