DEMOCRATIC AND POPULAR REPUBLIC OF ALGERIA
MINISTRY OF HIGHER EDUCATION AND SCIENTIFIC RESEARCH
UNIVERSITY OF 08 MAY 1945 - GUELMA

Graduation Thesis
For Master's Degree

Field : Computer Science
Option : STIC : Science And Technology Of Information And Communication

Topic

From Pixels to Polymers : A Deep-Hybrid Spectral Framework for Precise Plastic Detection Using Hyperspectral Imaging and AI

Presented by :

Mohammed CHAIB
Abderrahman BARKACHE
Achraf YAHIAOUI
Wassim MESSIOUD

Supervised by :

Dr. Samir HALLACI

Co-Supervised by :

Pr. Marco Balsi

PhD. Soufyane Bouchelaghem

Academic Year : 2024/2025

Acknowledgements

We would like to express our deep gratitude to all those who contributed to the completion of this work.

We extend our sincere thanks to our supervisor, **Dr. Samir HALLACI**, for his guidance, valuable advice, and continuous support throughout this project.

Our gratitude also goes to the entire faculty of the University of 08 May 1945 - Guelma for their dedication and knowledge, which have greatly enriched our academic journey.

This project was carried out within the framework of the **National Program 1275**, aimed at fostering innovation and entrepreneurship among students. We are especially grateful for the collaboration with the eminent professors from Sapienza University of Rome, **Prof. Marco BALSI** and **PhD. Soufyane BOUCHELAGHEM**, whose insightful exchanges and support contributed to the quality and relevance of this work.

We would also like to express our sincere gratitude to the **members of the jury** for the honor they grant us by agreeing to examine this work. Their presence reflects the interest they have in our research and their commitment to promoting scientific rigor. We warmly thank them for the time they will devote to evaluating this thesis, as well as for the richness of their remarks, constructive criticisms, and insightful suggestions, which will undoubtedly enhance and deepen the scope of this study.

A special thanks to our families for their unwavering support, encouragement, and patience during this challenging period. Their belief in us has been a constant source of motivation.

Finally, we acknowledge our friends and colleagues who have provided both moral and technical support, making this journey even more rewarding.

CHAIB Mohammed's Dedication

Praise be to Allah, Lord of the Worlds, for granting me the strength, patience, and clarity to reach this academic milestone. By His grace, I conclude a journey of 17 years of education — a journey filled with challenges, growth, and unforgettable memories that have shaped the person I am today.

I am deeply grateful to my family. To my father and mother — thank you for raising me with love, discipline, and the values that shaped who I am. Your unwavering support through every challenge, your belief in me even during moments of doubt, and your encouragement in both triumph and hardship have meant everything to me.

To my dear siblings — **Zaki**, **Rahma**, **Meriem**, and **Marwa** — you have always been pillars of strength and sources of wisdom. Your guidance, affection, and inspiration have played a vital role in my journey.

My gratitude also extends to my entire extended family — both those still with us and those who have passed on — for their love, prayers, and constant presence in my life.

Throughout this path, I have had the privilege of learning, evolving, and collaborating with exceptional individuals. I express my highest and most sincere gratitude to my mentor and supervisor, **Dr. Samir Hallaci**, whose wisdom, vision, and unwavering support have been instrumental in the success of this thesis. Your guidance has left a profound impact on both my academic and personal journey.

To my brothers-in-arms — the **SM/ART Team** : **Yahiaoui Achraf**, **Messioud Wassim**, and **Abderrahmane Barkach** — it was an honor to work alongside you. Your intelligence, dedication, and friendship made this final project one of the most rewarding experiences of my life.

I extend my heartfelt appreciation to my classmates and all my fellow students who shared this academic path with me. Special thanks go to my dear friends : **Mohammed Biri**, **Amine Kouadria**, **Ben Ziyadi Nabil Iskandar**, **Kheireddine Chiaoui**, **Sarah Doughman**, **Maya Bouzidi (jinx)**, and **Hala Hammouchi** — thank you for your support, friendship, and the many moments we shared.

I am deeply grateful to all the professors who guided me throughout this scientific journey. I would especially like to thank **Pr. Kouahla Nadjib**, **Pr. Benabdellah Ahsan**, **Pr. Hanous Hakim**, **Pr. Farou Brahim**, and **Pr. Abdelmoumen Hiba** for their dedication and valuable teaching.

A special note of thanks to the leadership of the Department of Computer Science. To **Pr. Kouahla Zine Eddine**, Head of Department — thank you for your commitment and for facilitating access to the essential resources that enabled our scientific projects. I also extend my thanks to **Dr. Khaled Halimi** for his technical support and for helping deploy our work to the university server.

And to the Eagles : Islam, Amjad, Hachim, Ahmed, and Mohammed Aouamri thanks for all the good vibes and unforgettable moments. You made this journey even better.

Finally, I express my deepest appreciation to **Prof. Marco Balsi** and **Prof. Piotr Rynkowski**, whose insights, presence, and contributions enriched this project. Meeting and learning from you has been a true honor.

To every person who played a role in this journey — whether by teaching, supporting, encouraging, or simply believing in me — thank you. This achievement is yours as much as it is mine.

YAHIAOUI Achraf's Dedication

Almighty Allah, thank You for always being by my side. I dedicate this thesis to the dearest beings in my heart.

First, to my parents especially my Mom, source of love, generosity, and affection. She who taught me to value hard work and good conduct, who supported and encouraged me throughout all these years of study. May she find here the expression of my gratitude and my deep love.

To my sister Ibtissam, for her patience and support. I wish you every success, joy, and happiness.

To my brother Abdennour, may God protect you. I wish you all the happiness you deserve, my dear brother. I love you very much.

To my dear Uncle Kamel, for his kindness, unwavering support, and steadfast belief in me.

To my beloved grandparents, for their unconditional love, heartfelt prayers, and the values they instilled in me. May this humble work be a reflection of my deepest gratitude and respect.

I would like to express my deepest gratitude to my supervisor, Mr. Hallci, for his constant availability, his invaluable guidance, and the numerous opportunities he offered me throughout this journey. His unwavering support, insightful advice, and remarkable human and moral qualities have greatly contributed to the success of this work and have left a lasting impact on me.

I would like to express my heartfelt gratitude to Pr. Seridi Hamid, head of the laboratory, and Miss Madiha for their unwavering support and for granting us daily access to the lab without any obstacles. My thanks also go to Pr. Kouahla Zine Eddine, head of the department, as well as to Pr. Kouahla Nadjib, Pr. Farou Brahim, Dr. Halimi Khaled, and all the professors and staff of the department for their guidance and encouragement throughout this journey.

A special note of appreciation goes to Pr. Marco Balsi, whose precise and insightful guidelines greatly facilitated the progress of our project. Your expertise and structured approach had a significant impact on our work.

I am also profoundly thankful to Pr. Piotr Rynkowski for the memorable and enriching moments we shared. Working with you was not only an academic privilege but also a deeply human experience from which I gained invaluable lessons.

Shoutout to my teammates from SM/ART team Mohammed Chaib, Wassim, Abderrahmane, and Mayah (aka Jinx), working with you was such a blast.

Big thanks to my friends Mohammed Biri, Amine Kouadria, nabil, Amine Semsar and Sarra Doughman for always being there.

And to the Eagles : Islam, Amjad, Hachim, Ahmed, and Mohammed Aouamri thanks for all the good vibes and unforgettable moments. You made this journey even better.

To all the students of my class—thank you for the good times we shared together. And to everyone who has helped me, whether from near or afar, throughout my studies.

MESSIOUD Wassim's Dedication

First and foremost, I would like to thank God for the destiny He has written for me, for all the strength, patience, and blessings that have guided me throughout my life.

I am deeply grateful to my family — to my father and my mother, who raised me with love and values, who have always been there for me, supporting me through every challenge, believing in me even in moments of doubt, and encouraging me through both highs and lows.

I also wish to sincerely thank my older brothers : Mouhamed, Samir, and Chouayb. You have always been pillars of support and sources of inspiration. Your maturity and guidance have shaped the person I am becoming.

My gratitude extends to my entire extended family — both those who are still with us and those who have passed on — for their love and their presence in my life.

A special thanks goes to Dr. **Samir Hallaci**, for all the knowledge he has shared with us, for the invaluable opportunities we experienced thanks to him, for guiding us throughout this project and many others, and for the lessons not only in academics but also in life and humanity.

I would also like to thank Pr. **Seridi Hamid**, the head of the laboratory, and Miss. **Madiha** for allowing us to work in the lab every day without any issues. I extend my gratitude as well to Pr. **Kouahla Zine Eddine**, the department head, and to Pr. **Kouahla Nadjib**, Pr. **Farou Brahim**, Dr. **Halimi Khaled**, and all the professors and staff of the department.

I would like to express my sincere thanks to Pr. **Marco Balsi** for providing us with clear and helpful guidelines that greatly facilitated our project work. Your expertise and structure made a real difference.

I am also deeply thankful to Pr. **Piotr Rynkowski** for the unforgettable moments we shared. It was a truly enriching experience, both academically and personally, and I learned many valuable lessons through our time together.

To my teammates in the **SM/ART team** — Mouhamed, Achraf, and Abdou — thank you for the unforgettable moments we shared, from joy and excitement to stress and struggle. We accomplished so much together and built memories that I will carry with me forever.

To my dear colleagues — Amine Samsar, Nabil, Amine Kouadria, Mouhamed Biri, Jinx, and many others — thank you for your friendship, support, and companionship throughout this journey.

To the **FAUCON** group — Islem, Amjed, Hachim, Mouhamed Aouamri, and Ahmed — I am grateful for all the moments we've shared.

A heartfelt thank you to **Tonton le P** for all the delicious meals he prepared for us over the years.

To my childhood friends — Aymen, Mehdi, Lokman, Marwan, Benou, Jaber, and Mouha — thank you for more than 15 years of true friendship, for all the memories we've made, and for always being there. I trust and value each one of you more than words can express.

Last but not least, I want to thank me for believing in me, for never giving up, for staying determined and continuing to chase my dreams to the very end.

I want to thank everyone who has helped me in any way — whether near or far, whether through a gesture or a word, and all those who have loved me. I will be forever grateful.

BARKACHE Abderrahman's Dedication

I thank Allah for helping me in this work and this wonderful specialization, alongside all my colleagues in study and support.

I extend my heartfelt gratitude to my entire family,my uncles and my cousins, especially my mother and father, who are my unwavering support in life.

To my dear partners—Achraf, Mohamad, Wassim—for the success of this project : we worked hand in hand, sharing ideas, skills, and determination.

To our **AI House** boss, and my professor and supervisor **Hallaci Samir**, for standing by me in every situation—good and bad—with exceptional mentorship, patience, and inspiring guidance that shaped my growth.

I also express my sincere gratitude to all my favorite teachers who left a lasting impact on me : **Seridi Hamid**, **Farou Brahim**, the **Kouahla brothers**, **Hadjris Morad**, **Halimi Khaled**, **Benabdellah Ahcene**, **Kefali Abderrahman**, **Djakhjakha Linda**, **Mehnaoui Zahra**, **Abdelmoumen Hiba**, **Boursasse Hassina**, and **Hanousse Hakim**. Your knowledge, kindness, encouragement, and dedication have truly inspired me throughout my academic journey.

A special word of thanks to **Pr. Marco Balsi** for his invaluable co-supervision, for the exceptional opportunity he granted me to collaborate on this project, and for recommending me for further academic pursuits. His trust, guidance, and vision have been instrumental in shaping the quality and ambition of this work.

I am also deeply grateful to **Pr. Piotr Rynkowski**, to whom I had the pleasure of presenting a side project. His positive feedback and recognition of the value of our main project were truly encouraging. Beyond the academic exchange, I treasure the wonderful moments and lasting memories shared with him—an experience that enriched both my professional and personal journey.

To my friends : Dhiya, Taky, Iyed, Khaled, Didou, Mbarek, Said, Nafaa, Salim, Yousef, Ali, Redwan,Islam,Bouakaz brothers,Doudou,Amjed,Mossab... your camaraderie means the world.

To all students of our cohort, especially : Mohamed Chaib and Biri, Achraf, Wassim, Amine Samsar and Kouadria, Nabil, Ilyes, Ayoub, Khayrou... your collaboration made this journey memorable.

I deeply thank every member of **Teckgeeks** for their brilliant insights, tireless teamwork, and innovative spirit that elevated our computer science department.

To all who contributed, directly or indirectly, to my academic path : I express my profound appreciation for your invaluable support throughout my studies.

الملخص

أصبحت مسألة الكشف عن النفايات البلاستيكية وتصنيفها من الأولويات الحرجة في ظل تزايد المخاوف البيئية، لا سيما في النظم البيئية المائية والبرية. إن تقنيات التصوير التقليدية والفحص اليدوي غالبًا ما تفتقر إلى الدقة والكفاءة وقابلية التوسع اللازمة للتعامل مع هذه التحديات في ظروف العالم الحقيقي الديناميكية. وتبرز تقنية التصوير الطيفي الفائق HSI كحل واعد، نظرًا لقدرتها على التقاط التواقيع الطيفية التفصيلية عبر مئات النطاقات الضيقة من الأطوال الموجية، مما يتيح التمييز بين أنواع البوليمرات البلاستيكية والعناصر الطبيعية.

ومع ذلك، فإن التطبيق العملي لهذه التقنية يواجه عدة عراقيل، من بينها التعقيد الحسابي، والتكرار الطيفي، والتغيرات البيئية، والندرة في البيانات الموسومة. في هذا البحث، نقترح إطار عمل شامل وتكيفي يجمع بين تقنيات التعلم العميق والتعلم الآلي الكلاسيكي للتغلب على هذه الحواجز البيئية والحسابية.

نقترح ونقيّم أربع منهجيات متكاملة: (١) تحليل المركبات الرئيسية PCA يليه مصنف آلة المتجهات الداعمة (SVM) لتصنيف طيفي فعال منخفض الأبعاد؛ (٢) نموذج الغابة العشوائية (Random Forest) باستخدام الإحصائيات الكاملة للنطاقات لتعزيز المتانة وقابلية التفسير ؛ (٣) شبكة عصبية تلافيفية (2D U-Net) قادرة على إجراء تقسيم طيفي-مكاني مباشر دون الحاجة إلى تقليل الأبعاد؛ و(٤) خط أنابيب هجين مكوّن من مرحلتين يدمج بين تقسيم قائم على U-Net لصور RGB مع تصنيف طيفي باستخدام SVM لتحقيق تمييز دقيق لأنواع البوليمرات.

تم اختبار هذه المنهجيات على مجموعة من قواعد البيانات الطيفية الواقعية، بما في ذلك صور جوية مأخوذة بواسطة طائرات بدون طيار وتجارب مخبرية في قنوات مائية مضبوطة، وتتميز بخلفيات بيئية مختلفة (مثل الماء، والرمل، والطحالب) وظروف إضاءة متعددة وأنواع متنوعة من البوليمرات.

أظهرت النتائج أن الخط الهجين يتفوق على النماذج المنفردة من خلال الاستفادة من الدقة المكانية للشبكات التلافيفية والحساسية الطيفية للمصنفات القائمة على النواة. فقد حققت مرحلة التقسيم دقة F1 بنسبة 99.42%، ودقة كلية بلغت 99.35%، ودرجة تداخل (IoU) بلغت 98.75%، ومتوسط دقة (mAP) قدره 87.91%، مع خسارة تحقق منخفضة لا تتجاوز 0.016. أما في مرحلة التصنيف، فقد سجل مصنف SVM دقة كلية بلغت 91.5% عبر ستة أنواع من البوليمرات: PET و HDPE و LDPE و PP و EPSF والبلاستيك المتقادم، مع دقة كاملة بنسبة 100% لكل من EPSF و LDPE.

