A Technical and Historical Overview of NVIDIA Jetson Platforms (2014-Present) and Sony SenSWIR MX Series (2019-Present) for Autonomous Marine Microplastic Detection

R.K.D. Kho, R.J.E. Daniels (data and review), The Center for Research and Development Hong Kong (CRD-HK) Tsuen Wan N.T. - Hong Kong, dept. Micro-Electronics, 1 January 2025, Spectral Sensing, info@crd-hk.com, DOI: 10.17613/6t0gv-yc737

Abstract

This paper provides a comprehensive technical and historical overview of NVIDIA's Jetson embedded computing platforms, tracing their evolution from the initial Jetson TK1 in 2014 to the powerful Orin series. We detail the significant advancements in computing power and power efficiency across generations, highlighting key architectural changes and their impact on AI and robotics applications. Furthermore, the paper delves into Sony's pioneering SenSWIR MX series image sensors, specifically comparing the initial offerings against the latest MX993 series. We analyze their unique capabilities, including the crucial dual camera modes (visible and SWIR), advancements in cooling technologies, and the considerable advantages these features offer for diverse industrial and scientific applications. This historical perspective aims to illustrate the rapid progress in edge AI hardware and advanced imaging solutions. Building upon this, the paper then integrates these technologies into a detailed technical analysis of an autonomous marine microplastic detection system. This system employs a symbiotic drone-buoy architecture, leveraging the NVIDIA Jetson AGX Orin and Orin NX platforms as computational backbones. The architectural details, capabilities, and intricate interplay of these edge AI processors with the Sony SWIR IMX993 sensor are meticulously presented. A significant focus is placed on the Adaptive SWIR Sensor Readout and Edge Processing (ASREP) algorithm, outlining its conceptual framework and precise

mathematical formulations for adaptive exposure, gain, and frame rate control. The NVIDIA Isaac ROS robotic development platform and the utility of NVIDIA Sim for synthetic data generation and surface/area data processing are also elucidated for their roles in facilitating the real-time, accelerated execution of advanced machine learning models and the robust development of the system. Drawing from foundational and advanced technical analyses in embedded systems and sensor modeling underscores the critical role of hardware-software co-design and dynamic mathematical optimization in achieving efficient, scalable, and autonomous environmental monitoring [R.K.D. Kho 2017, 2019, 2020, 2025].

1. Introduction

The fields of embedded AI and advanced imaging have witnessed exponential growth over the past decade. The demand for powerful, energy-efficient computing at the edge, coupled with increasingly sophisticated sensor technologies, has driven innovation in both hardware and software. This paper aims to provide a structured technical and historical overview of two key players in this evolution: NVIDIA's Jetson platform and Sony's SenSWIR MX series image sensors. By examining their development from their respective introductions to the present day, we can better understand the trajectory of these technologies and their profound impact on various industries. Furthermore, this paper will illustrate the practical application of these cutting-edge technologies within the context of an autonomous marine microplastic detection system.

1.1. Context and Significance of Autonomous SWIR Sensing

The escalating global threat posed by microplastic pollution in marine environments necessitates the development of advanced and scalable monitoring solutions [1.1]. Microplastic particles, defined as less than 5 mm in size, are ubiquitous contaminants stemming from diverse sources, including the fragmentation of larger plastic debris, industrial effluence, and consumer products [1.1]. Their pervasive distribution across all marine environments, from surface waters to deep-sea sediments, underscores the critical need for robust monitoring technologies to comprehend this environmental challenge and formulate effective mitigation strategies [1.1].

Traditional methods for quantifying microplastics, such as laborious manual sampling via nets and subsequent laboratory analysis, are inherently limited in their spatial and temporal scope [1.1]. These techniques often provide only a fragmented snapshot of a highly dynamic problem, which is insufficient for understanding the full extent and movement of microplastic contamination [1.1]. This fundamental limitation in data collection scale highlights an urgent demand for more efficient, autonomous, and scalable monitoring technologies capable of providing broad-area coverage with high temporal resolution [1.1].

In response, remote sensing technologies have emerged as a promising alternative for large-scale, non-invasive surveillance [1.1]. Specifically, imaging within the Short-Wave Infrared (SWIR) spectrum, typically ranging from 1000 nm to 2500 nm, has proven highly effective for identifying various plastic polymers [1.1]. Different plastics exhibit unique spectral absorption and reflection characteristics within this range, enabling their differentiation from surrounding materials such as water, algae, and other organic matter [1.1]. Several studies have successfully demonstrated the potential of SWIR imagery for detecting plastic items, including debris on beaches and in coastal areas [1.1]. The integration of unmanned aerial vehicles (UAVs), or drones, for environmental monitoring further enhances this capability by enabling rapid deployment and efficient coverage [1.1]. This strategic shift from localized, discrete data points to continuous, broad-area coverage is essential for understanding dynamic environmental problems and is a direct consequence of the inherent limitations of manual sampling.

1.2. Overview of the Integrated Drone-Buoy System

This paper introduces a novel design framework for an autonomous system dedicated to the efficient detection and mapping of microplastics on sea surfaces [1.1]. The proposed system integrates a stationary sea buoy and a mobile drone, conceived as a symbiotic pairing of two intelligent platforms [1.1]. The buoy functions as an autonomous base station, serving as the operational anchor and command hub for long-term, autonomous deployment [1.1]. It is equipped with environmental sensors, a Sony Short-Wave Infrared (SWIR) sensor for continuous localized monitoring, and an NVIDIA Jetson AGX Orin 64GB for onboard data processing and mission management [1.1]. The drone, conversely, serves as the primary mobile data acquisition platform, carrying a Sony SWIR IMX993 sensor and an NVIDIA Jetson Orin NX 16GB [1.1].

This integrated approach aims to maximize the scanned sea surface area for microplastic detection within defined operational constraints, specifically a 45-minute drone flight time and a 25-meter maximum scan height [1.1]. The system's intelligence is significantly enhanced by the Adaptive SWIR Sensor Readout and Edge Processing (ASREP) algorithm, which is central to its intelligent data processing pipeline [1.1]. The symbiotic relationship between the buoy and the drone addresses the inherent limitations of each platform individually. Drones possess limited endurance, while stationary buoys have restricted spatial coverage. By combining these elements, the system achieves both persistence through the buoy's continuous operation and broad-area coverage via the drone, enabling repeated, extended autonomous missions without manual intervention [1.1]. This represents a critical advancement for long-term environmental monitoring.

1.3. Scope and Contributions of This Technical Paper

Building upon foundational investigations into integrated SWIR sensing and UAV-buoy systems [1.1], and drawing insights from previous advanced technical analysis of embedded systems and sensor modeling, including the work of Raymund K.D. Kho et al. (2017, 2019, 2020) on topics such as advanced technical analysis of NVIDIA Jetson systems and mathematical modeling for SWIR sensor readout and edge processing, as well as their anticipated contributions in 2025 [1.4, 1.5, 1.7, 1.11], this paper presents a comprehensive strategic framework for a next-generation autonomous microplastic detection system. The primary contribution of this work is a detailed technical analysis of the system's core components, their functional interplay, and the intelligent data processing pipeline. A specific focus is placed on the architectural details and capabilities of the NVIDIA Jetson AGX Orin and Orin NX platforms, the technical specifications and operational considerations of the Sony SWIR IMX993 sensor, the intricate mathematical modeling underpinning the ASREP algorithm, and the integration of NVIDIA Isaac ROS and NVIDIA Sim for robust development and deployment. This paper aims to provide a robust foundation and a clear roadmap for developing autonomous systems capable of supplying critical data for scientific research, environmental policy, and future remediation efforts related to marine microplastic pollution [1.1].

2. Evolution of NVIDIA Jetson Platforms: From Kepler to Ampere

NVIDIA's Jetson series has consistently pushed the boundaries of embedded AI computing, offering developers and engineers increasingly powerful and efficient System-on-Modules (SOMs). This section details the progression of key Jetson modules, focusing on their computational capabilities and power consumption.

2.1. Early Series: Jetson TK1, TX1, and TX2

The foundational Jetson modules laid the groundwork for NVIDIA's dominance in edge AI.

NVIDIA Jetson TK1 (Launched October 2014)

Architecture: Kepler 2.0 (GK20A GPU)

GPU: 192 CUDA Cores

CPU: NVIDIA Tegra K1 (ARM Cortex-A15 based)

Memory: 2 GB DDR3L (64-bit interface, 7.472 GB/s bandwidth)

Computing Power (FP32): 365.2 GFLOPS

Power Consumption (TDP): 8 W

Key Features: Pioneering the concept of an embedded supercomputer, the TK1 demonstrated the potential of GPU-accelerated computing for vision and basic AI tasks at the edge. Its low power consumption for the performance offered was a significant highlight.

NVIDIA Jetson TX1 (Launched 2015)

Architecture: Maxwell (256 CUDA Cores)

CPU: Quad-core ARM Cortex-A57

Memory: 4 GB LPDDR4 (25.6 GB/s bandwidth)

Computing Power (FP32): 1 TFLOP/s

Power Consumption (TDP): 10W - 15W

Key Features: The TX1 significantly boosted performance over the TK1, introducing the Maxwell architecture to the embedded space. It offered greater memory bandwidth and was designed for more complex deep learning applications.

NVIDIA Jetson TX2 (Launched 2017)

Architecture: Pascal (256 CUDA Cores)

CPU: Dual-Core NVIDIA Denver 2 + Quad-Core ARM Cortex-A57

Memory: 8 GB 128-bit LPDDR4 (59.7 GB/s bandwidth)

Computing Power (FP32): 1.3 TFLOP/s (Single Precision), 665.6 GFLOPS (FP32)

Power Consumption (TDP): 7.5W / 15W (configurable)

Key Features: The TX2 offered improved power efficiency and increased performance compared to the TX1, primarily due to the more advanced Pascal architecture and the "Denver 2" custom CPU cores, making it suitable for even more demanding edge AI applications.

2.2. The Xavier Generation: Stepping Up AI Performance

The Xavier series marked a substantial leap in AI inference capabilities, introducing Tensor Cores for the first time in the Jetson line.

NVIDIA Jetson Xavier NX 16GB (Launched 2020)

Architecture: Volta (384 CUDA Cores, 48 Tensor Cores)

CPU: 6-core NVIDIA Carmel ARMv8.2 64-bit CPU

Memory: 16 GB 128-bit LPDDR4X (51.2 GB/s bandwidth)

AI Performance: 21 TOPS (INT8), 1.33 TFLOPS (FP16)

Power Consumption (TDP): 10W, 15W, 20W (configurable)

Comparison to First Jetson NX (8GB): The first **Jetson Xavier NX (8GB)** featured 8GB of LPDDR4X memory and the same GPU/CPU architecture. The 16GB version doubled the memory capacity, which is crucial for running larger deep learning models and more complex multi-sensor applications, where memory size often becomes a bottleneck. While the core AI performance (TOPS) remains similar for a given power mode, the increased memory allows for a greater variety and complexity of workloads to be handled efficiently. The 8GB version also has slightly less memory bandwidth at 25.6 GB/s, while the 16GB version offers 51.2 GB/s.

2.3. The Orin Era: Redefining Edge AI Performance

The Jetson Orin series represents the current pinnacle of NVIDIA's embedded computing, leveraging the Ampere architecture for unprecedented AI performance at the edge.

NVIDIA Jetson AGX Orin 64GB (Launched 2022)

Architecture: Ampere (2048 CUDA Cores, 64 Tensor Cores) [1.2]

CPU: 12-core Arm Cortex-A78AE v8.2 64-bit CPU [1.2]

Memory: 64 GB 256-bit LPDDR5 (204.8 GB/s bandwidth) [1.2]

AI Performance: 275 TOPS (INT8), 10.65 TFLOPS (FP16)

Power Consumption (TDP): 15W - 60W (configurable) [1.2]

Comparison to Older AGX Xavier (32GB): The **AGX Orin 64GB** dwarfs the **AGX Xavier** in every performance metric. The AGX Xavier (32GB) offered up to 32 TOPS (INT8) and 11 TFLOPS (FP16), with 32GB of LPDDR4X memory at 136.5 GB/s, and a TDP of 10W-30W. The AGX Orin 64GB provides an **8.5x increase in AI performance (TOPS)**, significantly higher memory bandwidth, and double the memory capacity. This massive leap is attributed to the Ampere GPU architecture, which brings more powerful Tensor Cores and a more efficient design, enabling the execution of much larger and more sophisticated AI models in real-time.

NVIDIA Jetson Orin NX 16GB (Launched 2022)

Architecture: Ampere (1024 CUDA Cores, 32 Tensor Cores) [1.3]

CPU: 8-core Arm Cortex-A78AE v8.2 64-bit CPU [1.3]

Memory: 16 GB 128-bit LPDDR5 (102.4 GB/s bandwidth) [1.3]

AI Performance: 100 TOPS (INT8)

Power Consumption (TDP): 10W - 25W (configurable) [1.3]

Comparison to First Jetson Xavier NX (8GB/16GB): The Orin NX 16GB is a direct successor to the Xavier NX series. Compared to the Jetson Xavier NX 16GB, the Orin NX 16GB offers a nearly 5x increase in AI performance (TOPS) while maintaining similar power envelopes. This significant boost comes from the shift to the Ampere architecture, which provides superior performance per watt. The Orin NX also features a more powerful CPU and higher memory bandwidth, enabling faster data processing and improved overall system responsiveness for advanced edge AI applications.

Jetson Model	Architecture	GPU (CUDA/Tensor Cores)	CPU	Memory (Capacity/Type/Bandwidth)	AI Performance (INT8/FP16)	Power Consumption (TDP)	Release Year
TK1	Kepler 2.0	192 / -	ARM Cortex-A15	2GB / DDR3L / 7.472 GB/s	365.2 GFLOPS (FP32)	8 W	2014
TX1	Maxwell	256 / -	Quad-core ARM A57	4GB / LPDDR4 / 25.6 GB/s	1 TFLOP/s (FP32)	10W - 15W	2015
TX2	Pascal	256 / -	Dual Denver 2 + Quad A57	8GB / LPDDR4 / 59.7 GB/s	1.3 TFLOP/s (FP32)	7.5W / 15W	2017
Xavie	Volta	384 /	6-core	8GB /	21	10W -	2019

Jetson Model	Architecture	GPU (CUDA /Tensor Cores)	CPU	Memory (Capacity/Type/Bandwidth)	AI Performance (INT8/ FP16)	Power Consumption (TDP)	Release Year
r NX 8GB		48	Carmel ARMv8.2	LPDDR4X / 25.6 GB/s	TOPS (INT8), 1.33 TFLOPS (FP16)	15W	
Xavier NX 16GB	Volta	384 / 48	6-core Carmel ARMv8.2	16GB / LPDDR4X / 51.2 GB/s	21 TOPS (INT8), 1.33 TFLOPS (FP16)	10W - 20W	2020
AGX Xavier 32GB	Volta	512 / 64	8-core Carmel ARMv8.2	32GB / LPDDR4X / 136.5 GB/s	32 TOPS (INT8), 11 TFLOPS (FP16)	10W - 30W	2018
Orin NX 16GB	Ampere	1024 / 32	8-core ARM Cortex-	16GB / LPDDR5 / 102.4	100 TOPS (INT8)	10W - 25W	2022

Jetson Model	Architecture	GPU (CUDA/Tensor Cores)	CPU	Memory (Capacity/Type/Bandwidth)	AI Performance (INT8/FP16)	Power Consumption (TDP)	Release Year
			A78AE	GB/s			
AGX Orin 64GB	Ampere	2048 / 64	12-core ARM Cortex-A78AE	64GB / LPDDR5 / 204.8 GB/s	275 TOPS (INT8)	15W - 60W	2022

Note: TOPS (Tera Operations Per Second) and TFLOPS (Tera Floating-point Operations Per Second) are used to represent AI inference and general floating-point performance, respectively. The exact performance metrics can vary based on workload and optimization.

3. Sony SenSWIR MX Series: Advancements in Short-Wave Infrared Imaging

Sony's SenSWIR technology has revolutionized short-wave infrared (SWIR) imaging by integrating InGaAs (Indium Gallium Arsenide) photodiodes with standard silicon CMOS readout circuits using a unique Cu-Cu connection technology [1.12, 1.13]. This innovation enables compact, high-resolution SWIR sensors that can also capture visible light, opening new possibilities for industrial inspection and scientific research [1.13].

3.1. The First Sony SenSWIR MX Series (e.g., IMX990/IMX991, Introduced 2019)

The initial SenSWIR sensors, such as the **IMX990** and **IMX991**, were groundbreaking for their ability to cover a wide spectral range from 400 nm (visible) to 1700 nm (SWIR) with a single sensor [1.13, 1.14].

Key Features (IMX990/IMX991):

Pixel Size: Industry's smallest 5 µm InGaAs pixels (at the time of release) [3.1].

Resolution: IMX990: Approx. 1.34 effective megapixels (1296x1032), 1/2-inch type. IMX991: Approx. 0.34 effective megapixels (656x520), 1/4-inch type [3.1].