يتمكن الإطار الكامل من معالجة 10 ملفات خلال 12.7 ثانية فقط باستخدام وحدة معالجة الرسوميات (GPU)، مما يؤكد إمكانياته للتطبيق في الوقت الحقيقي ضمن الأنظمة المدمجة. وتثبت هذه الدراسة جدوى الكشف الآلي عن النفايات البلاستيكية في البيئات الميدانية والمخبرية، وتفتح آفاقًا لتطبيقات قابلة للتوسع في المراقبة الجوية والفضائية، والبنى التحتية الذكية لإعادة التدوير، ومنصات الاستشعار البيئي المدمجة.

Abstract

The detection and classification of plastic waste have become critical priorities amid escalating environmental concerns, particularly in aquatic and terrestrial ecosystems. Conventional imaging and manual inspection techniques often lack the accuracy, efficiency, and scalability required to address these challenges in dynamic, real-world conditions. Hyperspectral imaging (HSI), with its ability to capture detailed spectral signatures across hundreds of narrow wavelength bands, emerges as a promising solution for identifying and distinguishing various plastic polymers from natural elements. However, the practical application of HSI is hindered by computational complexity, spectral redundancy, environmental variability, and the limited availability of labeled data.

This thesis presents a comprehensive, adaptive framework that synergizes deep learning and classical machine learning techniques to overcome both environmental and computational barriers. We propose and evaluate four complementary methodologies : (1) a Principal Component Analysis (PCA) followed by Support Vector Machine (SVM) classifier for efficient low-dimensional spectral classification ; (2) a Random Forest model utilizing full-band statistics to enhance robustness and interpretability ; (3) a 2D U-Net convolutional neural network capable of direct spectral-spatial segmentation without prior dimensionality reduction ; and (4) a hybrid two-phase pipeline that couples U-Net-based RGB segmentation with hyperspectral SVM classification to achieve fine-grained polymer-level discrimination.

These methods are benchmarked on a range of diverse, real-world hyperspectral datasets, including aerial drone imagery and controlled flume-based laboratory acquisitions, which feature varying environmental backgrounds (e.g., water, sand, algae), lighting conditions, and polymer types. Results demonstrate that the hybrid pipeline outperforms standalone models by leveraging the spatial precision of convolutional architectures and the spectral sensitivity of kernel-based classifiers. Specifically, the segmentation phase achieved an F1-score of 99.42%, accuracy of 99.35%, Intersection over Union (IoU) of 98.75%, and mean Average Precision (mAP) of 87.91%, with a validation loss of only 0.016. In the classification phase, the SVM attained an overall precision of 91.5% across six polymer types—PET, HDPE, LDPE, PP, EPSF, and weathered plastics, with perfect (100%) precision for EPSF and LDPE.

The full framework processes 10 files in just 12.7 seconds on GPU, highlighting its potential for real-time deployment in embedded systems. This work confirms the feasibility of automated plastic detection in both field and lab environments and lays the groundwork for scalable applications in UAV and satellite-based monitoring, smart recycling infrastructures, and embedded environmental sensing platforms.

Résumé

La détection et la classification des déchets plastiques sont devenues des priorités critiques face à l'aggravation des enjeux environnementaux, en particulier dans les écosystèmes aquatiques et terrestres. Les techniques d'imagerie conventionnelle et d'inspection manuelle manquent souvent de précision, d'efficacité et de capacité à passer à l'échelle pour traiter ces défis dans des conditions réelles complexes. L'imagerie hyperspectrale (HSI), grâce à sa capacité à capturer des signatures spectrales détaillées sur des centaines de bandes étroites, s'impose comme une solution prometteuse pour identifier et distinguer divers polymères plastiques des éléments naturels. Toutefois, son application pratique est limitée par la complexité computationnelle, la redondance spectrale, la variabilité environnementale et la rareté des données annotées.

Ce travail propose un cadre méthodologique complet et adaptatif combinant l'apprentissage profond et les techniques classiques de machine learning afin de surmonter ces obstacles à la fois environnementaux et computationnels. Nous présentons et évaluons quatre approches complémentaires : (1) une classification spectrale efficace par ACP suivie d'un classifieur SVM ; (2) un modèle Random Forest exploitant l'ensemble des bandes spectrales pour une prise de décision robuste ; (3) un réseau U-Net 2D permettant une segmentation spectro-spatiale directe sans réduction de dimension ; et (4) une méthode hybride en deux phases combinant segmentation RGB par U-Net et classification hyperspectrale par SVM pour une discrimination fine des polymères.

Ces méthodes sont testées sur plusieurs jeux de données hyperspectrales issus de contextes réels : imagerie par drone, bassins expérimentaux, conditions lumineuses variées, arrière-plans complexes (eau, sable, algues) et différents types de plastiques. Les résultats montrent que l'approche hybride surpasse les modèles isolés en combinant la précision spatiale des architectures convolutionnelles avec la sensibilité spectrale des classifieurs à noyaux. En particulier, la segmentation atteint un F1-score de 99,42%, une précision de 99,35%, un IoU de 98,75%, un mAP de 87,91% et une perte de validation de seulement 0,016. Pour la classification, le SVM obtient une précision globale de 91,5% sur six types de polymères — PET, HDPE, LDPE, PP, EPSF et plastiques altérés — avec un taux de précision parfait (100%) pour EPSF et LDPE.

L'ensemble du pipeline traite 10 fichiers en seulement 12,7 secondes sur GPU, ce qui en fait une solution adaptée aux systèmes embarqués en temps réel. Ce travail confirme la faisabilité de la détection automatisée des déchets plastiques dans des environnements de terrain et de laboratoire, et jette les bases d'applications évolutives pour la surveillance par drones ou satellites, les infrastructures de recyclage intelligentes et les plateformes de suivi environnemental embarquées.

Table of Contents

List of Tables

List of Figures

General Introduction

Plastics are among the most widely used materials in modern society, found in everything from food containers, shopping bags, and bottles to electronics, furniture, and household goods [6, 33]. Their widespread adoption is largely due to their versatility, low cost, light weight, corrosion resistance, and durability [11, 26]. However, these very properties that make plastics economically attractive also render them environmentally persistent. The mass production, uncontrolled usage, and inherently low biodegradability of plastics have led to their accumulation in nearly every corner of the planet, from city streets and public parks to beaches, rivers, and oceans contributing not only to visible pollution but also to elevated carbon emissions and serious health concerns [8, 20]. As plastic materials have become indispensable across countless industries, the global production of plastics has soared and with it, the volume of unmanaged plastic waste. This has escalated plastic pollution into one of the most pressing environmental challenges of our time, posing a severe threat to ecosystems, marine biodiversity, and human health, and making the detection and classification of plastic debris an urgent priority for sustainable development.

Traditional methods of plastic waste identification, typically relying on manual sorting or conventional imaging are labor-intensive, error-prone, and often incapable of keeping pace with the scale of pollution. These limitations highlight the need for smarter, faster, and more reliable technologies. In this context, *Hyperspectral Imaging (HSI)* has emerged as a revolutionary tool. Unlike standard RGB imaging, HSI captures rich spectral information for each pixel, enabling the precise differentiation of materials based on their unique spectral signatures. This capability is particularly effective in distinguishing between various plastic polymers such as *Polypropylene (PP)*, *Polyethylene (PE)*, *Polystyrene (PS)*, *High-Density Polyethylene (HDPE)*, *Low-Density Polyethylene (LDPE)*, and *Polyvinyl Chloride (PVC)*.

To address these challenges, this thesis proposes a comprehensive, adaptive framework that integrates both classical machine learning and modern deep learning methods to enhance the accuracy, robustness, and automation of plastic detection using hyperspectral data. The proposed framework comprises four complementary approaches, each designed to overcome specific environmental and computational barriers. The first approach combines Principal Component Analysis (PCA) for spectral dimensionality reduction with a Support Vector Machine (SVM) classifier to enable efficient, low-dimensional classification. The second approach employs a Random Forest algorithm to provide robustness to noisy labels. The third approach introduces a 2D U-Net convolutional neural network capable of direct pixel-wise spectral-spatial segmentation, bypassing the need for prior feature extraction. Finally, a novel two-phase hybrid method is developed, where in a U-Net model is first used to perform segmentation on RGB-projected hyperspectral images, followed by an SVM applied to the original spectral data to accurately classify plastic polymers.

These methodologies are validated using a variety of hyperspectral datasets collected under diverse conditions, including aerial drone acquisitions and controlled laboratory environments with varying backgrounds such as water, sand, and vegetation. By combining the strengths of machine learning and deep learning techniques with the spectral richness of hyperspectral data, this research lays the groundwork for intelligent, automated systems capable of addressing plastic pollution in both natural and industrial settings. The ultimate goal of this work is to develop a scalable solution that not only improves detection accuracy but also facilitates real-time deployment in platforms such as unmanned aerial vehicles (UAVs), satellite-based monitoring systems, and smart recycling infrastructures.

Structure of the Thesis :

- **Chapter 1 : Introduction** — Introduces the research context, defines the problem, outlines the objectives, and presents an overview of the methodology.
- **Chapter 2 : State of the Art** — Reviews existing work in hyperspectral imaging, plastic detection, and deep learning-based classification systems.
- **Chapter 3 : System Design and Implementation** — Describes the data collection process, model architectures, segmentation and classification strategies, and training pipeline.
- **Chapter 4 : Results and Discussion** — Presents and analyzes experimental results, evaluates performance metrics, and discusses the strengths and limitations of the approach. Summarizes the contributions of the thesis and proposes future research directions for scaling and improving the system.

— **Annex :** — As part of our graduation project under Ministerial Resolution 1275, this section addresses the strategic and entrepreneurial dimensions of our work. It includes detailed technical documentation, supplementary experiments and visualizations. Furthermore, it presents a market study and business model proposal, emphasizing the project's potential for innovation, scalability, and real-world impact.

By merging artificial intelligence with remote sensing technologies, this research contributes to a growing body of work aimed at addressing one of the planet's most urgent environmental problems. The findings of this thesis offer promising insights for future advancements in intelligent waste management and large-scale plastic pollution monitoring.

CHAPTER 01

INTRODUCTION

1 Background

Plastic pollution has escalated into a global crisis, with projections forecasting a dramatic surge in plastic waste generation over the coming decades. Over the past 70 years, global plastic production has grown exponentially—from just two million tonnes in 1950 to over 516 million tonnes today. While plastic has brought immense benefits to industries due to its affordability, durability, and sterility, its widespread and often indiscriminate use has led to serious environmental consequences [w1].

Figure 1.1 shows the estimated per capita plastic waste entering the ocean in 2019. The data reflects domestic emissions per country, measured in kilograms per person, and excludes exported waste, which may pose an even higher risk of ocean leakage. The visualization is based on estimates by Meijer et al. (2021), as processed by Our World in Data [w2].

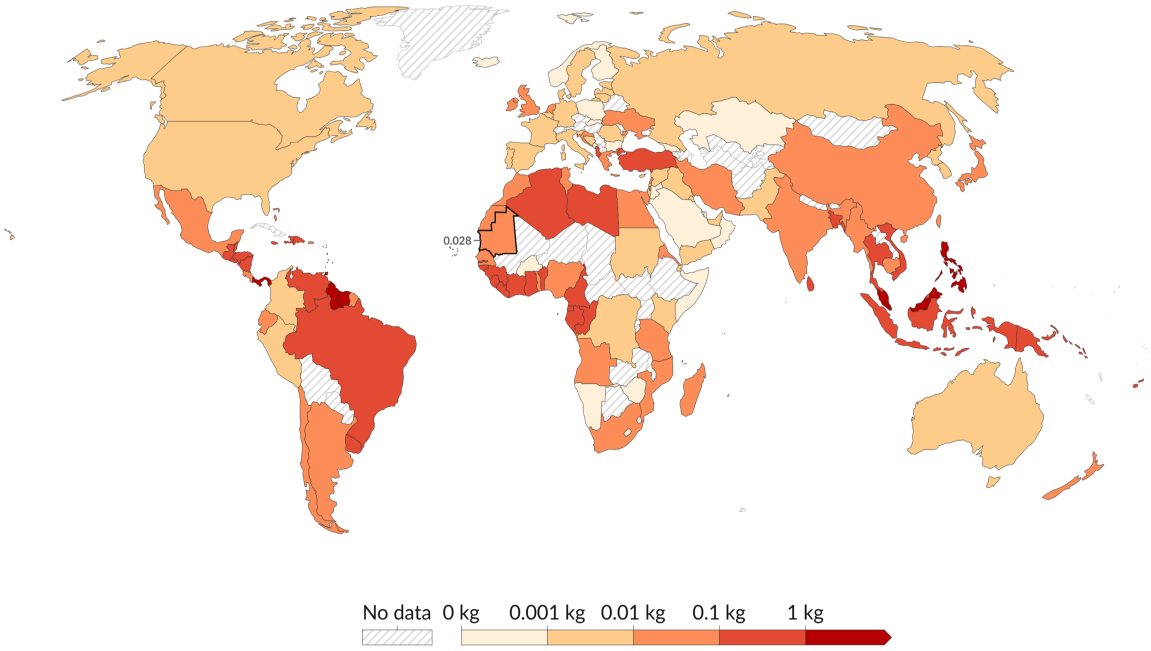

Figure 1.1 – Plastic waste emitted to the ocean per capita (2019).
Estimated domestic emissions per country, excluding exported waste. Data source : Meijer et al. (2021)[w1].

Without decisive intervention, global plastic leakage into the environment is expected to double, reaching an alarming 44 million tonnes annually. The total accumulation of plastic waste in aquatic environments—lakes, rivers, and oceans—is projected to more than triple, increasing from 460 million tonnes in 2019 to over 1,014 million tonnes by 2060. This surge threatens marine biodiversity, disrupts food chains, and exacerbates pollution in ecosystems already under stress.

A key contributor to this crisis is macroplastics—large plastic debris such as bottles, bags, and packaging that degrades very slowly in nature. Meanwhile, microplastics—synthetic polymer particles smaller than 5 mm—are emerging as a growing concern. These originate from industrial plastic pellets, synthetic textiles, tire wear, and the breakdown of larger plastic items, and have been found in marine organisms, drinking water, and even human tissues [w2].

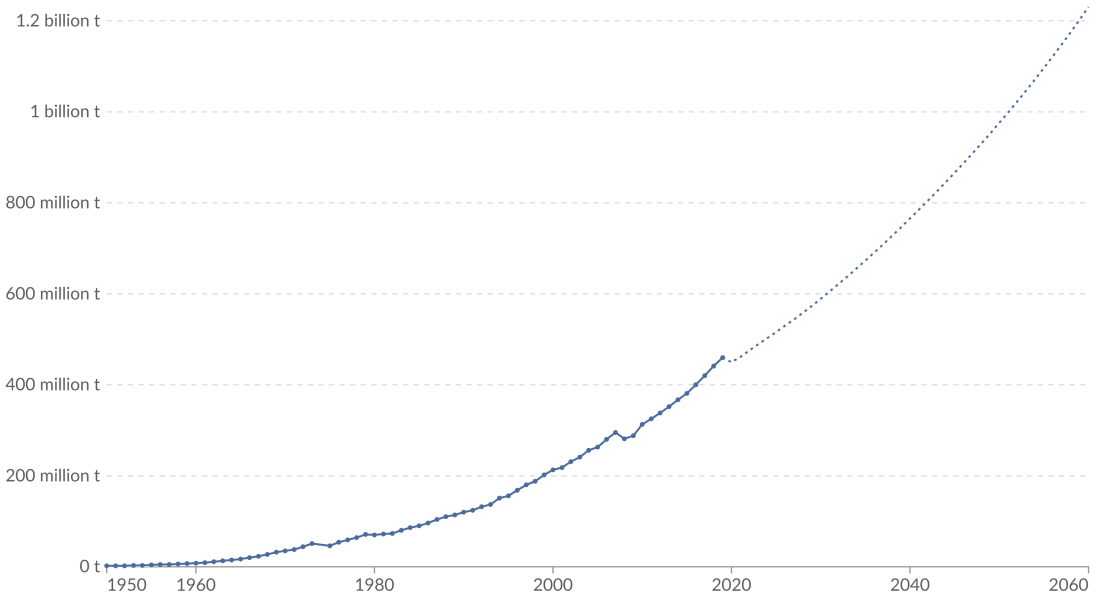

Figure 1.2 – Global plastic production and future projections, 1950–2060.
Annual production of polymer resins and fibers. *Projections follow a "business-as-usual" scenario assuming no major policy changes[w2].*

As illustrated in **Figure** 1.2, global plastic production has shown a steady upward trajectory, and forecasts indicate a significant rise if current consumption patterns continue. These trends are supported by Geyer et al. (2017) and the OECD Global Plastics Outlook (2022), which warn that, in the absence of systemic changes, plastic waste generation will reach unprecedented levels [17].

Plastic Transport Pathways and Current Detection Challenges

A significant portion of plastic waste enters marine environments via **river systems**, which serve as primary conduits, channeling land-based debris into coastal waters and ultimately into the open ocean [w3].

Community-led initiatives to collect plastics directly from rivers, as illustrated in **Figure 1.3**, underscore the urgency of mitigating pollution at its source before it reaches fragile marine ecosystems [w3]. Meanwhile, **Figure 1.4** presents diverse manifestations of ocean plastic pollution, emphasizing the widespread impact across various marine and coastal environments.

Once plastic waste reaches aquatic ecosystems, efforts to detect and remove it are hindered by **traditional detection methods** that rely on manual sorting or conventional imaging. These approaches are often labor-intensive, time-consuming, and prone to human error—factors that severely limit their scalability and accuracy in large-scale or dynamic marine conditions.

This growing crisis necessitates innovative, automated solutions for plastic waste detection and recycling. Traditional recycling technologies often struggle to distinguish chemically similar plastic polymers, reducing the overall effectiveness of waste sorting systems. To address this limitation, advanced techniques such as **Hyperspectral Imaging (HSI)** coupled with deep learning have emerged as powerful tools for accurate and scalable polymer classification. By capturing both spectral and spatial information, these methods can significantly enhance identification performance.

In this study, we propose a novel two-phase framework for automatic plastic waste detection and classification based on HSI. In the first phase, a U-Net-based deep learning model is employed to segment plastic debris from complex backgrounds with high precision. In the second phase, the segmented regions are fed into a Support Vector Machine (SVM) classifier to identify six major plastic polymers : **Polystyrene (PS)**,

Figure 1.3 – Human intervention to collect plastic waste from rivers in an effort to prevent marine pollution [w3].

Figure 1.4 – Examples of ocean plastic pollution in diverse marine and coastal contexts [w3].

Polypropylene (PP), **Polyvinyl Chloride (PVC)**, **High-Density Polyethylene (HDPE)**, **Low-Density Polyethylene (LDPE)**, and **Polyethylene Terephthalate (PET)**.

By combining the strengths of deep learning for semantic segmentation with the robust decision boundaries of machine learning classification, the proposed framework improves overall accuracy, robustness, and generalization. Furthermore, the approach addresses critical challenges such as spectral overlap, spatial heterogeneity, and class imbalance, making it suitable for real-time deployment in various environmental scenarios. As such, this work contributes to the development of intelligent, automated recycling systems that are essential for combating plastic pollution and promoting circular economy principles.

The Seven Types of Plastic

The Society of the Plastics Industry (SPI) categorizes plastics into seven distinct types, each identified by a standardized recycling code. These classifications, illustrated in **Figure 1.5**, are essential for facilitating the sorting, recycling, and effective management of plastic waste, thereby supporting global efforts toward

sustainability and environmental protection [w4].

Figure 1.5 – The seven standard types of plastic, categorized by recycling codes [w4].

2 Overview of Remote Sensing

Remote sensing is a multidisciplinary scientific field focused on the observation, measurement, and interpretation of Earth's surface and phenomena without direct physical contact. It plays a pivotal role in environmental monitoring by capturing data through the detection of electromagnetic radiation reflected or emitted from terrestrial objects and surfaces. Typically, remote sensing systems are deployed via space-borne satellites or airborne platforms, enabling large-scale, high-resolution data acquisition across extensive spatial and temporal domains [2].

These systems are broadly classified based on the origin of the detected radiation : **passive** or **active**. Passive remote sensing depends on naturally occurring energy sources, predominantly solar radiation, which is either reflected by surfaces (in the visible to near-infrared spectrum) or emitted as thermal infrared radiation. Such systems operate effectively under clear-sky conditions and during daylight hours. Common passive sensors include optical cameras, multispectral scanners, and microwave radiometers.

Conversely, active remote sensing involves sensors that emit their own energy toward a target and measure the returned signal. This approach enables continuous data collection regardless of lighting or weather conditions, making it suitable for all-weather and night-time applications. Examples include Radio Detection and Ranging (RADAR) and Light Detection and Ranging (LiDAR), which are widely employed in topographical mapping, vegetation analysis, and environmental assessment.

The effectiveness of a remote sensing system is determined by four key resolutions [2] :

— **Spatial resolution** – the smallest discernible unit on the ground,
— **Spectral resolution** – the sensor's ability to detect fine spectral variations,
— **Radiometric resolution** – the sensitivity to differences in signal intensity,
— **Temporal resolution** – the frequency of data acquisition over time.

These parameters dictate the sensor's suitability for specific applications, ranging from land-use classification and disaster management to environmental pollution monitoring. Of particular relevance to this thesis is the use of **hyperspectral imaging (HSI)**, a passive sensing modality, for the identification and classification of plastic materials based on their unique spectral signatures.

As global environmental challenges intensify, the integration of advanced remote sensing technologies with artificial intelligence offers significant potential for automated, accurate, and scalable environmental monitoring and decision-making systems.

3 Conceptual Framework for a Remote Marine Litter Detection System

Marine litter, particularly plastic debris, has emerged as a critical environmental concern due to its persistence, mobility, and ecological impact on aquatic ecosystems. Effectively addressing this issue requires large-scale, frequent, and reliable monitoring systems capable of detecting and characterizing litter across diverse spatial and temporal scales. Remote sensing offers a powerful solution for such monitoring by enabling non-contact observation of marine environments via platforms such as satellites, aircraft, and unmanned aerial vehicles (UAVs).

As illustrated in Figure 1.6, the conceptual framework for a remote marine litter detection system is typically structured into several core stages. These include data acquisition, where sensors capture imagery or spectral data from the ocean surface ; preprocessing, which involves atmospheric correction, geometric calibration, and noise reduction ; and target detection, where algorithms are applied to isolate potential litter from the background. Following detection, classification techniques are employed to categorize detected objects, often based on spectral, spatial, or textural features. Finally, the outputs are analyzed and interpreted to support environmental decision-making and policy development.

This framework highlights the importance of integrating multi-platform and multi-resolution data sources to improve detection accuracy under varying environmental conditions. Furthermore, it underscores the critical role of sensor selection—particularly regarding spatial, spectral, and radiometric resolution—in influencing overall system performance. While traditional multispectral and optical systems provide valuable insights into surface features, their relatively coarse spectral resolution may limit their capacity to discriminate between materials with similar reflectance characteristics, such as different types of plastics.

Hyperspectral imaging (HSI), which captures continuous spectral information across hundreds of narrow and contiguous bands, offers a substantially richer dataset that enables precise material identification through unique spectral signatures. This increased spectral granularity makes HSI a highly promising tool for advanced environmental monitoring tasks, including the detailed detection and classification of marine litter and plastic debris.

4 Introduction to Hyperspectral Imaging (HSI)

4.1 Hyperspectral Imaging Technology

Hyperspectral Imaging (HSI), also known as imaging spectrometry, is an advanced remote sensing technology that enables the detection and identification of materials based on their unique spectral properties. The concept of HSI was first introduced in the 1980s by Alexander F. H. Goetz and colleagues at NASA's Jet Propulsion Laboratory (JPL) through the development of the Airborne Imaging Spectrometer (AIS). This pioneering instrument was designed to acquire laboratory-quality reflectance spectra from airborne platforms, laying the foundation for modern imaging spectrometry [22].

While multispectral imaging systems—typically capturing data in 3 to 10 broad spectral bands—have been widely used since the 1970s, HSI marks a significant advancement by collecting data across hundreds of narrow, contiguous spectral bands. Each pixel in a hyperspectral image contains a complete reflectance spectrum, offering detailed insights into the physical and chemical composition of surface materials. This higher spectral resolution allows for more accurate material classification and target detection than traditional multispectral approaches [31].

Hyperspectral data is typically represented as a three-dimensional cube, with two spatial dimensions (x, y) and one spectral dimension (λ). This structure enables the examination of spectral signatures at the pixel

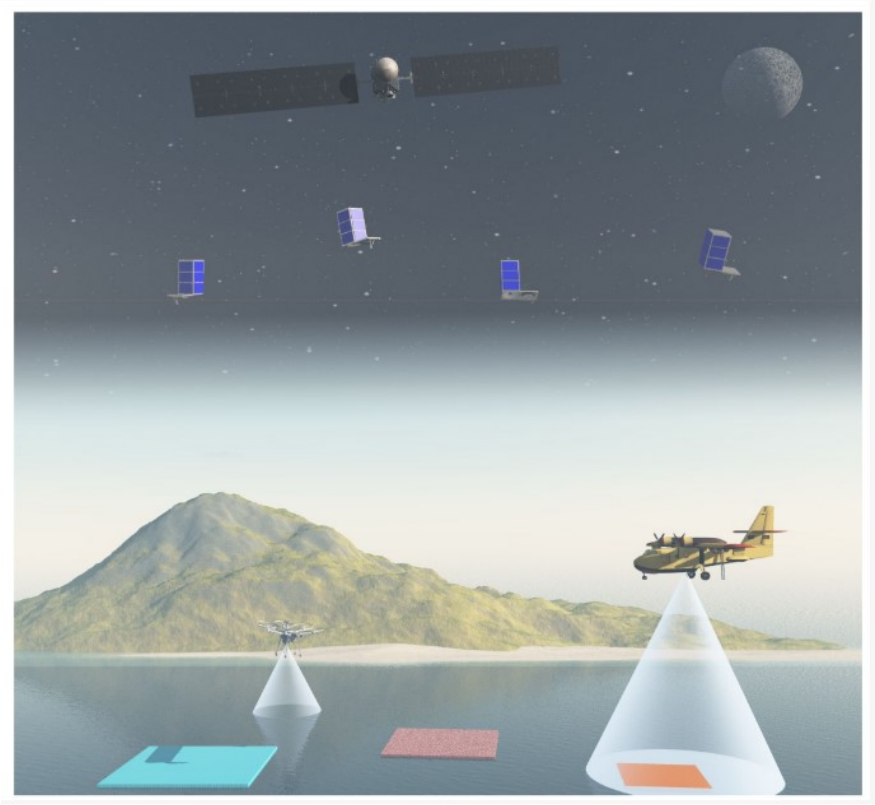

Figure 1.6 – Conceptual framework for a remote marine litter detection system [16].

level, allowing researchers to detect subtle variations in reflectance associated with material properties. The spectral resolution, defined as the smallest discernible difference in wavelength—facilitates the identification of specific features such as absorption bands that are vital for various applications, including vegetation monitoring, plastic waste detection, mineral exploration, and water quality assessment.

Despite its potential, hyperspectral imaging remains underutilized in many operational remote sensing contexts. However, the growing availability of airborne and spaceborne HSI sensors—such as NASA's AVIRIS (Airborne Visible/Infrared Imaging Spectrometer) and ESA's upcoming missions, has renewed interest in its application across environmental and industrial sectors.

The exceptional capabilities of HSI sensors allow for the acquisition of imagery in which each pixel is represented as a high-dimensional spectral vector. This unprecedented spatial-spectral resolution has enabled a wide array of applications in both civilian and military domains. These include agricultural assessment, land use mapping, ecological monitoring, mineral exploration, detection of anthropogenic materials, change detection, surveillance, ground cover classification, and target recognition. All of these applications are based on the fundamental principle that materials scatter electromagnetic energy in distinctive patterns across wavelengths, directly linked to their molecular composition [14].

Understanding the principles, advantages, and limitations of hyperspectral imaging is essential for advancing remote sensing research and applications. With ongoing technological advancements enhancing sensor performance and reducing operational costs, HSI is poised to become a cornerstone in the next generation of remote sensing systems.

4.2 Scope of Hyperspectral Imaging

Hyperspectral Imaging (HSI) has emerged as a powerful and versatile tool for a wide range of remote sensing applications due to its capacity to capture high-resolution spectral information across hundreds of narrow,

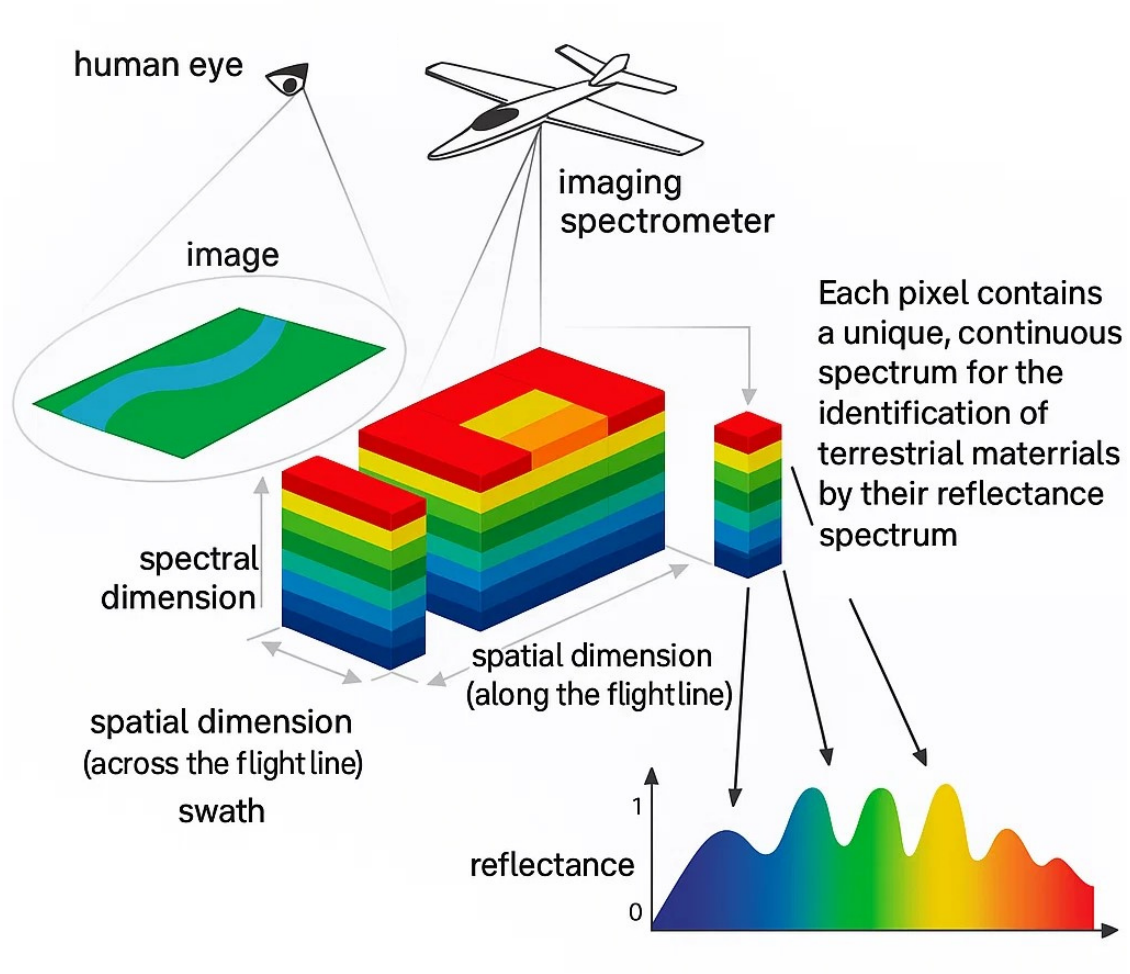

Figure 1.7 – Principle of hyperspectral airborne remote sensing [19].

contiguous wavelength bands. This rich spectral detail enables advanced image analysis capabilities that extend far beyond those offered by conventional imaging systems. The principal analytical tasks in HSI can generally be grouped into three core categories :

1. **Detection :** Identifying the presence of known or unknown materials or objects within a scene, typically by leveraging their distinctive spectral signatures.
2. **Classification and Segmentation :** Assigning each pixel to a specific class or region based on its dominant spectral characteristics, enabling automated mapping of land cover, materials, or surface features.
3. **Spectral Unmixing :** Decomposing each pixel's spectrum into a mixture of constituent materials (known as endmembers) and estimating their respective fractional abundances, which is essential in scenarios where spatial resolution does not allow for pure pixels.