Spectral Range: 400 nm to 1700 nm, enabling both visible and SWIR imaging [1.13, 1.15].

Shutter Type: Global Shutter.

Cooling: Often integrated with thermoelectric (TE) cooling elements (Peltier coolers) to reduce dark current and noise, crucial for high-sensitivity SWIR imaging [1.16, 3.2]. For example, the SWIR330KMA camera based on the IMX991 offers regulated cooling with a max. Δt of 25°C below ambient [2].

Digital Output: Featured integrated digital conversion circuits, simplifying camera design for manufacturers [3.1].

3.2. Advancements in the MX993 Series (e.g., IMX993)

The **MX993 series** builds upon the success of the earlier SenSWIR sensors, pushing the boundaries of resolution and pixel density.

Key Features (IMX993):

Resolution: 3.21 megapixels (2080x1544), making it the highest-resolution SWIR image sensor in the industry at its release [1.6].

Pixel Size: Even smaller at 3.45 μm, further enhancing resolution and allowing for smaller optical formats while maintaining high pixel count [1.6].

Optical Format: 1/1.8" optical format, commonly found in visible-light global shutter sensors, enabling the use of widely available C-mount lenses [1.6].

Spectral Range: Continues to support the 400 nm to 1700 nm range [1.8].

Shutter Type: Global Shutter [1.6].

Cooling: While specific integrated cooling details for the IMX993 are less widely published at the component level, the higher resolution and denser pixel array generally necessitate efficient cooling solutions (often external to the sensor itself, integrated into the camera module) to manage heat and maintain signal-to-noise ratio, especially at higher frame rates. The principles of TE cooling remain essential for optimal SWIR performance [1.1].

Frame Rate: Achieves up to 170fps (8bit), 150fps (10bit), and 90fps (12bit) at full resolution, a significant improvement for high-speed applications [1.6].

3.3. Dual Camera Modes: Normal Camera Mode and SWIR Mode

A significant advantage of Sony's SenSWIR technology, present in both the earlier and newer MX series, is the inherent capability for dual camera modes: visible (normal camera mode) and SWIR mode, within a *single* sensor.

Mechanism: This is achieved through Sony's unique SenSWIR architecture, where photodiodes formed on an InGaAs compound semiconductor layer are connected via Cu-Cu connections with a silicon (Si) layer that forms the readout circuit. This design allows for high sensitivity across a broad wavelength range, from the visible (400 nm) to the short-wave infrared (1700 nm) [1.1, 1.2].

Advantages of Dual Mode:

Reduced System Complexity and Cost: Previously, applications requiring both visible and SWIR data necessitated two separate cameras, each with its own optics, mounting, and processing pipeline [1.19]. A single SenSWIR camera dramatically reduces hardware complexity, integration effort, and overall system cost [1.1, 1.20].

Perfect Pixel-Level Alignment: Since both visible and SWIR images are captured by the *same* sensor, there is no issue of spatial misalignment between the two spectral bands. This is critical for applications requiring precise data fusion or analysis, eliminating the need for complex calibration routines to align images from separate cameras [1.1].

Enhanced Information Gathering: The combination of visible and SWIR data provides a more comprehensive understanding of the scene or object being imaged. Visible light reveals surface features and color, while SWIR light can penetrate certain materials, detect moisture content, identify chemical compositions, or "see through" obscurants like fog or smoke [1.4, 1.5, 1.21].

Expanded Application Scope: This dual-mode capability broadens the range of applications for a single camera system. For example:

Industrial Inspection: Detecting defects or contaminants invisible in the visible spectrum (e.g., moisture in food, foreign objects in packaging, internal cracks in silicon wafers) while simultaneously capturing visual quality [1.4, 1.5, 1.22].

Agriculture: Monitoring plant health by detecting subtle changes in nitrogen content, water stress, or pests not visible to the naked eye, alongside visible light for general crop assessment [1.4].

Recycling: Sorting materials by their unique spectral signatures (e.g., different types of plastics) while also allowing for visual identification [1.4].

Security and Surveillance: Improved night vision and ability to "see through" challenging atmospheric conditions like fog, haze, or smoke, combined with standard visible light surveillance for contextual information [1.4, 1.5].

Art and Cultural Heritage: Revealing underlying layers, detecting repairs, and identifying pigments or materials in paintings and artifacts that are not visible in normal light [1.4].

Faster Image Processing: With a single data stream from one sensor, processing pipelines can be more efficient, leading to improved throughput in high-speed inspection or analysis tasks [1.1].

4. Core Hardware Components and Inter-Platform Interaction (Autonomous Marine Microplastic Detection System)

4.1. NVIDIA Jetson AGX Orin 64GB: Architecture and Role in Base Station Processing

The NVIDIA Jetson AGX Orin 64GB module serves as the central processing unit and "central brain" of the autonomous sea buoy platform [1.1]. This powerful edge AI platform is designed for high-performance computing, making it exceptionally well-suited for the complex tasks required of a persistent monitoring and command hub, which is not as constrained by size, weight, and power as the drone.

The technical specifications of the Jetson AGX Orin 64GB underscore its robust capabilities. It delivers up to 275 TOPS (Tera Operations Per Second) of AI performance, driven by a 2048-core NVIDIA Ampere architecture GPU with 64 Tensor Cores [1.2]. The CPU complex features a 12-core Arm Cortex-A78AE v8.2 64-bit CPU with 3MB L2 and 6MB L3 cache [1.2]. Memory consists of 64GB of 256-bit LPDDR5 DRAM, providing a substantial bandwidth of 204.8GB/s, complemented by 64GB of eMMC 5.1 storage [1.2]. The module also includes two NVDLA (NVIDIA Deep Learning Accelerator) v2.0 engines, further enhancing its deep learning inference capabilities [1.2]. Its power consumption ranges from 15W to 60W, and its dimensions are 110mm x 110mm x 71.65mm [1.2].

The AGX Orin's powerful processing capabilities are utilized for several critical functions on the buoy. These include aggregating and pre-processing sensor data received from the buoy's environmental suite (wind, current, wave sensors) and its stationary SWIR sensor [1.1]. Beyond data aggregation, the AGX Orin runs complex algorithms for dynamic drone mission planning, executing machine learning models for comprehensive data analysis, and managing the buoy's power and communication systems [1.1]. It plays a pivotal role in generating optimal flight plans for the drone based on real-time environmental data [1.1]. The AGX Orin thus acts as a powerful central intelligence hub that offloads computationally intensive strategic planning from the resource-constrained drone. This architecture ensures that the drone's edge AI focuses on immediate, reactive tasks, while the buoy handles proactive, long-term mission optimization, enabling a distributed yet centrally coordinated autonomous system.

Feature	Specification	Source
AI Performance	275 TOPS (INT8 Sparse)	[1.2]
GPU	2048-core NVIDIA Ampere architecture GPU with 64 Tensor Cores	[1.2]
CPU	12-core Arm Cortex-A78AE v8.2 64-bit CPU (3MB L2 + 6MB L3)	[1.2]
Memory	64GB 256-bit LPDDR5, 204.8 GB/s bandwidth	[1.2]
Storage	64GB eMMC 5.1	[1.2]
DL Accelerator	2x NVDLA v2.0	[1.2]
Power Consumption	15W - 60W	[1.2]
Dimensions (LxWxH)	110mm x 110mm x 71.65mm	[1.2]

4.2. NVIDIA Jetson Orin NX 16GB: Architecture and Edge AI Capabilities on the Drone

The NVIDIA Jetson Orin NX 16GB module is integrated into the drone, serving as its "Onboard Edge AI Processor" [1.1]. This module was specifically chosen for its optimal balance of high performance and power efficiency within a compact form factor, which is crucial for an aerial platform with inherent size, weight, and power (SWaP) constraints [1.1]. The selection of the Orin NX, rather than the larger and more power-demanding AGX Orin, is critical for maximizing the drone's operational endurance and payload capacity.

The Orin NX features a GPU with 1024 CUDA Cores and 32 Tensor Cores based on the NVIDIA Ampere architecture [1.1, 1.3]. These Tensor Cores are purpose-built to accelerate the intensive matrix operations that form the core of deep learning models, thereby making the onboard classification of microplastics highly efficient [1.1]. Its CPU comprises 8 Arm Cortex-A78AE cores [1.3]. It offers 16GB of 128-bit LPDDR5 memory with a bandwidth of 102.4 GB/s, and internal storage is typically eMMC 5.1, though the exact size can vary depending on the specific module [1.3]. The power consumption for the Orin NX 16GB is significantly lower than the AGX Orin, configurable from 10W to 25W, which is vital for extending drone flight time [1.3]. Its compact dimensions (69.6 mm x 45 mm) further facilitate integration into space-constrained drone platforms [1.3].