Despite its advantages, HSI also presents challenges due to its high-dimensional data structure and pixel-based format. Each pixel corresponds to a localized spectral measurement of a small spatial area in the observed scene. While this pixel-level granularity provides detailed spectral information, it often lacks contextual or structural information, which can hinder performance in tasks that require higher-level semantic understanding. Furthermore, the fine spatial resolution may not always align with the broader scales necessary for practical decision-making in real-world applications.

As previously discussed, an HSI dataset is typically structured as a three-dimensional data cube, comprising two spatial dimensions (x, y) and one spectral dimension (λ). Each layer (or band) of the cube corresponds to the scene observed at a specific wavelength, while each pixel vector within the cube contains a reflectance spectrum that captures the light absorption and scattering characteristics of the material at that location. These spectral profiles form the foundation for distinguishing between materials based on their unique optical properties. Figure 1.8 illustrates examples of hyperspectral image cubes and highlights how the spectral information is structured across bands.

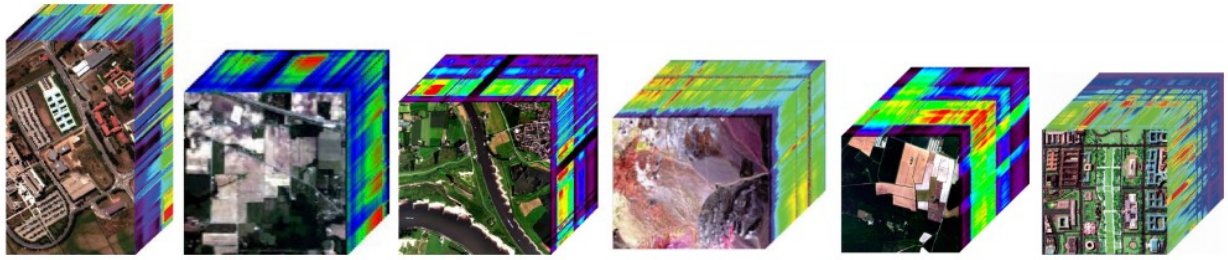

Figure 1.8 – Hyperspectral Image Cubes [2].

5 What is Hyperspectral Imagery ?

Hyperspectral imagery (HSI) provides a rich spectral dataset that enables the identification and differentiation of materials with similar spectral characteristics. Unlike traditional multispectral imaging, which captures reflectance data across a limited number of broad and widely spaced spectral bands, HSI records a continuous and finely resolved spectrum for each pixel. This spectral continuity facilitates more precise and detailed information extraction, significantly enhancing the analytical capabilities of remote sensing systems.

Most multispectral sensors, such as Landsat, SPOT, and AVHRR, collect reflectance data from the Earth's surface in a small number of spectral bands, typically ranging from 3 to 10, with gaps between the measured wavelengths. In contrast, hyperspectral sensors measure reflected radiation across hundreds of narrow and contiguous spectral bands. The resulting spectral profile of each pixel in an HSI image closely resembles a high-resolution laboratory spectrum obtained through traditional spectroscopy. This level of spectral granularity enables the detection of subtle differences in surface composition that are often undetectable with conventional multispectral imaging systems [31].

It is important to note that the distinguishing feature of hyperspectral sensors is not merely the number of spectral bands but the narrowness and contiguity of those bands. Hyperspectral sensors effectively oversample the spectral properties of a scene, thereby improving the discrimination of materials with similar reflectance properties. While the exact number and spacing of spectral bands depend on the specific application and target materials, hyperspectral systems commonly operate at wavelength intervals of 10 to 20 nanometers [31].

To better understand the structural composition of hyperspectral data, it is useful to conceptualize it as a three-dimensional image cube composed of two spatial dimensions (x, y) and one spectral dimension (λ). In this representation, each spatial pixel is associated with a complete spectrum across the measured wavelengths, enabling detailed analysis of spectral signatures for material identification, classification, and unmixing.

Figure 1.9 illustrates this concept using an AVIRIS image of the Leadville mining district in Colorado. The front face of the cube displays a true-color composite, while the top and side faces reveal the spectral dimension. Areas impacted by acid mine drainage, characterized by distinct mineral compositions, appear in red, orange, and yellow hues due to their unique spectral responses.

Figure 1.9 The cube shows an AVIRIS hyperspectral image of the Leadville mining district in Colorado, with the spectral dimension shown as the top and right faces of the cube. The front of the cube is a true color

composite, with areas containing secondary minerals from acid mine drainage highlighted in red, orange and yellow. This cube was processed using ENVI.

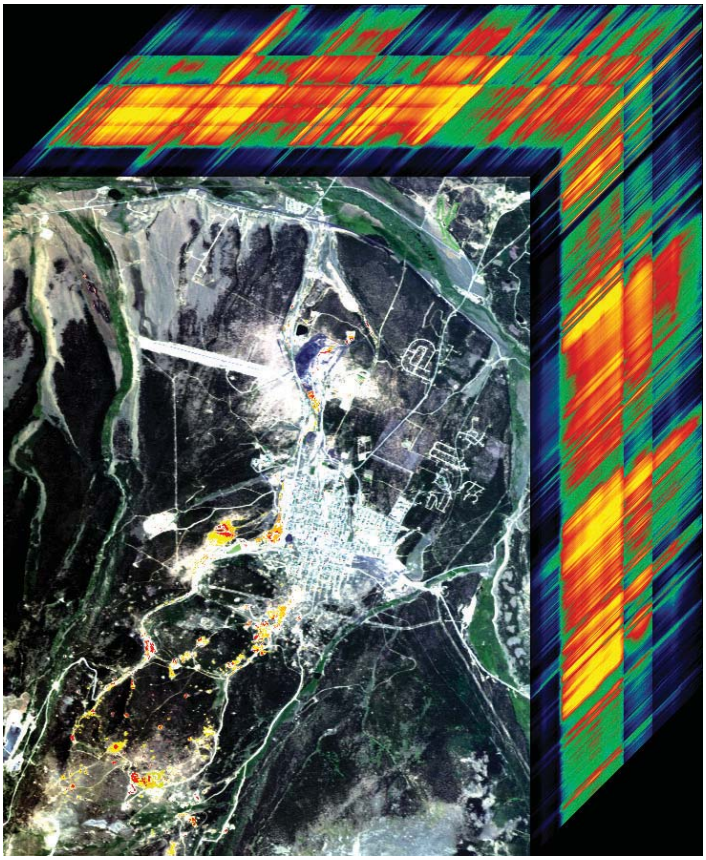

Figure 1.9 – Hyperspectral images are sometimes referred to as "image cubes" because they have a large spectral dimension as well as the two spatial dimensions. [31]

6 Spectral Analysis of Hyperspectral Imagery

Hyperspectral imaging (HSI) provides a highly detailed spectral dataset that enables precise identification and discrimination of materials with similar reflectance properties. In this study, we perform spectral analysis on data extracted from the hyperspectral dataset `cubo_cut.mat`, which was acquired using a drone-mounted sensor operating in the short-wave infrared (SWIR) region. The dataset comprises an image of dimensions 13081×441 pixels, where each pixel is represented by a spectral vector consisting of 60 bands, spanning wavelengths from 980 nm to 1680 nm. A subset of spectral bands was removed due to high noise levels, particularly at wavelengths affected by weak sunlight reflectance.

6.1 Extracting Spectral Signatures

To investigate plastic classification using HSI, five distinct regions of interest (ROIs) were manually selected, corresponding to different plastic types : **Polyvinyl Chloride (PVC)**, **Polypropylene (PP)**, **Polyethylene Terephthalate (PET)**, **Polystyrene (PS)**, as well as a representative sample of a non-plastic background. For each ROI, the spectral reflectance values were averaged across all enclosed pixels to derive characteristic spectral signatures. These signatures serve as reference spectra for subsequent material classification and discrimination tasks.

6.2 Results and Visualization

Figure 1.10 presents the spectral reflectance profiles extracted from the selected ROIs. Each curve illustrates the variation in reflectance across the 60 spectral bands for a given material. The differences in spectral shape and intensity among the plastic types highlight the potential of hyperspectral data for material-specific classification, even when materials exhibit similar visual appearances.

The curves represent averaged pixel spectra for PVC, PP, PET, PS, and a non-plastic background across the SWIR range (980–1680 nm).

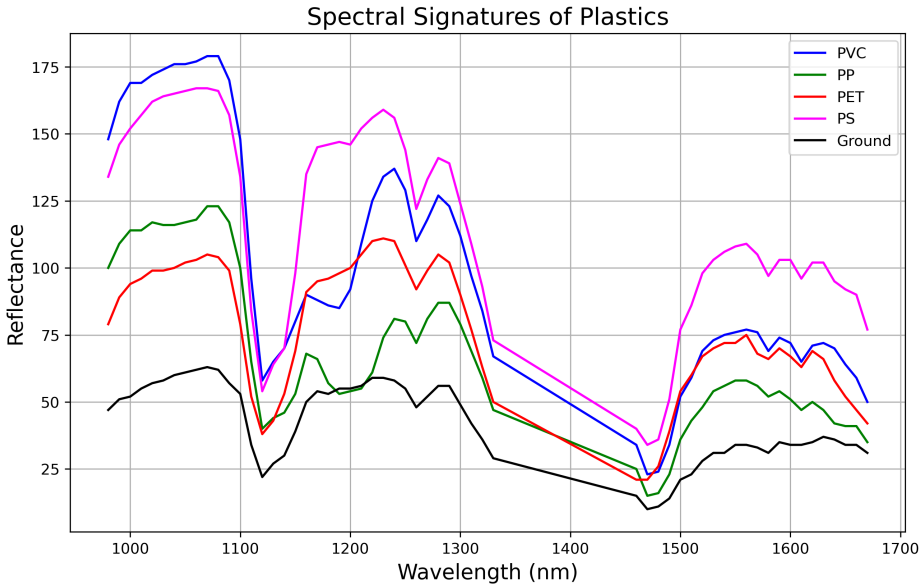

Figure 1.10 – Spectral reflectance profiles of selected plastic materials extracted from the `cubo_cut.mat` hyperspectral dataset.

Note : The title and analysis presented in this section are original contributions developed by the author. They are based on the `cubo_cut.mat` dataset, provided through a collaborative research initiative between the University of 08 May 1945 Guelma and Sapienza University of Rome, with support from Professor Marco Balsi.

7 What Information Does Hyperspectral Imagery Provide ?

Hyperspectral remote sensing offers a substantial advantage in distinguishing materials with similar spectral characteristics, providing a level of spectral detail far beyond that of conventional multispectral imaging systems. A notable example is the discrimination between the minerals kaolinite and alunite, as observed using NASA's Airborne Visible/Infrared Imaging Spectrometer (AVIRIS). Although both minerals exhibit absorption features near 2.2 µm, kaolinite is characterized by a distinctive doublet absorption feature, whereas alunite presents only a single absorption dip in the same region.

Traditional multispectral sensors, such as Landsat ETM+, capture this entire portion of the electromagnetic spectrum within a single, broad spectral band, rendering them unable to resolve such fine spectral differences. In contrast, hyperspectral sensors acquire data over numerous narrow and contiguous bands, enabling detailed analysis of the 2.2 µm region and facilitating the detection of kaolinite's unique double-dip feature. This enhanced spectral resolution is critical for accurately identifying and differentiating materials with overlapping or adjacent spectral features.

Beyond mineralogy, hyperspectral imaging has demonstrated utility across diverse application domains, including vegetation classification, plant health assessment, analysis of building and construction mate-

rials, camouflage detection, and the identification of spectrally similar surface features. These applications underscore the versatility and effectiveness of hyperspectral imagery in delivering high-precision material characterization for both scientific and operational purposes [31].

8 Applications of Hyperspectral Remote Sensing

Hyperspectral imaging (HSI) captures an extensive range of spectral information, enabling the detailed characterization and discrimination of materials based on their unique spectral properties. However, the sheer volume and complexity of hyperspectral data pose significant challenges for conventional image processing techniques. Consequently, the field has evolved to include a suite of specialized algorithms designed to process and analyze this high-dimensional data for a wide variety of real-world applications.

HSI has demonstrated utility across numerous civilian and military domains, including agricultural monitoring, land use and land cover mapping, environmental and ecological assessment, mineral exploration, detection and identification of anthropogenic materials, change detection, target recognition, surveillance, and natural mineral characterization [14, 15]. The primary advantage of HSI in these applications lies in its ability to detect subtle spectral variations associated with the molecular composition of materials, thereby enabling highly precise material identification.

To extract actionable insights from hyperspectral data, a range of preprocessing and postprocessing techniques has been developed. These analytical methods are generally categorized into five principal tasks :

— **Classification :** Assigning each pixel in the hyperspectral cube a discrete label corresponding to a specific material or land cover class [21].
— **Dimensionality Reduction :** Reducing the number of spectral bands while retaining essential information, thereby improving computational efficiency and minimizing the effects of noise and redundancy [1].
— **Spectral Unmixing :** Decomposing mixed pixels into constituent materials (endmembers) and estimating their fractional abundances, particularly important when spatial resolution leads to mixed spectral responses [7].
— **Target and Anomaly Detection :** Identifying rare, distinctive, or unexpected spectral signatures that may correspond to known targets or anomalies in the observed scene [14].
— **Change Detection :** Comparing hyperspectral images of the same area acquired at different times to identify significant spectral or structural changes, frequently used in environmental monitoring, agriculture, and defense applications [13].

This thesis focuses in particular on **Hyperspectral Image Classification (HSIC)**, a foundational task that facilitates the semantic interpretation of HSI data. Emphasis is placed on the integration of supervised and semi-supervised learning techniques to improve classification accuracy and generalization. These approaches leverage a diverse set of classifiers—ranging from generative and discriminative models to ensemble-based and parametric techniques—aimed at efficiently labeling each pixel, including those representing mixed or ambiguous spectral responses, with high precision [34].

9 Special Issues When Working with Hyperspectral Imagery

Despite the remarkable potential of hyperspectral remote sensing, several challenges hinder its widespread adoption and practical implementation. These issues primarily relate to atmospheric correction requirements, the reliance on spectral libraries, and the limited availability of high-quality hyperspectral datasets.

9.1 Atmospheric Distortions and Correction Methods

Many hyperspectral analysis algorithms depend on accurately corrected reflectance data to ensure reliable interpretation of spectral features. Atmospheric effects, including absorption and scattering by gases and aerosols, can distort the recorded spectral signatures. To mitigate these effects, advanced correction algorithms

have been developed to estimate atmospheric parameters directly from the hyperspectral data itself, often removing the need for external ancillary information. These corrections are typically applied on a pixel-wise basis by leveraging the rich spectral content available at each spatial location.

Several commercial remote sensing software packages, such as ENVI and ERDAS IMAGINE, offer robust atmospheric correction tools. However, certain algorithms—such as the *BandMax* tool developed by the Galileo Group—can be effectively applied to radiance imagery without requiring atmospheric corrections, making them useful for rapid or preliminary analyses.

9.2 Reference Spectra and Library Limitations

A critical component of hyperspectral image analysis is the use of spectral libraries—databases that store reference spectra of known materials. These libraries serve as ground truth for classification, target detection, and material identification tasks. Some studies require collecting material spectra in the field using portable spectroradiometers, yielding high-fidelity and context-specific reference data. Alternatively, publicly available spectral libraries (e.g., the USGS Spectral Library) provide a comprehensive set of reference spectra suitable for many applications.

In cases where pre-existing spectral libraries are unavailable or insufficient, users may derive spectral signatures directly from the hyperspectral image under analysis. This is commonly achieved using endmember extraction algorithms—often integrated into commercial software—that identify representative spectra within the image. This approach ensures that the reference spectra are closely aligned with the actual scene content and sensor characteristics.

9.3 Constraints in Hyperspectral Sensor Deployment

One of the major constraints in hyperspectral remote sensing is the limited availability of operational spaceborne sensors. Unlike traditional multispectral platforms (e.g., Landsat or Sentinel), the number of satellites equipped with hyperspectral instruments remains relatively small. However, several noteworthy missions have contributed valuable datasets, including :

— **Hyperion** on NASA's Earth Observing-1 (EO-1) satellite
— **CHRIS** on the European Space Agency's PROBA platform
— **FTHSI** on the U.S. Air Force Research Lab's MightySat II

Moreover, the Earth Resources Observation and Science (EROS) Center offers Hyperion imagery at relatively low cost, improving accessibility for research and educational purposes. In addition to satellite data, airborne platforms—such as NASA's AVIRIS (Airborne Visible/Infrared Imaging Spectrometer)—provide high-quality hyperspectral datasets with flexible acquisition capabilities. A more exhaustive inventory of available hyperspectral sensors is discussed in [31].

9.4 Dimensionality and the Curse of High Feature Spaces

Hyperspectral imaging (HSI) fundamentally operates in a high-dimensional feature space, where each pixel is characterized by a spectral signature spanning hundreds of bands. This results in a multidimensional vector called a feature vector, and all possible such vectors span what is known as the feature space. Visualizing this space is feasible in two or three dimensions using scatter plots, but the real complexity emerges in dozens or even hundreds of dimensions.

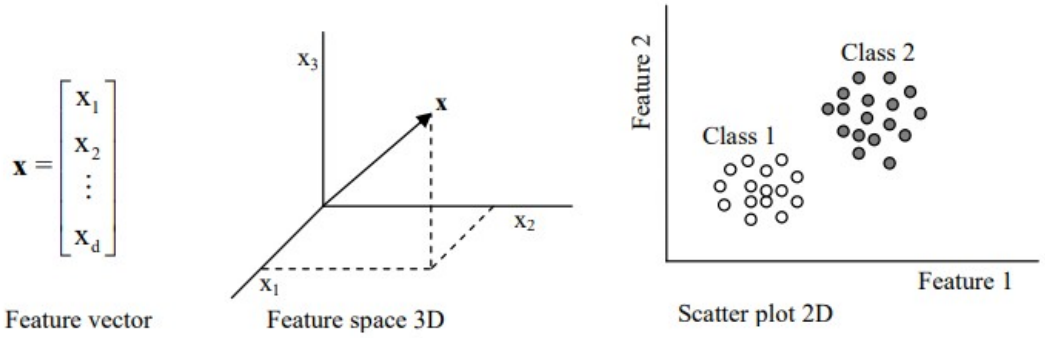

Figure 1.11 – Relationship between feature vector and feature space [4]

Figure 1.11 illustrates a simplified view of the feature space in 2D, where two features (e.g., spectral intensities at two wavelengths) define the axes. Each point represents an object in our case, potentially a pixel belonging to a type of plastic. As we scale this to hundreds of dimensions, visual interpretation becomes impossible, and classification challenges arise due to sparsity.

Let us consider each hyperspectral image cube as a Euclidean space, where the number of spectral bands defines the dimensions. Pixels (representing materials such as various types of plastic debris) are plotted as individual points in this high-dimensional space. In supervised classification, these points are assigned class labels such as PET, PP, or PVC—based on training samples collected via spectroradiometric fieldwork or reference spectral libraries.

However, a well-known limitation emerges when dealing with high-dimensional data : the Hughes phenomenon, also known as the curse of dimensionality. As the number of spectral bands (dimensions) increases, so does the potential for more discriminative power. Yet paradoxically, it also raises the number of statistical parameters required to model the data, and with limited training samples, classifier accuracy can degrade. This leads to data sparsity, where the distance between meaningful data points grows, reducing the statistical significance of any inference drawn from small neighborhoods.

To illustrate this with a more domain-specific analogy : imagine attempting to classify different types of plastic waste (e.g., PET, PE, PVC) floating in a river. Suppose you increase your sensor resolution so much that you are now analyzing hundreds of spectral bands. Initially, this helps distinguish the plastics better—maybe PET has a sharp absorption at 1700nm, while PE peaks near 1200nm. But beyond a certain point, the added detail becomes redundant or noisy, especially when training data is scarce. It's like adding too many dimensions to a plot : your "plastic signatures" spread out thinly and become harder to cluster meaningfully.

Mathematically, if we inscribe a hypersphere (representing real data density) within a hypercube (representing the full dimensional space), the volume occupied by the hypersphere rapidly shrinks relative to the cube as dimensions increase. That is, most of the "volume" of the data lies in the corners, away from where meaningful clustering happens. This sparsity increases classifier uncertainty and affects performance, particularly for parametric models like Mahalanobis distance classifiers.

In contrast, non-parametric methods like Support Vector Machines (SVMs) have shown resilience against the curse of dimensionality in HSI, as demonstrated in the work of Alonso et al. (2011) [4]. Their results on synthetic and real datasets confirmed that SVM maintains classification robustness even with minimal training samples and high dimensional input.

Hence, careful feature selection, dimensionality reduction, or deep models with embedded regularization become essential for plastic detection pipelines, especially when deploying high-resolution hyperspectral systems in real-world applications.

10 Conclusion

The escalating issue of plastic pollution in terrestrial and aquatic environments necessitates scalable and intelligent monitoring systems that surpass the limitations of traditional manual approaches. Hyperspectral imaging (HSI) emerges as a promising technology in this context, offering detailed spectral information at the pixel level that enables precise material identification. However, transforming this rich spectral data into actionable insights presents numerous technical challenges.

Plastic materials often exhibit subtle spectral differences—not only when contrasted with natural backgrounds such as vegetation, soil, or foamy water—but also among polymer types themselves. For example, polymers like polyethylene (PE) and polypropylene (PP) can produce nearly indistinguishable reflectance patterns, particularly under varying lighting conditions or after weathering. These spectral ambiguities are further compounded by environmental noise, illumination variability, and background clutter, which collectively hinder accurate classification.

In addition to material discrimination challenges, the high dimensionality of hyperspectral data introduces significant computational and analytical complexity. The presence of hundreds of contiguous spectral bands leads to redundancy and increased processing demands. Moreover, the scarcity of large, well-annotated, and balanced hyperspectral datasets exacerbates the difficulty of training robust, generalizable models. Developing approaches that are both computationally efficient and resilient to noise and variability is therefore critical.

Addressing these intersecting challenges requires advanced algorithms that effectively combine spectral sensitivity with spatial contextualization, and that are capable of learning discriminative features under constraints such as limited supervision and class imbalance. The next chapter explores how recent literature has approached these issues, laying the groundwork for a comparative analysis of state-of-the-art methods and the development of a hybrid detection strategy tailored to the specific demands of hyperspectral plastic waste detection.

CHAPTER 02

Literature Review and Background

1 Introduction

This chapter presents a comprehensive review of the state of the art in plastic detection using hyperspectral imaging (HSI), focusing on both laboratory and real-world environments. The discussion spans across classical machine learning and recent deep learning approaches, highlighting the methods, materials, and datasets employed in the literature.

Specifically, we examine prior work related to hyperspectral imaging for material identification, with an emphasis on plastic polymers. We explore the spectral characteristics of commonly encountered plastic types and analyze the techniques proposed to address key challenges, including spectral similarity, environmental variability, sensor noise, and limited labeled data. Our review encompasses a broad spectrum of contributions, from physics-based spectral analysis to data-driven classification frameworks.

In doing so, we critically assess both the successes and shortcomings of previous efforts, drawing attention to algorithmic strategies that have proven effective and common pitfalls that hinder generalization and robustness. By synthesizing these insights, we outline methodological best practices and design considerations that inform the development of our proposed system for plastic detection in hyperspectral imagery.

This analysis establishes a foundation for the next stages of this research, where we build on proven techniques while addressing the limitations revealed in earlier studies.

2 Overview of Relevant Literature

2.1 Various Environments (Review)

1. **Deep Learning Approaches**

— **Real-Time Object Detection Based on UAV Remote Sensing : A Systematic Literature Review** (3 October 2023)[12].

This review surveys deep learning frameworks—particularly YOLO variants, SSD, and Faster R-CNN—for real-time object detection in UAV imagery across diverse environments, utilizing RGB, multispectral, hyperspectral, and thermal-infrared sensors on GPU-accelerated edge platforms such as NVIDIA Jetson Nano, TX2, and Xavier NX. The authors screened 222 records from Web of Science and 255 from Scopus, evaluating model optimization strategies including lightweight architecture design, parameter pruning, quantization, and image downsampling. Reported performance metrics span from 1 to over 70 FPS, including precision, recall, F-score, mAP, latency, and onboard energy consumption. However, hyperspectral imaging is only superficially addressed, and no methods are specifically tailored to the spectral characteristics of plastic waste. Our proposed lightweight U-Net model, which incorporates spatial–spectral feature learning, offers a promising solution for pixel-wise plastic detection in hyperspectral UAV data under real-time or near-real-time constraints.

2.2 Laboratory Environment

1. **Machine Learning Approaches**

— **A New Remote Hyperspectral Imaging System Embedded on an Unmanned Aquatic Drone for the Detection and Identification of Floating Plastic Litter Using Machine Learning** (8 July 2023)[3].

This study integrates a remote hyperspectral imaging system (900–1700 nm, NIR–SWIR) onboard an unmanned aquatic drone (ROV-ULCO platform) under laboratory-controlled conditions. The dataset includes reflectance spectra for nine plastic types (HDPE, LDPE, PET, PP, PVC, PS, ABS, PUR, POM) and various non-plastic materials (e.g., wood, paper, rubber, vegetation). Machine learning classifiers—K-Nearest Neighbors, Support Vector Machines, and Artificial

Neural Networks—were trained on these spectral signatures, achieving overall classification accuracy close to 90%. However, the controlled setting and reliance on handcrafted spectral features may limit generalization to complex, real-world aquatic environments. Our proposed lightweight U-Net with spatial–spectral feature learning seeks to overcome these limitations, enabling robust real-time plastic detection in diverse conditions.

— **Systematic Reduction of Hyperspectral Images for High-Throughput Plastic Characterization** (7 December 2023)[18].

This work employs hyperspectral imaging (900–1700 nm) on a laptop (Intel(R) Core(TM) i7-10850H @ 2.70 GHz) to analyze four plastic polymers—PP, PE, PS, and PET. Techniques such as Nonnegative Matrix Factorization (NMF), Essential Spectral Pixels (ESPs), and Essential Spatial Variables (ESVs) were used to reduce the original dataset (144,000 $\times$ 173) to smaller representations (48 $\times$ 12 and 149 $\times$ 45), reducing processing time from approximately 220 seconds to 0.003 seconds. However, no accuracy or performance metrics were reported.

2. **Deep Learning Approaches**

— **Municipal Solid Waste Segregation with CNN** (2019)[32].

Using a balanced dataset of 9,200 images covering 20 item-level waste classes (e.g., foam containers, snack foils, batteries) and four waste-type classes (general, compostable, recyclable, hazardous), this study implemented CNN architectures—VGG-16, ResNet-50, MobileNet V2, and DenseNet-121—pre-trained on ImageNet. These were fine-tuned via transfer learning for classification tasks. ResNet-50 achieved 94.86% accuracy for waste-type classification, with compostable waste classification peaking at 98.72%.

— **Plastic Classification with NIR Hyperspectral Images and Deep Learning** (5 January 2023)[30].

This study employs a Specim FX17 hyperspectral camera, illuminated with four 150 W lamps over a conveyor belt, in a recycling-facility testbed. It captures 1,491 images (4,500 plastic pieces ; ABS, acrylic, PC, PET, PP, PS, POM, and non-plastics), totaling 161 billion sample pixels. A GPU-powered workstation processes data using Approximate Nearest Neighbors, FAISS Mask, LightGBM, and N-BEATS. The best-performing model achieved an F1-score of 0.79 and overall accuracy of 0.79—reaching 0.90 for colored plastics and 0.67 for black plastics—at a processing rate of two images per minute.

— **Low-Cost Recognition of Plastic Waste Using Deep Learning and a Multi-Spectral Near-Infrared Sensor** (28 April 2024)[24].

Utilizing the SparkFun AS7265x spectroscopy sensor (18 bands ; 410–940 nm ; GBP 56) in a light-sealed, matt-black enclosure, the study collected 1,200 spectra from 423 household plastic samples (PET, HDPE, PVC, LDPE, PP, PS). Ten ML/DL models (CNN, MLP, SVM, kNN, RF) were tested using PCA/LDA for dimensionality reduction and 10-fold cross-validation. The PCA-CNN achieved the best performance with 72.5% mean accuracy and 72.38% F1-score. Highest recognition accuracy was seen for PS (83.5%), with PET/PVC confusion at 66%, and HDPE/PP misclassifications.

2.3 Real-World Environment

1. **Machine Learning Approaches**

— **Detecting the Great Pacific Garbage Patch Floating Plastic Litter Using Satellite Imagery** (October 2021)[28].

Using WorldView-3 multispectral imagery (eight Visible-NIR bands), this study analyzed two images from the Great Pacific Garbage Patch (17 September 2018 and 19 August 2019), with The Ocean Cleanup System 001 Wilson as a reference target. A spectral anomaly-based detection approach was implemented using Python, ENVI, and MATLAB, with the Reed-Xiaoli Detector (RXD) algorithm. Plastic debris was successfully detected by identifying NIR spectral anomalies, despite challenges from cloud shadows, ocean glint, and spectral overlaps with whitecaps.

— **A Preliminary Study on the Utilization of Hyperspectral Imaging for the On-Soil Recognition of Plastic Waste Resulting from Agricultural Activities** (18 October 2023)[9].

This work employs an ImSpector™ N17E NIR-HSI device (1000–1700 nm, 161 bands) to collect spectra from plastic fragments (PE, EVA, PP, biodegradable mulch) and agricultural soils at six locations in Sabaudia, Italy. Preprocessing involved SNV, mean-centering, and PCA, followed by PLS-DA classification on 241,538 spectra (98,863 plastic vs. 142,675 soil). Processing in MATLAB (PLS_Toolbox v9.0 and MIA_Toolbox v3.1) yielded classification sensitivity and specificity of 92% and overall accuracy between 95% and 98%.