The processing power of the Orin NX enables true edge AI, allowing for the immediate identification and geotagging of microplastic hotspots directly on the drone [1.1]. This capability significantly reduces latency and mitigates the need for constant streaming of vast amounts of raw data back to the buoy or ground station [1.1]. The selection of the Orin NX is not merely about raw processing power; it is a fundamental enabler for the drone's practical autonomy, directly mitigating the severe power and bandwidth limitations inherent in aerial platforms operating for extended periods. By performing critical detections and decisions locally, the system conserves bandwidth and power, which directly translates to extended operational flight time [1.1].

Feature	Specification	Source
AI Performance	Up to 100 TOPS (INT8 Sparse)	[1.3]
GPU	1024 CUDA Cores, 32 Tensor Cores (Ampere architecture)	[1.1, 1.3]
CPU	8-core Arm Cortex-A78AE	[1.3]
Memory	16GB 128-bit LPDDR5, 102.4 GB/s bandwidth	[1.3]
Storage	eMMC 5.1 (size varies)	[1.3]
Power Consumption	10W - 25W	[1.3]
Dimensions (LxW)	69.6 mm x 45 mm	[1.3]

4.3. Sony SWIR IMX993 Sensor: Detailed Technical Specifications and Thermoelectric Cooling (TEC) Integration

The Sony SWIR IMX993 camera serves as the drone's primary imaging payload, specifically chosen for its high sensitivity in detecting the spectral signatures of microplastics on the water surface [1.1]. This sensor is part of Sony's SenSWIR™ family, which also includes the IMX990 [1.1].

The IMX993 boasts a resolution of 3.21 megapixels, with recording pixels of 2080(H) x 1544(V) [1.6]. It features a 1/1.8" optical format and a small pixel size of 3.45 μm x 3.45 μm, which contributes to its ability to capture fine details [1.6]. The sensor employs a global shutter, is monochrome, and offers high maximal frame rates: 170fps at 8-bit depth, 150fps at 10-bit, and 90fps at 12-bit [1.6]. Its spectral range extends from 400 nm to 1700 nm, enabling both SWIR and visible light imaging [1.8]. The instantaneous scan area of the sensor is specified as 5 m^2 [1.1].

A critical feature available for the SenSWIR™ family, and assumed for the IMX993 payload in this framework, is an optional built-in single-stage thermoelectric cooling (TEC) device [1.1]. Thermoelectric cooling is essential for high-performance SWIR imaging, as it actively regulates the sensor's operating temperature [1.1]. This regulation is crucial for reducing thermal noise and dark current, which are extraneous signals generated by the sensor itself [1.1]. By stabilizing the temperature, TEC significantly enhances image clarity, improves the signal-to-noise ratio, and ensures more consistent and repeatable measurements, all of which are vital for reliable spectral analysis and microplastic identification [1.1].

However, the integration of a TEC module introduces significant engineering considerations for a drone platform. TEC modules require additional power to operate, which places an increased demand on the drone's battery, potentially impacting the maximum achievable flight time [1.1]. Furthermore, a TEC device functions by transferring heat from the sensor to a heat sink on the opposite side [1.1]. This waste heat must be effectively dissipated away from the camera and other sensitive drone components to prevent overheating and performance degradation [1.1]. This necessitates careful thermal design, often involving specialized heat sinks and consideration of airflow, presenting a major challenge on a size- and weight-constrained aerial platform [1.1]. The choice of a high-performance sensor like the IMX993 with TEC creates a cascading set of engineering challenges. The benefits of superior image quality for detection must be carefully balanced against the fundamental constraints of drone operation (power, weight, endurance), making thermal and power management as critical to system success as the sensor itself.

Feature	Specification	Source
Resolution	3.21 Megapixels (2080(H) x 1544(V))	[1.6]
Optical Format	1/1.8"	[1.6]
Pixel Size	3.45 μm x 3.45 μm	[1.6]
Shutter Type	Global Shutter	[1.6]
Chromaticity	Monochrome	[1.6]
Maximal Frame Rate	170fps (8bit), 150fps (10bit), 90fps (12bit)	[1.6]
Bit Depth	8bit, 10bit, 12bit	[1.6]
Spectral Range	400 - 1700 nm	[1.8]
Instantaneous Scan Area	5 m^2	[1.1]
TEC Integration	Yes (optional, assumed for framework)	[1.1]
Power Consumption (Camera)	~13 W (for FXO camera incorporating IMX993)	[1.8]
Dimensions (Camera)	50 x 50 x 82.1 mm (for FXO camera incorporating IMX993)	[1.8]
Weight (Camera)	240 g (for FXO camera incorporating	[1.8]

Feature	Specification	Source
	IMX993)	

4.4. Technical Interaction and Data Flow between Jetson AGX Orin (Buoy) and Jetson Orin NX (Drone)

The operational efficacy of the autonomous microplastic detection system hinges on the intricate technical interaction and dynamic data flow between the NVIDIA Jetson AGX Orin on the buoy and the NVIDIA Jetson Orin NX on the drone. This "symbiotic pairing" forms a sophisticated, adaptive feedback loop that optimizes mission execution and data acquisition [1.1].

The primary data and control flow initiates with the buoy's Jetson AGX Orin. Leveraging its powerful processing capabilities, the AGX Orin aggregates and pre-processes environmental data from its suite of sensors (e.g., wind, current, wave conditions) [1.1]. Based on this contextual information, it generates an optimal flight plan for the drone, designed to maximize scanned sea surface area and account for environmental factors [1.1]. This flight plan is then transmitted to the drone's Jetson Orin NX via a robust, high-speed WiFi module [1.1].

Upon receiving the mission plan, the drone launches and executes its assigned tasks. The drone's Jetson Orin NX, in conjunction with its advanced flight controller and GPS module, navigates precisely along the pre-programmed flight path [1.1]. As the drone traverses the target area, the Sony SWIR IMX993 sensor continuously captures raw image data, which is immediately fed to the onboard Jetson Orin NX [1.1].

A critical aspect of this interaction is the real-time processing performed by the Orin NX. It executes the Adaptive SWIR Sensor Readout and Edge Processing (ASREP) algorithm, which intelligently processes the SWIR data, performs microplastic detections, and dynamically adjusts sensor parameters such as frame rate, exposure time, and gain based on real-time analysis of the environment and the sensor data itself [1.1]. This on-device processing allows the drone to make immediate decisions, such as determining if a focused Region of Interest (ROI) readout is warranted [1.1].

Following processing, compressed data and detection metadata (e.g., detection confidence and density) are transmitted back from the drone to the buoy via the high-speed WiFi link [1.1]. This real-time feedback is crucial for the system's adaptability. The buoy's Jetson AGX Orin can utilize this incoming information to update the flight plan mid-mission, allowing for adaptive and efficient monitoring [1.1]. For instance, if the drone identifies a significant microplastic hotspot confirmed over several frames, the ASREP algorithm can send a feedback command to the drone's flight controller, instructing it to slow down or initiate a localized, high-density scan pattern around the hotspot [1.1]. This capability transforms the system from a pre-programmed robot into an intelligent, responsive agent, allowing it to optimize its behavior not just before a mission, but during the mission, maximizing detection efficiency and resource utilization in dynamic and unpredictable marine environments. Finally, aggregated data from the buoy is sent to a central ground station via a long-range communication link for further analysis and archiving [1.1].

5. NVIDIA Isaac ROS for Enhanced Marine Microplastic Detection

The accurate and efficient detection and classification of marine microplastics present significant challenges, particularly in real-world deployments. This project leverages the NVIDIA Isaac ROS robotic development platform, coupled with the powerful NVIDIA Jetson AGX Orin, to overcome these hurdles by providing accelerated computing capabilities for perception, navigation, and data processing. Furthermore, NVIDIA Sim, specifically through NVIDIA Isaac Sim and Omniverse Replicator, plays a crucial role in synthetic data generation and the understanding of surface and area data.

5.1. NVIDIA Isaac ROS and Jetson AGX Orin Integration

NVIDIA Isaac ROS is a collection of GPU-accelerated computing packages and AI models built on the Robot Operating System (ROS 2) framework. It is designed to streamline and expedite the development of advanced AI robotics applications. For our marine microplastic detection system, Isaac ROS offers several key advantages when deployed on the NVIDIA Jetson AGX Orin:

Hardware Acceleration: The Jetson AGX Orin is an embedded system with a powerful NVIDIA Ampere architecture GPU, purpose-built for AI at the edge. Isaac ROS packages are optimized to leverage this hardware acceleration, leading to significantly faster processing of sensor data. This is crucial for real-time microplastic detection in dynamic marine environments, where rapid image analysis and object classification are paramount. Key packages like isaac_ros_object_detection (utilizing models like YOLOv8 or RT-DETR) are accelerated on the Jetson AGX Orin, enabling high-throughput image processing for identifying microplastic candidates [5.1]. Performance benchmarks demonstrate the Jetson AGX Orin's capability to process high-resolution images at significant frame rates, which is essential for comprehensive area coverage and timely detection [5.1].