— **Hyperspectral Imaging for Detecting Plastic Debris on Shoreline Sands to Support Recycling** (9 December 2024)[27].

This study uses near-infrared hyperspectral imaging (900–1700 nm) along with micro X-ray fluorescence and FT-IR analysis to identify plastic debris on three Italian beaches. Preprocessing included SNV, Savitzky–Golay smoothing, derivatives, and mean-centering. PCA, cascade detection, and PLS-DA were used for classification. Plastic classes (PP, PE, PS, PVC, PET) were distinguished from pollutants like glass, paper, cigarette butts, and wood, achieving sensitivity and specificity up to 1.00 and sand-class sensitivity between 0.802 and 0.885.

— **Plastic Litter Detection in the Environment Using Hyperspectral Aerial Remote Sensing and Machine Learning** (6 March 2025)[5].

This research integrates a hyperspectral push-broom system (Xenics Bobcat 320 SWIR + ImSpector NIR17 + RGB camera + VN-200 INS) on a DJI Matrice 600 UAV. Spectral cubes (81 bands, 980–1340 nm and 1460–1680 nm) were labeled into nine classes (six plastic types, white reference, non-plastic, background), totaling 6,350 balanced spectra. Feature selection used continuum removal and mRMR ; classification was done using LDA and RBF-SVM with leave-one-out cross-validation and Cohen's Kappa evaluation. Seven temporal acquisitions (Jan 2020–Feb 2024) showed SVM Kappa scores between 0.61 and 0.92, indicating robust performance.

2. **Deep Learning Approaches**

— **Detection of River Plastic Using UAV Sensor Data and Deep Learning** (25 June 2022)[23].

This study deployed a DJI Phantom 4 UAV (0.82 cm GSD) over two riverine locations in Laos and Thailand, acquiring 500 ortho-image tiles (256256 px) per site. Plastic debris was manually annotated using bounding boxes. YOLO models (v2–v5) were trained from scratch and via transfer learning (feature-freezing and fine-tuning) using Anaconda (Intel i7-10750H, NVIDIA RTX 2060) and Google Colab Pro. YOLOv5s yielded the best performance (mAP 0.81, F1 0.78, inference time 0.019 s/img, 13.6 MB model). Transfer learning improved cross-site performance by 2.7–22%. Detection challenges included bright, transparent, and soil-covered plastics.

2.4 Comparative Articles Table

Article (Reference)	Strengths	Weaknesses
Real-Time Object Detection Based on UAV Remote Sensing : A Systematic Literature Review [12]	Increases detection efficiency ; optimized for on-board edge computing.	**Limited onboard resources** ; difficult real-time processing for large multi-/hyperspectral datasets.
A New Remote Hyperspectral Imaging System Embedded on an Unmanned Aquatic Drone for the Detection and Identification of Floating Plastic Litter Using Machine Learning [3]	90% detection accuracy ; demonstrates potential for automated litter identification.	**Limited spectral range** restricts material discrimination ; environmental variability and data redundancy slow processing.
Systematic Reduction of Hyperspectral Images for High-Throughput Plastic Characterization [18]	Strong data dimensionality reduction ; enhances HSI analysis efficiency.	**Difficulty with multilayer/-complex plastics** ; challenges in characterizing heterogeneous materials.
Municipal Solid Waste Segregation with CNN [32]	High classification accuracy (up to 94.9%) ; transfer learning improves generalization.	**Limited dataset size** ; RGB input insufficient for fine plastic-type discrimination.
Plastic Classification with NIR Hyperspectral Images and Deep Learning [30]	High overall F1-score (0.79) ; demonstrates feasibility for NIR HSI classification.	**Reduced performance for difficult classes** (e.g., black plastics) ; slow real-time inference speed.
Low-Cost Recognition of Plastic Waste Using Deep Learning and a Multi-Spectral Near-Infrared Sensor [24]	Ultra-low-cost sensor ; portable design for field applications ; generalizable pipeline.	**Moderate accuracy (72.5%)** ; significant confusion between PET/PVC and HDPE/PP ; limited PVC/PS samples.
Detecting the Great Pacific Garbage Patch Floating Plastic Litter Using Satellite Imagery [28]	High spatial resolution detection of small debris ; enables large-scale monitoring ; proxy method reduces need for full atmospheric correction.	**Cloud shadows/glint interference** degrades detection ; requires field validation ; high cost of high-resolution imagery.
A Preliminary Study on the Utilization of Hyperspectral Imaging for the On-Soil Recognition of Plastic Waste Resulting from Agricultural Activities [9]	Rapid, non-invasive, reliable "in situ" or lab-scale application ; provides plastic aging information.	**Limited by debris size and detection depth** ; field acquisition conditions and soil chemistry interfere ; poor detection of carbon-black polymers.
Hyperspectral Imaging for Detecting Plastic Debris on Shoreline Sands to Support Recycling [27]	Enables polymer-type identification ; robust across varying sandy backgrounds.	**Complex environmental backgrounds** degrade accuracy ; struggles with weathered microplastics.
Plastic Litter Detection in the Environment Using Hyperspectral Aerial Remote Sensing and Machine Learning [5]	Adapts to changing environmental conditions without per-mission calibration ; real-time onboard detection feasible.	**Degrades on poorly structured surfaces (especially water)** ; lacks georeferencing ; manual labeling needed.
Detection of River Plastic Using UAV Sensor Data and Deep Learning [23]	Efficient YOLOv5s model (mAP 0.81, fast inference, small model size) ; transfer learning improves cross-site adaptation.	**Struggles with soil-covered, bright, or transparent plastics** ; requires manual labeling and significant training resources.

Table 2.1 – Summary of reviewed articles' strengths and weaknesses.

Our work intends to bridge these gaps by addressing the main recurring limitation across the literature : the inability to jointly exploit high-resolution spatial features and full-range spectral information in variable environmental conditions. Existing methods either rely on limited spectral bands, simplistic models, or lack

robustness in real-world settings. Our approach enhances generalization across plastic types and dynamic environments, overcoming the core limitations identified in previous studies.

3 Conclusion

This chapter has provided a comprehensive review of the current landscape in plastic detection using Hyperspectral Imaging (HSI), spanning various application domains including laboratory experiments, in-field assessments, UAV-based missions, and satellite remote sensing. We critically analyzed a spectrum of methodological approaches, ranging from classical machine learning algorithms to cutting-edge deep learning architectures, and assessed their performance across key evaluation metrics such as detection accuracy, spectral and material discrimination, computational efficiency, and adaptability to challenging environmental conditions.

Our comparative analysis highlights the growing maturity of HSI-based plastic detection systems, while also revealing persistent gaps that hinder operational scalability and generalization. These include limitations in handling complex or multilayered materials, reduced robustness in real-world, unstructured environments, and trade-offs between model accuracy and real-time feasibility. This synthesis sets the stage for identifying opportunities where hybrid models, sensor fusion, and intelligent preprocessing pipelines can significantly improve detection reliability and field applicability.

CHAPTER 03

Methodology and Experimental Framework

(Description of the developed approach and experiments)

1 Introduction

The rapid advancement of hyperspectral imaging (HSI) technology has opened new avenues for the automatic detection, classification, and segmentation of plastic waste in diverse environmental settings. Hyperspectral data, which captures rich spectral information across hundreds of contiguous bands, enables precise material discrimination that surpasses conventional imaging techniques. However, the high dimensionality and volume of hyperspectral data present unique computational and algorithmic challenges that necessitate robust preprocessing, dimensionality reduction, and intelligent modeling strategies.

This chapter presents the comprehensive methodological framework adopted in this study for plastic waste detection using HSI. Our work explores multiple machine learning and deep learning paradigms tailored to the structure of the dataset and the nature of the detection task. To rigorously evaluate different modeling strategies, we utilized a diverse collection of hyperspectral datasets, including both traditional `.mat` (MATLAB matrix) spectral cubes and `.nc4` (NetCDF) formats. These datasets contain annotated scenes of plastic debris under controlled and naturalistic conditions, offering a rich testbed for algorithmic experimentation.

The proposed experimental framework comprises three main methodological directions. First, traditional machine learning approaches were investigated, involving spectral feature extraction followed by classification using Support Vector Machines (SVM) and Random Forest (RF) classifiers. These models were applied to pixel-level spectral signatures derived from the hyperspectral cubes, leveraging classical statistical learning principles to establish performance baselines.

Second, we designed and implemented deep learning architectures based on the U-Net model, which is widely recognized for its efficacy in biomedical and remote sensing segmentation tasks. The U-Net was trained on individual band combinations and spatial patches. This model was applied to the high-resolution hyperspectral cubes to produce fine-grained plastic segmentation masks.

Third, a hybrid two-phase approach was introduced, integrating deep learning-based segmentation with traditional classification. In this architecture, a U-Net model is first used to segment potential plastic regions, followed by a dedicated SVM classifier trained to predict specific plastic types. This pipeline was particularly applied to the NetCDF-based dataset, allowing us to exploit the spatial precision of U-Net alongside the spectral discriminative power of SVM.

In the following sections, we elaborate in detail on the dataset characteristics, preprocessing pipeline, model architectures, training procedures, and evaluation strategies, thus laying the groundwork for reproducible and comparative experimentation across different detection paradigms.

2 Problem Statement

The goal of Chapter 3 is to develop and articulate a coherent methodological framework for detecting and classifying plastic waste using hyperspectral imaging data. Building on the introductory discussion, we identify the central challenges and research questions that guide our choice of data processing steps, model architectures, and evaluation strategies. Specifically, this chapter seeks to address the following interrelated issues :

1. **High Dimensionality of Hyperspectral Data.**
   - — How can we process and organize hyperspectral cubes (with hundreds of contiguous bands) to enable efficient model training without losing critical spectral information ?
   - — What preprocessing steps (e.g., spectral standardization, patch extraction, optional dimensionality reduction) are necessary to balance computational feasibility and classification accuracy ?
2. **Heterogeneous Data Formats and Acquisition Conditions.**
   - — In what ways do controlled-environment scenes (laboratory-acquired cubes) differ from naturalistic scenes (field-acquired cubes) in terms of label quality, background variability, and illumination effects ? How do these differences impact model selection and training protocols ?
3. **Selection of Learning Paradigms.**
   - — What are the relative strengths and limitations of traditional machine learning classifiers (SVM, Random Forest) versus deep learning-based segmentation (U-Net) when applied to pixel-level spectral signatures ?
4. **Integration of Spatial and Spectral Information.**
   - — How do we determine the optimal patch size, stride, and overlap to ensure that each model captures sufficient spatial context without incurring prohibitive GPU memory usage ?

5. **Model Generalization.**
   — Which performance metrics (accuracy, per-class F1-score, IoU) best capture the trade-offs between overall detection rates and the ability to distinguish among different plastic types ?
6. **Computational Resources and Practical Constraints.**
   — Given the high dimensionality of hyperspectral input and the depth of U-Net architectures, what memory–computation trade-offs (e.g., batch size, number of filters, use of PCA) must we consider to enable training on a single GPU ?
   — How do we structure the code to facilitate reproducibility, parameter tuning, checkpointing, and logging of intermediate results ?

By systematically addressing these questions, Chapter 3 aims to (1) justify the choice of algorithms and hyperparameters ; (2) describe implementation details in a reproducible manner ; and (3) lay out the evaluation protocols that will be used in Chapter 4 to compare performance across different approaches. This problem statement thus provides the scaffold for the detailed methods that follow, ensuring that each design decision is motivated by one or more of the challenges enumerated above.

3 System Overflow

Figure 3.1 provides an overview of the entire methodological pipeline developed in this chapter. The flowchart is organized into four sequential stages :

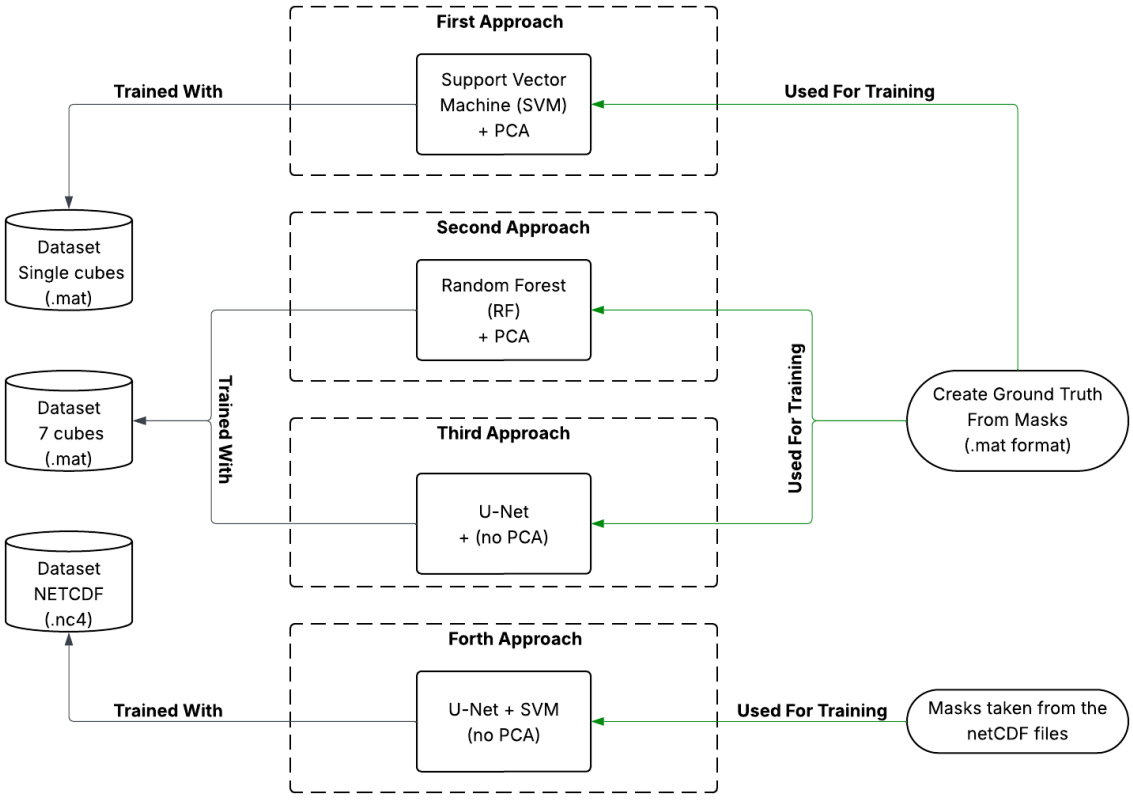

Figure 3.1 – the Overview of the entire methodological pipeline

3.1 Data Description

3.1.1 3.1.1 `cubo_cut` Data Cube

Within the framework of Algerian–Italian cooperation, and as part of the collaboration between our university (08 May 1945 Guelma) and Sapienza University of Rome, in partnership with Professor Marco Balsi, we were provided with a hyperspectral data cube representing various plastic and non-plastic materials for research purposes.

The dataset used in this visualization corresponds to a case study in which samples of different materials were laid on the ground and surveyed by a hyperspectral sensor operating in the SWIR band, mounted on a drone, as shown in Figure 3.2. The file `cubo_cut.mat` contains a hyperspectral cube (variable `cubo`) consisting of an array of 13081 441 pixels, each represented by a spectral vector of 60 bands. These bands are 10 nm wide and are centered at wavelengths ranging from 980–1340 nm and 1460–1680 nm. Some wavelengths were removed due to weak sunlight signals causing high noise levels. The cube corresponds to a central stripe of the surveyed scene, cropped to make the dataset more manageable.