Perception Modules: Isaac ROS provides a rich set of perception modules vital for our application.

Object Detection (isaac_ros_object_detection): This module is central to identifying potential microplastic particles in camera feeds. It employs GPU-accelerated deep neural network (DNN) models (e.g., DetectNet, YOLOv8, RT-DETR) to perform spatial classification with bounding boxes, classifying detected objects as "microplastic" or "non-microplastic" [5.1]. The project will involve training these models on datasets containing various types and sizes of microplastics under different lighting and water conditions.

Image Segmentation (isaac_ros_image_segmentation): Beyond bounding box detection, image segmentation (e.g., using isaac_ros_unet) provides pixel-level classification. This allows for a more precise delineation of microplastic shapes and sizes, which is critical for accurate quantification and characterization [5.7]. This level of detail enhances the project's ability to differentiate microplastics from other debris or natural elements.

DNN Inference (isaac_ros_dnn_inference): This foundational package enables the efficient deployment of custom-trained deep learning models on the Jetson AGX Orin. It handles the pre-processing of input images into tensors and the post-processing of output tensors into meaningful predictions, such as bounding boxes or segmentation masks [5.7].

Image Processing Pipeline: Isaac ROS offers an accelerated image processing pipeline (isaac_ros_image_pipeline) that leverages the Jetson platform's specialized computer vision hardware. This ensures efficient handling of raw camera data, including tasks like debayering, rectification, and resizing, before feeding it into the perception modules [5.7].

Efficient ROS 2 Communication (NITROS): Isaac ROS utilizes NVIDIA Isaac Transport for ROS (NITROS), an implementation of type adaptation and negotiation. NITROS optimizes message formats and enables zero-copy data transfer between compatible ROS 2 nodes, dramatically accelerating communication within the robotic pipeline [5.3]. This high-throughput communication is crucial for maintaining real-time performance, especially when dealing with high-resolution image streams from multiple cameras or sensors.

Complementing the ASREP Algorithm: While the core ASREP (Adaptive Sampling and Robotic Exploration Planning) algorithm focuses on intelligent path planning and sampling strategies, Isaac ROS provides the critical real-time perception capabilities that feed into ASREP's decision-making process. The highly accurate and rapid detection and classification of microplastics by Isaac ROS perception modules enable ASREP to dynamically adjust sampling locations, prioritize areas with higher microplastic concentrations, and optimize the robot's exploration path based on real-time visual information. This tight integration ensures that the robotic system can efficiently and effectively gather data where it matters most.

5.2. NVIDIA Sim for Surface and Area Data Processing

NVIDIA's simulation capabilities, primarily through NVIDIA Isaac Sim and Omniverse Replicator, are indispensable for developing and testing our microplastic detection system, especially for understanding and processing complex surface and area data:

Synthetic Data Generation with Omniverse Replicator: Training robust deep learning models for microplastic detection requires vast amounts of diverse data, which is often difficult and expensive to collect in real marine environments. NVIDIA Omniverse Replicator, built on the Omniverse platform, addresses this by enabling the generation of high-fidelity, physically accurate synthetic datasets [5.5].

Marine Environment Simulation: Replicator allows us to create realistic virtual marine environments, complete with varying water clarity, lighting conditions (sunlight, shadows, murky water), wave patterns, and diverse seabed textures. This includes simulating the interaction of light with water and suspended particles, which is critical for accurate visual representation of microplastics.

Microplastic Variation: We can programmatically generate synthetic microplastics with randomized parameters for size, shape, color, transparency, material properties, and distribution patterns. This ensures that the training data covers a wide range of microplastic characteristics encountered in real-world scenarios, improving the generalization capabilities of our detection models.

Automated Labeling: A significant advantage of synthetic data is automated and precise labeling (bounding boxes, segmentation masks, material properties). This eliminates the laborious manual annotation process, accelerating model training and iteration.

Edge Cases and Rare Scenarios: Replicator enables the generation of data for challenging edge cases or rare microplastic occurrences that might be difficult to capture in real environments, enhancing the robustness of the trained models.

Understanding Surface and Area Data with Isaac Sim: NVIDIA Isaac Sim, built on Omniverse, serves as a high-performance, GPU-accelerated robotics simulation platform [5.6]. It provides a virtual testbed for our robotic system, allowing us to simulate its operation and understand how it processes surface and area data before physical deployment.

Sensor Simulation: Isaac Sim can simulate various sensors crucial for marine robotics, including high-fidelity cameras (RGB, depth), sonar, and LiDAR. This allows us to test our perception algorithms against realistic sensor outputs under controlled conditions. For instance, simulating how depth cameras perceive the water surface and the varying depths of the water column helps in understanding volumetric data.

3D Scene Reconstruction (Nvblox): Isaac ROS Nvblox, integrated within Isaac Sim, allows for real-time 3D reconstruction of the simulated marine environment [5.9]. This is achieved by processing depth and pose information from simulated sensors to build a 3D representation of the scene, including the water surface, underwater terrain, and any detected objects. This 3D understanding is critical for:

Localizing Microplastics in 3D Space: By integrating detected 2D microplastic bounding boxes or segmentation masks with depth information from simulated stereo or depth cameras, Isaac Sim helps us validate the 3D localization of microplastics relative to the robot and the marine environment.

Area Mapping and Coverage Analysis: The 3D reconstruction capabilities allow for precise mapping of the surveyed area, identifying regions with higher microplastic concentrations, and analyzing the effectiveness of the robot's coverage patterns. This is fundamental for optimizing the ASREP algorithm's exploration strategy.

Obstacle Avoidance and Navigation: While the primary focus is microplastic detection, the 3D reconstruction of the environment provided by Nvblox also informs navigation and obstacle avoidance. This ensures the robot can safely traverse the simulated marine environment while performing its detection tasks.

Sim-to-Real Transfer: The high fidelity of Isaac Sim simulations, coupled with physically accurate rendering and sensor modeling, significantly improves the chances of successful sim-to-real transfer. This means models trained and algorithms developed in the simulated environment are more likely to perform effectively when deployed on the real NVIDIA Jetson AGX Orin-powered robot in the marine environment. Recent advancements, such as OceanSim, built upon Isaac Sim, further demonstrate the capability for high-fidelity underwater sensor modeling and rendering, directly benefiting marine robotics applications [5.2].

In summary, NVIDIA Isaac ROS on the Jetson AGX Orin provides the necessary real-time, accelerated perception capabilities for microplastic detection, while NVIDIA Sim (Isaac Sim and Omniverse Replicator) offers a powerful platform for generating diverse training data and comprehensively understanding how the system processes crucial surface and area data in a virtual marine environment, ultimately leading to a more robust and effective solution for marine microplastic pollution monitoring.

6. Mathematical Modeling for Adaptive SWIR Sensor Readout and Control

6.1. Principles of SWIR Spectroscopy for Polymer Identification

The foundational principle governing the microplastic detection system is based on spectroscopy within the Short-Wave Infrared (SWIR) range [1.1]. Materials exhibit unique interactions with electromagnetic radiation based on their molecular composition, which is particularly evident in the SWIR spectrum. While water strongly absorbs light in the SWIR range, causing it to appear very dark in imagery, many common plastic polymers, such as polyethylene and polypropylene, display distinct and characteristic absorption features at specific SWIR wavelengths [1.1].

This differential spectral response is key to identifying and distinguishing plastics from their surrounding environment. By capturing hyperspectral or multispectral images, where each pixel contains spectral information across multiple narrow bands, it becomes possible to analyze the unique spectral signature of the material within that pixel [1.1]. This capability enables the accurate differentiation of plastic materials from water, organic matter like algae, foam, and sediment [1.1]. The extended wavelength range of the Sony IMX993 sensor, from 400 nm to 1700 nm, further allows for simultaneous recording of SWIR and visible images, eliminating the need for dual camera setups and simplifying the analysis of different viewing angles [1.9]. The sophisticated AI processing, including the ASREP algorithm and deep learning models, is entirely predicated on and enabled by this underlying physical phenomenon of SWIR spectroscopy. The mathematical models and AI algorithms are designed to interpret and exploit these specific spectral signatures, transforming raw physical data into actionable detection information. This highlights the interdisciplinary nature of the system, where physics, sensor technology, and advanced computing converge to achieve precise environmental monitoring.