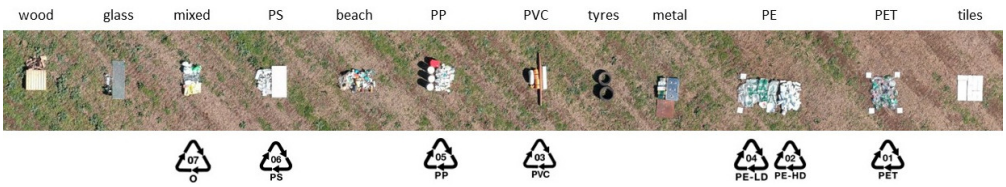

Figure 3.2 – Visible RGB image of the scene during hyperspectral drone acquisition.

Figure 3.2 presents an RGB image of the target scene captured during the drone-based hyperspectral acquisition campaign. While the actual hyperspectral data contains reflectance values across hundreds of spectral bands, this figure illustrates only the visible spectrum for visualization purposes. The RGB snapshot helps contextualize the scene layout, object distribution, and lighting conditions during the data collection process.

Binary masks (available in the `masks_cut` folder) of the same spatial dimensions are provided for class labeling. These masks identify pure polymers (PE, PET, PP, PS, PVC), mixed plastics, and non-plastic materials (background, white, other non-plastics). Two classification tasks are considered : (1) plastics vs. non-plastics, and (2) distinguishing between individual polymer types.

As we have seen earlier in Figure 1.10 in Chapter 1, the mean spectral profiles extracted from this hyperspectral cube allow us to clearly observe the unique reflectance signatures of different plastic materials. Each curve in that figure demonstrates how plastic types reflect light differently across the SWIR bands, which forms the foundation for accurate polymer classification in our proposed framework.

Figure 3.3 shows the acquired masks ; each mask represents a specific material or class extracted from the same hyperspectral cube, highlighting the extensive spatial and spectral information contained within a single acquisition.

Due to the large number of pixels and spectral bands, the hyperspectral cube is substantial in size and provides detailed reflectance characteristics across a wide spectral range. This vast amount of data requires careful preprocessing and management to be effectively utilized in machine learning and deep learning models. The primary challenge is to extract a minimal set of features while maintaining high classification performance and to reduce the number of wavelengths used without compromising predictive reliability.

(a) Mask for all materials

(b) Mask for all plastics

(c) Mask for mixed plastics

(d) Mask for other non-plastics

Figure 3.3 – Visualization of different masks extracted from the hyperspectral cube. Each subfigure shows a specific material or class present in the surveyed scene.

3.1.2 3.1.2 Balsi–Moroni–Bouchelaghem Hyperspectral Plastic Detection Dataset

Data Acquisition Setup

The dataset was acquired by mounting a shortwave-infrared (SWIR) push-broom hyperspectral camera on an unmanned aerial vehicle (UAV). Specifically, the Xenics Bobcat 320 camera (900–1700 nm, 320×256 px), equipped with a Specim ImSpector NIR17 spectrometer, was used [5]. The payload, which included the SWIR sensor and an auxiliary RGB camera, was flown on a DJI Matrice 600 drone at an altitude of 7–15 m, moving at a speed of 0.5–2 m/s with a frame rate of 12.5 Hz, yielding a ground sampling distance of approximately 1–2 cm. During each flight, a synchronized visible camera was used to mosaic overlapping frames, and inertial/GNSS data were logged to aid hyperspectral cube reconstruction. White reference targets (commercial white floor tiles) were placed within the scene for on-the-fly radiometric calibration checks. Figure 3.4 illustrates the aerial setup and the sample layout.

Surveys were conducted in natural outdoor environments over several campaigns (January 2020, March 2020, April 2022, and two days in February 2024) under varied conditions [5]. In most campaigns (except one conducted over greenhouses covered with transparent polypropylene foil), plastic objects were "planted" on grass or bare soil. Plastic litter was sorted by polymer type and secured (e.g., using hidden nets or threads) to minimize movement. In addition, unsorted weathered beach-plastic samples and various non-plastic debris (e.g., wood, metal, glass) were included, along with white tile targets for radiometric calibration. Flights were performed in different seasons and lighting conditions (sunny or cloudy), capturing a wide variety of exposures [5].

Figure 3.4 – Aerial hyperspectral-survey setup showing the SWIR camera mounted on the UAV, white reference tiles, and arranged plastic samples [5].

3.1.3 Dataset Structure and Files

The released dataset comprises **seven hyperspectral data cubes**, corresponding to the flights labeled `Jan20`, `Mar20`, `Apr22`, `7Feb24a`, `7Feb24b`, `21Feb24a`, and `21Feb24b` [5], as shown in the setup images in Figure 3.5. Each cube is a "rectangular" hyperspectral volume (parallelepiped) formed by stitching together 1-D scans along the flight path. In the provided cubes, the along-track dimension was zero-padded at the edges so that each cube appears as a full 2-D array (pixels set to zero outside the scanned swath).

in Figure 3.5, Material labels are self-explanatory. Unknown or unsorted plastic materials are indicated as mixed (this class also includes type-7 plastics according to international codes), unknown, or beach—the latter indicating weathered litter collected from a beach. High- and low-density PE (HDPE and LDPE) have different standard codes but are chemically identical and not distinguishable by reflected spectra ; they were thus grouped into a single PE class. These images were obtained by stitching together a few visible photographs taken during the same or separate flights. They are not necessarily rectangular and are not co-registered with the hyperspectral data, serving only to document the ground situation [5].

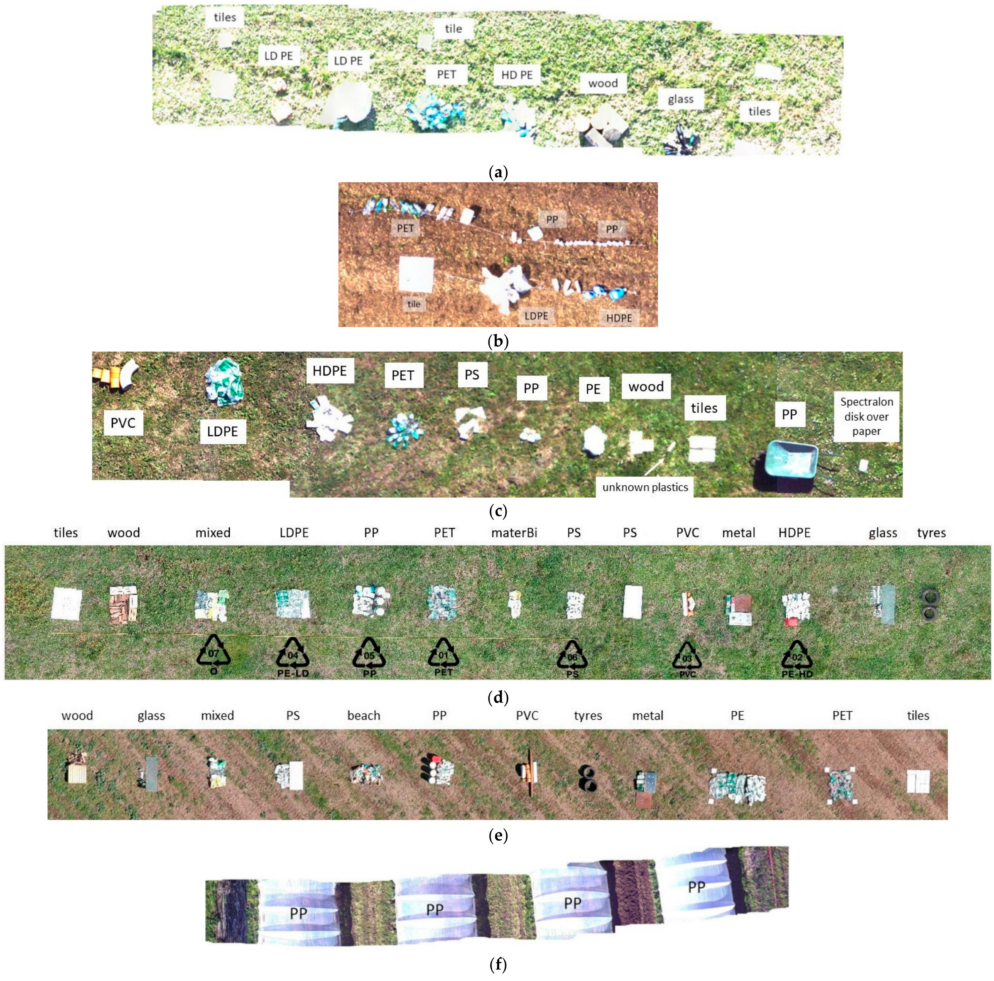

Figure 3.5 – Ground setup for the surveys used in plastic detection machine learning. (a) Jan20, (b) Mar20, (c) Apr22, (d) 7Feb24a and 7Feb24b, (e) 21Feb24a, (f) 21Feb24b. [5].

Each cube contains spectral data ranging from 900 to 1700 nm. The raw camera collected approximately 300 spectral bands per line (with a resolution of ≈2.5 nm), but the data were downsampled and binned to **81 bands** spaced at 10 nm intervals [5]. The first four bands (nominally 900–930 nm) are empty due to the sensor/spectrometer geometry. Additional processing removed noisy regions : bands at the spectral edges and around 1400 nm (due to strong water absorption) were discarded, leaving **60 usable bands** spanning 980–1340 nm and 1460–1680 nm. Thus, the hyperspectral cubes ultimately provide 60 effective SWIR channels of data per image.

The data files (archived on Mendeley Data) include the full hyperspectral cubes and related metadata [w5]. In addition to the raw or ENVI-format cubes, the dataset includes **manually drawn mask files** for each cube, providing per-pixel class labels (plastic polymer or other). Visible-light photographs of the scenes (showing the arrangement of samples and white tiles) are also provided, along with MATLAB-readable `.mat` versions of the cubes. The `.mat` files contain the trimmed spectral range (940–1340 nm and 1460–1680 nm ; 60 bands), consistent with the noise removal procedure. No proprietary formats are required : the core hyperspectral data are stored as numerical arrays, with masks provided as corresponding label arrays.

3.1.4 Plastic Types and Sample Categorization

Nine material categories were defined : **six plastic-polymer classes**, along with *Other Plastics*, *Non-Plastic Debris*, *Background*, and a separate *White Reference* class. The six primary polymer classes are listed in Table 3.1.

In addition, the label "Non-Plastic Debris" includes materials such as wood, metal, and glass. The "Background" class represents natural surfaces such as grass or soil. The "White Reference" class corresponds to commercial white floor tiles used for radiometric calibration. Table 3.1 summarizes the approximate number of labeled samples per cube for each class.

Class	Description/Material	Samples per Cube
PET	Polyethylene terephthalate (e.g., bottles)	100
PE (LDPE + HDPE)	Low-/High-density polyethylene (e.g., bags, containers)	100
PP	Polypropylene (e.g., foam items, caps)	100
PS	Polystyrene (e.g., insulation foam, packaging)	100
PVC	Polyvinyl chloride (e.g., pipes, rigid packaging)	100
Other Plastics	Mixed or unsorted plastic waste, bioplastics	100
Non-Plastic Debris	Wood, metal, glass, etc. (non-polymeric litter)	300
Background	Natural background (grass, soil)	300
White Reference	Commercial white tiles (radiometric reference)	100

TABLE 3.1 – Material classes and approximate sample counts per hyperspectral cube.

in the Table 3.1, Each plastic-polymer class includes roughly 100 labeled pixel samples per cube. Non-plastic and background classes were oversampled to provide diverse negative examples. Across all cubes, approximately 7300 total labeled spectra were initially collected, later reduced to ∼6350 after removing low-intensity pixels (DN < 32).

3.1.5 Labeling and Annotation

All hyperspectral cubes were annotated through manual delineation of polygons around each object or region of interest. For every cube, trained operators created pixel-wise mask images to identify the following classes : PET, PE, PP, PS, PVC, Other Plastics, White Reference, Non-Plastic Debris, and Background [5]. These mask files are provided alongside each corresponding hyperspectral cube.

For example, pixels corresponding to a transparent PET bottle were labeled as "PET," whereas regions of vegetation such as grass were assigned the "Background" label. The "White Reference" class was explicitly included to mark the white calibration tiles used for radiometric standardization.

Each mask is stored as a binary or integer-labeled image, spatially co-registered with its corresponding hyperspectral data cube. These annotations provide the ground truth required for supervised classification and model evaluation.

3.1.6 Preprocessing

The raw hyperspectral data underwent minimal but essential preprocessing :

1. **Low-signal filtering :** Spectra with an average digital number (DN) less than 32 were removed to eliminate saturated, noisy, or excessively dark pixels.
2. **Spectral downsampling and trimming :** The original ∼300 spectral bands (with ∼2.5 nm resolution) were binned to 81 bands with 10 nm spacing. Subsequently, noisy edge bands and those affected by strong

water absorption around 1400 nm were discarded, resulting in 60 usable bands covering 980–1340 nm and 1460–1680 nm.

Each pixel's spectrum is therefore represented by 60 effective shortwave infrared (SWIR) wavelengths, with intensity values in raw DN units (0–255). No atmospheric or radiometric correction to reflectance was applied, as classification tasks can proceed directly from DN values. However, white-tile spectra are available for optional post hoc calibration.

3.1.7 Environmental and Experimental Settings

The seven data cubes encompass a wide range of *real-world survey conditions*. Most acquisitions were conducted over short vegetation or bare soil, with plastic samples carefully arranged on the ground [5]. One unique case (21Feb24b) involved acquisition over greenhouses covered with transparent polypropylene (PP) foil.

Data collection spanned multiple seasons and lighting conditions (sunny and cloudy), introducing natural variations in illumination, shadows, and background brightness. Vegetated scenes typically show high SWIR reflectance, whereas samples placed over darker backgrounds (e.g., bare soil or greenhouse floors) yield different spectral contrast levels. These environmental factors increase the complexity of the detection task, requiring models to focus on intrinsic polymer-specific absorption features while being robust to contextual variations.

3.1.8 Challenges and Unique Characteristics

The dataset is particularly valuable due to its diversity, operational realism, and the challenges it presents for hyperspectral image analysis :

— **Multi-seasonal variability :** Flights were conducted from January 2020 to February 2024, allowing the dataset to capture temporal spectral variations across seasons. Models trained on data from a single flight often perform poorly under new lighting conditions, emphasizing the importance of generalization [5].
— **Background clutter :** Unlike ocean-based plastic detection, terrestrial backgrounds such as grass and soil exhibit strong SWIR signals. This complicates the separation of plastic materials from natural backgrounds, especially under variable lighting and occlusion.
— **Spectral band removal due to water absorption :** Bands near 1400 nm were removed because of atmospheric water vapor absorption, which would otherwise introduce significant noise and missing data points.

The Balsi–Moroni–Bouchelaghem dataset comprises seven SWIR hyperspectral data cubes (900–1700 nm) acquired by UAV-mounted sensors using a push-broom mechanism. Each cube is accompanied by per-pixel ground-truth annotations, covering six primary plastic polymers, other materials, and background. This dataset offers a richly annotated and environmentally varied benchmark for developing robust hyperspectral plastic detection and classification models under realistic conditions.

3.1.9 3.1.3 NetCDF Data

In addition to our own dataset, we incorporated a publicly available hyperspectral reflectance database developed by Olyaei et al. [25], titled *"A Hyperspectral Reflectance Database of Plastic Debris with Different Fractional Abundance in River Systems"*, published in *Optics Express* in November 2024. The dataset was created by the research team at the Saint Anthony Falls Laboratory, University of Minnesota, and integrated directly into our experimental framework. It provides high-resolution hyperspectral reflectance measurements, along with comprehensive metadata, and is specifically designed to support the automated classification of six common plastic polymers in fluvial environments.