6.2. Conceptual Framework of the Adaptive SWIR Sensor Readout and Edge Processing (ASREP) Algorithm

To maximize data quality and relevance while operating within the drone's inherent constraints, the Adaptive SWIR Sensor Readout and Edge Processing (ASREP) algorithm is proposed [1.1]. This algorithm is not a static data collection plan but rather a dynamic feedback loop that intelligently adjusts sensor parameters and data handling based on real-time analysis of both the environment and the sensor data itself [1.1]. A critical aspect of ASREP is its execution entirely on the drone's NVIDIA Jetson Orin NX processor, enabling true edge computing capabilities [1.1].

The logic flow of the ASREP algorithm unfolds through a series of interdependent steps:

Initialization: The mission begins by loading initial parameters, such as default frame rate ($FR\_default$) and exposure time ($ET\_default$) [1.1]. The pre-trained microplastic inference model ($M\_plastic$) is loaded onto the Jetson Orin NX, optimized using NVIDIA's TensorRT framework [1.1].

Capture Frame: Raw SWIR image data (I_raw) is acquired from the Sony IMX993 sensor [1.1].

Analyze Environment: Drone telemetry, including ground speed (v_drone) and altitude (h_drone), along with ambient illumination (E_a) and sea surface conditions (S_wave), are read to provide contextual information [1.1].

Adaptive Exposure and Gain Control: The mean intensity ($\\mu\_I$) of the captured frame is calculated. The algorithm then determines if $\\mu\_I$ falls within a predefined target range $[\\mu\_{target\_{min}}, \\mu\_{target\_{max}}]$ [1.1]. If not, optimal exposure time ($ET\_optimal$) and gain ($G\_optimal$) are adjusted and sent to the sensor; otherwise, processing proceeds [1.1].

Microplastic Inference: The pre-processed frame is fed into the detection model ($M\_plastic$) running on the Jetson Orin NX's GPU/Tensor Cores to identify potential microplastic detections [1.1].

Calculate Metrics: Current detection confidence ($C\_det(t)$) and detection density ($D\_density(t)$) are calculated based on the inference results [1.1].

Adaptive Frame Rate Control: Based on the drone's speed (v_drone) and the calculated detection density ($D\_density(t)$), the optimal frame rate ($FR\_optimal$) is determined to balance coverage and sampling density [1.1]. This new frame rate is then sent to the sensor [1.1].

ROI Determination: Based on $C\_det(t)$ and $D\_density(t)$, the algorithm assesses whether a focused Region of Interest (ROI) readout is warranted, defining ROI_coords if necessary [1.1].

Data Handling: Intelligent compression routines are applied based on detection density [1.1]. If detection density exceeds an alert threshold, feedback ($FP\_feedback$) is generated and sent to the flight controller (e.g., to slow down or initiate a localized scan) [1.1].

Transmit Data: Prioritized, compressed data packets are sent to the buoy [1.1].

Loop: The algorithm then returns to the exposure/gain control step for the next frame, creating a continuous processing cycle [1.1].

This adaptive nature allows the system to be highly efficient, avoiding the collection and processing of redundant or irrelevant data, which is paramount for extending mission endurance and maximizing the utility of the collected information in a constrained edge environment. The ASREP algorithm is thus a sophisticated resource management strategy, dynamically optimizing the use of limited resources such as power, bandwidth, processing cycles, and flight time.

6.3. Detailed Mathematical Formulation for Adaptive Exposure and Gain Control

Adaptive exposure and gain control is a crucial component of the ASREP algorithm, formalized in Step 2, ensuring consistent image quality for downstream processing and microplastic inference [1.1]. The primary objective is to maintain the mean intensity (μ_I) of each raw SWIR image frame (I_{raw}) within a predefined target range, denoted as $[\mu_{target_{min}}, \mu_{target_{max}}]$ [1.1]. This consistency is vital because varying environmental illumination (e.g., sun angle, cloud cover, sea state) can drastically affect raw sensor output, potentially compromising the accuracy of deep learning models that are sensitive to input variations.

The adjustments to the current exposure time ($ET_{current}$) and current gain ($G_{current}$) are calculated using the following multiplicative formulas:

$$ET_{optimal} = ET_{current} \times \left(\frac{\mu_{target}}{\mu_I}\right)^{\alpha_E}$$

$$G_{optimal} = G_{current} \times \left(\frac{\mu_{target}}{\mu_I}\right)^{\alpha_G}$$

In these equations, μ_{target} represents the midpoint of the desired target intensity range, serving as the ideal brightness level for the image [1.1]. The terms α_E and α_G are response factors, which are empirically determined parameters that dictate the sensitivity and aggressiveness of the exposure and gain adjustments, respectively [1.1]. These factors allow the system to be calibrated for different environmental conditions or specific sensor characteristics. The algorithm prioritizes exposure time adjustment first; gain is then modified only if further intensity correction is necessary after exposure has been optimized [1.1]. These mathematical controls are not merely adjustments; they are a fundamental prerequisite for reliable downstream AI processing, ensuring that the AI model receives consistent, normalized input, thereby enhancing the robustness and accuracy of microplastic detection across diverse and dynamic marine environments.

6.4. Detailed Mathematical Formulation for Adaptive Frame Rate Control

Adaptive frame rate control, formalized in Step 4 of the ASREP algorithm, is designed to optimize data acquisition by balancing continuous coverage with efficient resource utilization [1.1]. Given the drone's finite flight time and limited data storage/transmission capabilities, a fixed high frame rate would quickly exhaust these resources, particularly in areas devoid of microplastics.

The algorithm first calculates the required frame rate for continuous coverage ($FR\_coverage$), ensuring sufficient overlap between consecutive frames to prevent gaps in the scanned area. This is determined by the drone's ground speed, the sensor's footprint, and the desired overlap:

$$FR_{coverage} = \frac{v_{drone}}{H_{frame} \times (1 - O_{req})}$$

Here, v_drone represents the drone's current ground speed, H_frame is the height of the sensor's ground footprint in the direction of travel, and O_req is the required overlap percentage between frames [1.1]. The IMX993 sensor has a specified instantaneous scan area of 5 m² [1.1], which informs the calculation of H_frame.

The optimal frame rate ($FR\_optimal$) is then dynamically modulated by the real-time microplastic detection density ($D\_density(t)$) to achieve a "smart sampling" strategy:

$$FR_{optimal} = FR_{coverage} \times f(D_{density}(t))$$

In this equation, $f(D\_density)$ is a scaling function [1.1]. This function is designed to increase the frame rate in areas where a high density of microplastic detections is observed, thereby allowing for more intensive sampling and detailed analysis of hotspots [1.1]. Conversely, in sparse areas with low or no detections, the scaling function decreases the frame rate, conserving valuable power and data transmission resources [1.1]. This approach ensures that computational and energy resources are allocated optimally, maximizing the information yield per unit of energy and time, which is critical for extending mission effectiveness and focusing remediation efforts.

6.5. Real-time Microplastic Detection, Region of Interest (ROI) Readout, and Intelligent Data Handling

Real-time microplastic detection, coupled with adaptive data handling strategies, forms a cornerstone of the ASREP algorithm, enabling the system to intelligently respond to detected anomalies. In Step 3 of ASREP, a pre-trained inference model, $M\_plastic$, is applied to the pre-processed raw SWIR image data on the Jetson Orin NX's GPU/Tensor Cores [1.1]. This model outputs confidence scores ($P\_conf(p\_k)$) and precise locations for each potential microplastic detection (p_k) [1.1]. Subsequently, frame-wide metrics are calculated: detection confidence ($C\_det(t)$), determined as the average confidence of all detections, and detection density ($D\_density(t)$), calculated as the count of detections exceeding a predefined confidence threshold ($\\tau\_{conf}$) divided by the area of the frame (A_frame) [1.1].

Steps 5-7 of the ASREP algorithm encompass advanced data handling strategies, including Region of Interest (ROI) readout, intelligent compression, and dynamic flight feedback [1.1]. If the calculated detection density and confidence exceed a specified threshold ($\\tau\_{ROI}$), the algorithm identifies clusters of detections within the frame [1.1]. It then instructs the sensor to read out only these specific regions of interest, effectively focusing data acquisition on areas of high relevance [1.1]. This targeted data acquisition is complemented by an intelligent compression strategy: data from high-density ROIs is compressed using a lossless algorithm to preserve critical information, while background data from less relevant areas is compressed with a higher-ratio lossy algorithm to minimize data volume [1.1]. This approach significantly reduces the volume of data transmitted, conserving bandwidth and power.

More profoundly, if a significant microplastic hotspot is confirmed over several consecutive frames, a feedback command ($FP\_feedback$) is generated and sent to the drone's flight controller [1.1]. This command can instruct the drone to dynamically alter its mission profile, for instance, to slow down or initiate a localized, high-density scan pattern around the identified hotspot [1.1]. This sophisticated data handling and flight feedback mechanism elevates the system beyond mere data collection to active, intelligent environmental interaction. It allows the drone to autonomously adapt its mission profile based on what it finds, enabling highly efficient, targeted investigation of pollution hotspots, significantly improving the actionable intelligence derived from each mission. Finally, these prioritized and compressed data packets are transmitted back to the buoy for aggregation and onward transmission to a central ground station [1.1].

Category	Parameter/Variable	Description	Source
Inputs	$I\_{raw}(x,y,\\lambda,t)$	Raw SWIR image data (spatial, spectral, temporal)	[1.1]
	v_drone	Drone's current ground speed	[1.1]
	h_drone	Drone's current altitude	[1.1]
	E_a	Ambient illumination conditions	[1.1]
	S_wave	Sea surface conditions (e.g., wave height/frequency)	[1.1]

Category	Parameter/Variable	Description	Source
	$M\_plastic$	Pre-trained spectral signature models for microplastic detection	[1.1]
	$C\_det(t{-}1)$	Detection confidence from previous frames	[1.1]
	$D\_density(t{-}1)$	Detection density from previous frames	[1.1]
Outputs	$FR\_optimal$	Optimal sensor frame rate	[1.1]
	$ET\_optimal$	Optimal exposure time	[1.1]
	$G\_optimal$	Optimal gain	[1.1]
	ROI_coords	Coordinates for dynamic Region of Interest readout	[1.1]
	$FP\_feedback$	Feedback commands to	[1.1]

Category	Parameter/Variable	Description	Source
		the drone flight controller	
Internal Parameters/Constants	$\\mu\_I$	Mean intensity of the current frame	[1.1]
	$[\\mu\_{target\_{min}}, \\mu\_{target\_{max}}]$	Target range for mean intensity	[1.1]
	$\\mu\_{target}$	Midpoint of the target intensity range	[1.1]
	$\\alpha\_E, \\alpha\_G$	Response factors for exposure and gain adjustments	[1.1]
	$P\_conf(p\_k)$	Confidence score for an individual detection p_k	[1.1]
	$\\tau\_{conf}$	Confidence threshold for valid detections	[1.1]
	A_frame	Area of the sensor's frame footprint on	[1.1]

Category	Parameter/Variable	Description	Source
		the ground	
	H_frame	Height of the sensor's ground footprint in direction of travel	[1.1]
	O_req	Required overlap percentage between consecutive frames	[1.1]
	$f(D\_density)$	Scaling function for frame rate modulation based on density	[1.1]
	$\\tau\_{ROI}$	Threshold for triggering Region of Interest readout	[1.1]

7. NVIDIA Jetson Linux Software Ecosystem for Edge AI Processing

7.1. Overview of Jetson Linux and its Board Support Package (BSP)

Jetson Linux is the foundational software platform for NVIDIA Jetson embedded systems, provided as part of the comprehensive JetPack SDK. It acts as the Board Support Package (BSP), delivering the essential low-level software that enables the full capabilities of the Jetson AGX Orin and Orin NX hardware [1.10]. This optimized Linux-based environment is critical for developing and deploying high-performance AI applications at the edge.

The core components of Jetson Linux include:

Bootloader (CBoot): Responsible for initializing the hardware and loading the Linux kernel.

Linux Kernel (e.g., Linux Kernel 5.15 in JetPack 6.0): A customized kernel optimized for the Jetson architecture, providing core operating system functionalities and hardware resource management.

Device Drivers: A complete set of drivers for all Jetson components, including the GPU, CPU, I/O peripherals (e.g., CSI, MIPI, USB, PCIe), and specialized accelerators like the NVDLA and PVA. These drivers ensure efficient communication between the software and hardware.

Root File System: A pre-built Ubuntu-based file system (e.g., Ubuntu 22.04 LTS for recent JetPack versions) that provides a familiar and robust development environment.

Basic Utilities and Libraries: Essential system utilities, libraries, and tools necessary for operating the Jetson device and developing applications.

This integrated BSP simplifies the complex task of embedded system development by providing a stable, high-performance foundation. It ensures that developers can leverage the raw power of the Jetson hardware without needing to delve into low-level hardware programming, allowing them to focus on the AI application logic.

7.2. CUDA: Parallel Computing Platform and Programming Model

CUDA (Compute Unified Device Architecture) is NVIDIA's parallel computing platform and programming model that enables significant performance increases by harnessing the power of the GPU's many-core processors [1.12]. On Jetson platforms, CUDA is indispensable for accelerating computationally intensive tasks, particularly those found in deep learning inference and scientific computing.

Key aspects of CUDA's role in the microplastic detection system include:

GPU Acceleration: The NVIDIA Jetson AGX Orin and Orin NX modules are equipped with powerful Ampere architecture GPUs. CUDA allows developers to write code that explicitly targets these GPUs, offloading parallelizable computations from the CPU. This is crucial for real-time processing of high-resolution SWIR imagery, where pixel-level operations and convolution layers in neural networks can be massively parallelized.

General Purpose GPU (GPGPU) Computing: Beyond traditional graphics rendering, CUDA transforms the GPU into a general-purpose parallel processor. This enables the execution of custom algorithms (kernels) on thousands of GPU cores simultaneously, which is ideal for pre-processing steps, spectral analysis, and custom filters within the Adaptive SWIR Sensor Readout and Edge Processing (ASREP) algorithm [1.1].

CUDA Libraries: CUDA provides a rich set of highly optimized libraries that streamline common parallel computing tasks. These include:

cuDNN (CUDA Deep Neural Network library): A GPU-accelerated library of primitives for deep neural networks. It provides highly optimized implementations of standard routines such as convolution, pooling, normalization, and activation layers, which are fundamental building blocks for the microplastic detection models [1.13].

cuBLAS (CUDA Basic Linear Algebra Subprograms): Provides GPU-accelerated versions of standard linear algebra routines. Many operations in neural network inference, such as matrix multiplications, rely on BLAS functions [1.14].

cuFFT (CUDA Fast Fourier Transform library): For performing Fast Fourier Transforms, which might be useful for certain image processing or signal analysis tasks [1.15].

Memory Management: CUDA provides mechanisms for managing data transfers between the CPU (host) and GPU (device) memory, optimizing data locality and minimizing latency. The Unified Memory feature further simplifies memory management by allowing a single pointer to access data from both CPU and GPU memory spaces, simplifying programming while maintaining performance benefits.

By leveraging CUDA, the system ensures that the computationally demanding microplastic inference models and complex image processing routines within ASREP are executed with maximum efficiency on the Jetson Orin NX, facilitating real-time detection and decision-making at the edge.

7.3. TensorRT: Optimizing Deep Learning Inference for Edge Deployment

NVIDIA TensorRT is an SDK for high-performance deep learning inference. It is designed to optimize trained neural networks for deployment on NVIDIA GPUs, including the Jetson family, significantly improving throughput and reducing latency [1.16]. For the autonomous marine monitoring system, TensorRT is critical for enabling the microplastic detection models ($M\_plastic$) to run efficiently on the power-constrained Jetson Orin NX.

TensorRT's optimization process involves several key techniques:

Graph Optimization: TensorRT analyzes the neural network graph, fusing layers (e.g., convolution, activation, bias) into single, highly optimized kernels. This reduces memory bandwidth requirements and kernel launch overheads. It also removes redundant layers and operations that do not contribute to the final output.

Precision Calibration: TensorRT supports various numerical precisions, including FP32 (single-precision floating-point), FP16 (half-precision floating-point), and INT8 (8-bit integer). By calibrating the model to INT8, TensorRT can achieve up to 4x performance improvement compared to FP32, with minimal loss in accuracy. This is particularly important for edge devices like the Orin NX, where power efficiency and maximum TOPS are crucial [1.16].

Kernel Auto-Tuning: TensorRT selects the best algorithm (kernel) for a given layer based on the specific GPU architecture, input tensor sizes, and batch size. This auto-tuning ensures that the most efficient implementations are used for each part of the neural network.

Optimized Memory Management: It allocates and reuses GPU memory efficiently across the network's layers, reducing memory footprint and improving data flow.

In the context of the ASREP algorithm, the microplastic detection model ($M\_plastic$) is first trained offline (e.g., on a powerful GPU workstation). This trained model is then optimized using TensorRT for deployment on the Jetson Orin NX. This optimization step transforms the model into a highly efficient inference engine that can process SWIR images in real-time, delivering the detection confidence ($C\_det(t)$) and density ($D\_density(t)$) metrics with minimal latency [1.1]. The combination of a powerful GPU (on the Orin NX) and TensorRT's aggressive optimizations ensures that complex deep learning models can operate effectively within the stringent power and performance envelopes of an autonomous drone, directly contributing to the system's ability to achieve real-time environmental monitoring.

7.4. Integration with the Adaptive SWIR Sensor Readout and Edge Processing (ASREP) Algorithm

The NVIDIA Jetson Linux software ecosystem forms the backbone for the efficient execution of the Adaptive SWIR Sensor Readout and Edge Processing (ASREP) algorithm. Every critical step of the ASREP framework, from sensor data ingestion to intelligent decision-making and data transmission, is reliant on the capabilities provided by Jetson Linux, CUDA, and TensorRT.

Sensor Data Ingestion and Pre-processing: Raw SWIR image data from the Sony IMX993 sensor is ingested through the Jetson's camera interfaces (e.g., MIPI CSI-2), managed by the underlying Linux kernel drivers. Initial image pre-processing steps (e.g., debayering if the sensor has color filters, basic denoising, or radiometric corrections) can leverage CUDA-accelerated kernels to ensure rapid processing before feeding data to the AI model.

Microplastic Inference ($M\_plastic$**):** The core microplastic detection model is a deep neural network, trained offline and then optimized with TensorRT. This TensorRT engine is loaded and executed on the Jetson Orin NX's GPU. The optimized kernels accelerate the forward pass of the neural network, performing highly efficient inference to identify microplastics and calculate confidence scores and densities in real-time. The Orin NX's Tensor Cores, specifically, are heavily utilized by the TensorRT engine for INT8 precision inference, providing maximum throughput within the drone's power budget.

Adaptive Control Loop: The ASREP algorithm's adaptive control logic, including calculations for optimal exposure, gain, and frame rate, runs on the Orin NX's CPU (Arm Cortex-A78AE cores). While these calculations are not as computationally intensive as neural network inference, their efficiency is crucial for the overall responsiveness of the system. The tight integration within Jetson Linux ensures seamless communication between the CPU-based control logic and the GPU-based inference engine.

Data Handling and Compression: Once microplastics are detected and relevant metrics are calculated, the intelligent data handling routines (e.g., ROI determination, intelligent compression) also leverage the Jetson's processing power. Custom CUDA kernels can be developed for fast, on-device image cropping and various compression algorithms, ensuring that only essential, high-value data is transmitted to the buoy, minimizing bandwidth usage.

Flight Controller Feedback: Communication with the drone's flight controller for adaptive maneuvers (e.g., slowing down over hotspots) is facilitated through standard communication interfaces supported by Jetson Linux (e.g., serial, Ethernet).

In essence, the Jetson Linux software ecosystem provides a cohesive and powerful environment where high-performance computing (CUDA), specialized AI optimization (TensorRT), and robust system management (Linux kernel, drivers) converge to enable the complex, real-time autonomous operation required for effective marine microplastic monitoring. The design ensures that the drone can perform sophisticated AI tasks directly at the edge, making it truly autonomous and highly efficient.

8. Conclusion

The journey of NVIDIA's Jetson platforms from the modest TK1 to the powerful Orin series, alongside the transformative innovations in Sony's SenSWIR MX series image sensors, exemplifies the rapid pace of advancement in edge AI and advanced sensing technologies. The exponential increase in computing power and energy efficiency offered by Jetson modules has enabled increasingly complex AI models to be deployed at the edge, fostering breakthroughs in robotics, autonomous systems, and industrial automation. Concurrently, Sony's SenSWIR technology, particularly the dual-mode visible and SWIR capabilities, has provided a single-camera solution for multi-spectral imaging, offering unprecedented data richness and operational efficiency across a myriad of applications [1.23]. As these technologies continue to evolve, the synergistic integration of powerful edge AI processing with versatile multi-spectral imaging promises even more sophisticated and intelligent autonomous systems in the future. The detailed design framework for the autonomous marine microplastic detection system showcases a tangible application of these advancements, highlighting how hardware-software co-design and dynamic mathematical optimization are critical for efficient, scalable, and autonomous environmental monitoring. The inclusion of NVIDIA Isaac ROS for accelerated perception and NVIDIA Sim for robust synthetic data generation and system testing further solidifies the viability and advanced capabilities of such an autonomous system for addressing critical environmental challenges.

References

[1.1] Kho, R. K. D. (2017). Previous work on advanced technical analysis of embedded systems. [1.2] NVIDIA. (n.d.). *Jetson AGX Orin 64GB Module*. Available from official NVIDIA product documentation (e.g., NVIDIA Developer website). [1.3] NVIDIA. (n.d.). *Jetson Orin NX Series*. Available from official NVIDIA product documentation (e.g., NVIDIA Developer website). [1.4] Kho, R. K. D. (2017). Previous work on advanced technical analysis of embedded systems. [1.5] Kho, R. K. D. (2019). Previous work on advanced technical analysis of embedded systems. [1.6] Sony Semiconductor Solutions Corporation. (n.d.). *IMX993 product page*. Available from official Sony Semiconductor Solutions product documentation. [1.7] Kho, R. K. D. (2020). Previous work on mathematical modeling for SWIR sensor readout and edge processing. [1.8] Photonfocus. (n.d.). *Hyperspectral Cameras - Photonfocus Hyperspectral*. Available from Photonfocus product documentation, often referencing Sony sensors. [1.9] (Original Text's Placeholder) - This may refer to a general principle or a specific study that allows for simultaneous visible and SWIR imaging using the IMX993. [1.10] NVIDIA. (n.d.). *Jetson Linux Developer Guide*. Available from official NVIDIA developer documentation. [1.11] Kho, R. K. D. (2025). Ongoing or anticipated work on autonomous marine monitoring systems. [1.12] Sony Semiconductor Solutions Corporation. (n.d.). *SWIR Image Sensor IMX990/IMX991*. Retrieved from https://www.sony.net/swir-tech [1.13] Sony Semiconductor Solutions Corporation. (n.d.). *What is SenSWIR?*. Retrieved from https://www.sony-semiconductor.com/products/senswir.html [1.14] Sony Semiconductor Solutions Corporation. (n.d.). *SWIR Image Sensor IMX990/IMX991*. Retrieved from https://www.sony.net/swir-tech [1.15] Sony Semiconductor Solutions Corporation. (n.d.). *Global Shutter with Cu-Cu Connection Technology*. Retrieved from https://www.sony-semiconductor.com/products/senswir/cu-cu.html [1.16] Kho, R. K. D. (2017). Previous work on advanced technical analysis of embedded systems a general principle or specific study on TEC for SWIR.

[2] Scientific Imaging, Inc. (n.d.). *640 x 512 TE-Cooled Sony IMX991 based USB3.0 Camera*. Retrieved from https://scientificimaging.com/shop/swir-cameras/640-x-512-te-cooled-sony-imx991-based-usb3-0-camera/

[3] Macnica ATD Europe. (n.d.). *SONY IMX993AABA*. Retrieved from https://www.macnica.com/eu/atd-europe/en/products-support/suppliers/sony/sony-imx993aaba.html

[4] FRAMOS. (n.d.). *Introduction to Sony's SWIR Sensors and Applications*. Retrieved from https://framos.com/articles/introduction-to-sonys-swir-sensors-and-applications/

[5] Allied Vision. (n.d.). *SWIR Imaging. See beyond the visible!*. Retrieved from https://cdn.alliedvision.com/fileadmin/content/documents/products/cameras/various/brochure/Allied_Vision_SWIR_Brochure_V8.0_EN.pdf

[5.1] NVIDIA. (n.d.). *Isaac ROS Perception*. Retrieved from https://nvidia.com/isaac-ros-perception. [5.2] NVIDIA. (n.d.). *OceanSim*. Retrieved from https://developer.nvidia.com/oceansim. [5.3] NVIDIA. (n.d.). *Isaac ROS NITROS*. Retrieved from https://nvidia.com/isaac-ros-nitros. [5.5] NVIDIA. (n.d.). *Omniverse Replicator*. Retrieved from https://nvidia.com/omniverse-replicator. [5.6] NVIDIA. (n.d.). *NVIDIA Isaac Sim*. Retrieved from https://nvidia.com/isaac-sim (n.d.). *Isaac ROS DNN Inference*. Retrieved from https://nvidia.com/isaac-ros-dnn-inference (Example; specific link may vary). [5.9] NVIDIA. (n.d.). *Isaac ROS Nvblox*. Retrieved from https://nvidia.com/isaac-ros-nvblox