The data were acquired through controlled laboratory experiments simulating riverine conditions. Hyperspectral measurements were captured using an Analytical Spectral Devices (ASD) FieldSpec-4 spectroradiometer, which covers the full spectral range of 350–2500 nm, encompassing the ultraviolet (UV), visible near-infrared (VNIR), and shortwave infrared (SWIR) regions. The spectroradiometer was mounted on a mobile carriage positioned 110 cm above the water surface to ensure consistent angular geometry and spatial coverage.

To replicate realistic river conditions, measurements were conducted over a 14.6 m long, 0.9 m wide, and 0.6 m deep hydraulic flume designed to simulate subcritical flow. The flume bed was covered with coarse sand to mimic natural substrates, while black-painted sidewalls minimized stray reflections. Various water surface conditions—clear, turbid, and foamy—were created by adjusting sediment concentrations and using controlled aeration. Additionally, a high-resolution DSLR camera was synchronized with the spectroradiometer to capture co-registered RGB imagery, supporting visual verification and alignment with spectral data.

Figure 3.6 illustrates the experimental setup, showing the optical sensors, lighting, and environmental conditions under which the hyperspectral data were acquired.

Figure 3.6 – (a) The ASD FieldSpec-4 spectroradiometer with the synchronized DSLR camera ; (b) Illumination system and reference panel, showing black-painted sidewalls ; (c) Sediment feeder for turbid water conditions ; (d) Floating plastic debris in a turbid flow, with Acoustic Doppler Velocimeter (ADV) visible. [25]

3.1.10 Plastic Types and Environmental Context

The dataset includes six common polymer types, as well as mixed and weathered plastic specimens, to represent realistic riverine plastic pollution scenarios. In total, 2078 spectral samples were collected. Table 3.2 summarizes the number of samples per polymer type.

Polymer Type	**Number of Samples**
Polyethylene Terephthalate (PET)	192
High-Density Polyethylene (HDPE)	145
Low-Density Polyethylene (LDPE)	152
Polypropylene (PP)	234
Expanded Polystyrene Foam (EPSF)	92
Mixed Plastics	1008
Weathered Plastics	255

Table 3.2 – Sample distribution across polymer types.

Plastic specimens varied in shape, size, thickness, and color, introducing spectral variability that enhances the dataset's robustness. These variations better simulate real-world plastic debris found in aquatic environments. Spectral reflectance measurements were obtained for 21 subgroups—corresponding to 7 specimen types under 3 water background conditions. The conditions of the plastic samples are detailed in Table 3.3.

Environmental realism was further enhanced by recording spectral data under three distinct water background conditions, summarized in Table 3.4. These conditions help quantify the influence of water clarity and surface properties on hyperspectral reflectance and support the development of generalized detection algorithms.

Condition	Description
Virgin Plastics	Pristine polymer samples directly obtained from manufacturers.
Weathered Plastics	Samples exhibiting natural degradation, including oxidation, fragmentation, and discoloration, collected from real aquatic environments.
Mixed Plastics	Combinations of multiple polymer types, reflecting heterogeneous plastic waste typically found in polluted waterways.

TABLE 3.3 – Condition-based categorization of plastic specimens.

Water Background	Description
Clear Water	Low-sediment water collected from the Mississippi River, representing baseline optical conditions.
Turbid Water	Simulated using fine silt (median grain size $d_{50} = 0.0075\,\text{mm}$) to mimic high suspended sediment loads.
Foamy Water	Generated through controlled aeration to replicate surface foam and bubble accumulation observed in natural rivers.

TABLE 3.4 – Description of water background conditions during reflectance measurements.

3.1.11 Data Structure and Format

The dataset is archived in **Network Common Data Form (NetCDF)** and **Comma-Separated Values (CSV)** formats, ensuring compatibility with various machine learning and remote sensing applications. Each NetCDF file follows a standardized naming convention and contains components summarized in Table 3.5.

Component	Description
RGB Image	Coincident high-resolution imagery for visual validation.
Reflectance Values	Processed reflectance data across the full spectral range.
Labeled Image (Mask)	Pixel-wise segmentation masks identifying plastic debris, water, and foam regions.
Metadata	Includes polymer type, sediment concentration, flow discharge rates, and optical measurement conditions.

TABLE 3.5 – Contents of each NetCDF file.

In addition to the structured data, pixel-wise labeled segmentation masks are available, enabling the use of deep learning models for classification. Table 3.6 outlines the relevance of this dataset for machine learning applications.

Feature	**Machine Learning Relevance**
Segmentation Masks	Enables supervised training for semantic segmentation tasks.
U-Net Architecture Compatibility	Pixel-level labeling supports high-precision segmentation with U-Net models.
Extendability	Suitable for UAV and satellite-based remote sensing applications.

Table 3.6 – Machine learning relevance of dataset components.

3.1.12 Example of NetCDF Sample and What It Includes

As shown in Table 3.7 and **Figure 3.7**. the NetCDF file includes spectral, visual, and labeled data, along with comprehensive metadata describing the environmental conditions.

Component	**Description**
Reflectance	1D array of spectral reflectance values over 2151 bands (350–2500 nm).
RGB Image	Coincident RGB image (3 × 2992 × 2000), used for visual reference and spatial alignment.
Labeled Image	Pixel-level segmentation map identifying plastic, water, and foam regions.
Metadata	**Description**
Debris Polymer	PET
Debris Color	Green
Background Flow Status	Clear
Plastic Fraction (%)	6.8
Sediment Concentration (mg/L)	0.0
Flow Discharge (cms)	0.01
Quality Flag (1–5)	5.0

Table 3.7 – Example of a NetCDF hyperspectral file.

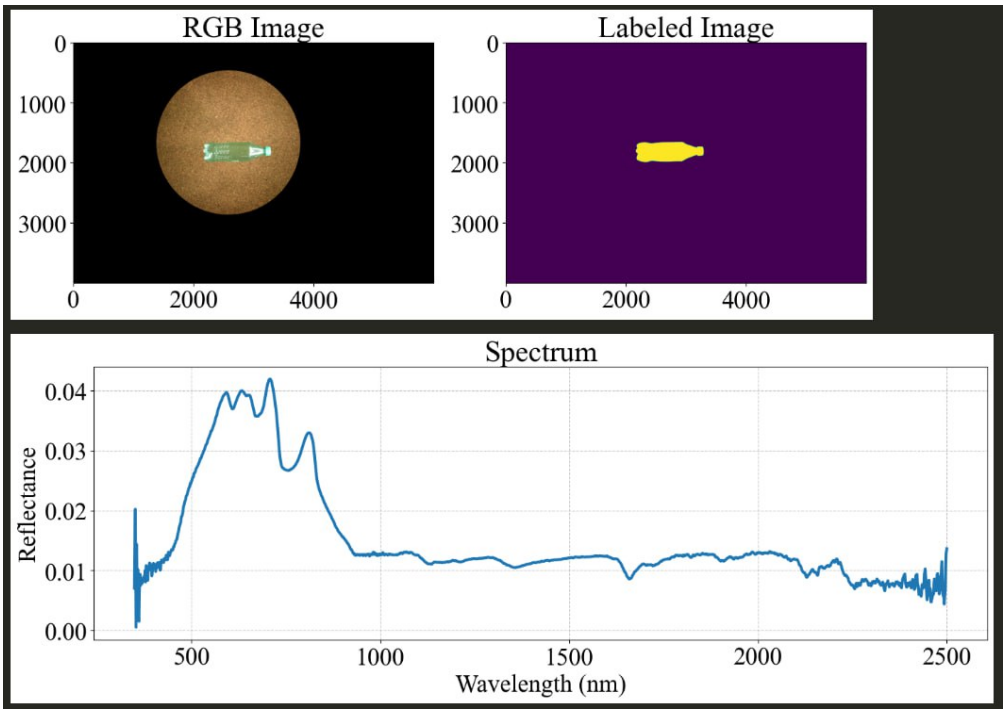

Figure 3.7 – NetCDF Content

3.1.13 Dataset and Preprocessing

For training and evaluation purposes, a total of 1500 samples were randomly selected for model training, while the remaining 578 samples were reserved for testing and validation.

Data preprocessing is a critical step to ensure data quality and optimize model performance. The pipeline comprises the following steps :

1. **RGB image resizing :** Each sample includes a high-resolution RGB image with original dimensions of (299220003). All images are resized to (2562563) using bilinear interpolation to standardize input dimensions and reduce computational overhead, while preserving key spatial features.
2. **Segmentation mask processing :** Pixel-wise labeled masks identify regions containing plastic debris. These masks are resized using nearest-neighbor interpolation to maintain class boundaries and ensure real-time processing efficiency. This method avoids the blending of labels that can occur with bilinear or bicubic interpolation.
3. **Reflectance normalization :** To standardize spectral intensity values and reduce the influence of acquisition variability, min–max scaling is applied to normalize all reflectance values to the $[0, 1]$ range. This enhances numerical stability during model training and ensures consistency across sessions.
4. **Class balancing :** The dataset exhibits class imbalance, with the 'Weathered' class (255 samples) overrepresented relative to the least frequent class, EPSF (92 samples). To mitigate this, the large 'Mix' class (1008 samples) is excluded, and the 'Weathered' class is randomly undersampled to 92 samples. This yields a balanced dataset of 907 samples used for support vector machine (SVM) classification.

In summary, the NetCDF dataset provides a total of 2078 hyperspectral reflectance samples acquired over a laboratory flume under three water background conditions : clear, turbid, and foamy. Each sample includes :

— A hyperspectral cube with 2151 spectral bands spanning 350–2500 nm ;
— A co-registered high-resolution RGB image ;
— A pixel-wise segmentation mask covering six polymer classes (PET, HDPE, LDPE, PP, EPSF, Weathered), water, and foam ;
— Metadata detailing polymer type, sediment concentration, and flow rate.

This richly annotated dataset enables both spectral pixel-level classification and spatial segmentation of plastic debris in varied aquatic conditions, supporting the development of robust machine learning models for plastic pollution monitoring.

3.2 Our Approaches

In this study, we evaluated four distinct approaches for plastic detection and classification using hyperspectral imaging data. The motivation behind implementing multiple methods lies in the desire to understand how different levels of model complexity and feature utilization affect segmentation and classification performance. Our goal was to compare classical machine learning techniques with modern deep learning methods, and to investigate how hybrid models could potentially combine their respective strengths.
We began with a PCA-SVM classifier to establish a spectral baseline, followed by a Random Forest model for faster classification. A 2D U-Net was then introduced to incorporate spatial-spectral segmentation. Finally, we developed a hybrid U-Net + SVM pipeline that combines spatial segmentation with refined spectral classification. This multi-approach strategy allowed us to compare classical and deep learning methods and highlight the strengths of hybrid modeling for precise plastic detection.

3.2.1 Used Metrics

To assess the performance our models, several standard evaluation metrics were employed. These metrics compare the predicted masks with the ground truth masks to quantitatively measure prediction quality.

1. **Accuracy**

   Accuracy represents the proportion of correctly classified pixels among all pixels :

   $$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{3.1}$$

   where :
   - TP : True Positives (plastic pixels correctly predicted),
   - TN : True Negatives (non-plastic pixels correctly predicted),
   - FP : False Positives (non-plastic pixels incorrectly predicted as plastic),
   - FN : False Negatives (plastic pixels incorrectly predicted as non-plastic).

2. **Precision**

   Precision measures the proportion of true positives among all predicted positives :

   $$\text{Precision} = \frac{TP}{TP + FP} \tag{3.2}$$

3. **Recall (Sensitivity)**

   Recall measures the model's ability to detect all actual plastic pixels :

   $$\text{Recall} = \frac{TP}{TP + FN} \tag{3.3}$$

4. **F1-Score**

   The F1-Score is the harmonic mean of precision and recall, providing a balance between the two :

   $$\text{F1-score} = 2\frac{\text{PrecisionRecall}}{\text{Precision} + \text{Recall}} \tag{3.4}$$

5. **Intersection over Union (IoU)**

   IoU, also known as the Jaccard Index, evaluates the overlap between predicted and ground truth masks :

   $$\text{IoU} = \frac{TP}{TP + FP + FN} \tag{3.5}$$

   IoU ranges from 0 to 1, where 1 indicates a perfect match of segmentation. IoU is particularly popular for segmentation tasks because it penalizes both false positives and false negatives and does not give undue weight to easily classified background pixels as shown in Figure 3.8 .

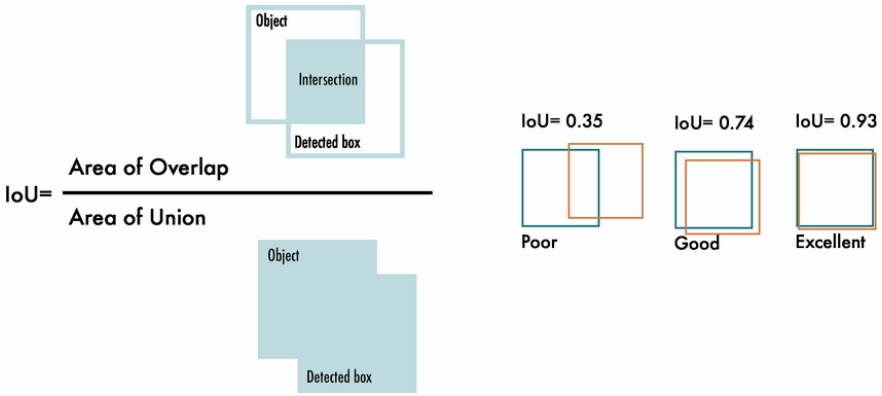

Figure 3.8 – Intersection over Union IoU [w6]

6. **Mean Average Precision (mAP)**

   Mean Average Precision (mAP) is a widely used metric for classification and object detection tasks. It measures the area under the precision-recall curve averaged over multiple classes or IoU thresholds :

$$\text{mAP} = \frac{1}{N} \sum_{i=1}^{N} \text{AP}_i \tag{3.6}$$

   where N is the number of classes or IoU thresholds, and AP_i is the Average Precision for class or threshold i, computed as the area under the precision-recall curve.

7. **Giga Floating Point Operations per Second (GFLOPS)**

   GFLOPS quantifies the computational performance of a system or model by measuring how many floating-point operations it can execute per second, in billions :

$$\text{GFLOPS} = \frac{\text{Total floating-point operations}}{\text{Execution time (s)} 10^9} \tag{3.7}$$

   For example, if a model performs 2.510^{11} floating-point operations in 0.5s, its throughput is

$$\text{GFLOPS} = \frac{2.510^{11}}{0.510^9} = 500\,\text{GFLOPS}.$$

8. **Confusion Matrix**

   In multiclass classification, the confusion matrix is an NN square matrix where N is the number of classes. Each row of the matrix represents the instances of the actual class, while each column represents the instances of the predicted class. The element at the i-th row and j-th column indicates the number of samples from class i that were predicted as class j.

   Formally, the confusion matrix C is defined as :

$$C = \begin{bmatrix} C_{11} & C_{12} & \cdots & C_{1N} \\ C_{21} & C_{22} & \cdots & C_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ C_{N1} & C_{N2} & \cdots & C_{NN} \end{bmatrix}$$

   where :

   — C_{ii} (diagonal elements) represents the number of correctly classified samples for class i.
   — C_{ij} (off-diagonal elements, $i \neq j$) represents the number of samples from class i misclassified as class j.

This matrix provides a detailed view of the classifier's performance, highlighting which classes are confused with each other. From this matrix, per-class precision, recall, and F1-scores can be computed by treating each class as the positive class in a one-vs-all manner.

The confusion matrix aids in diagnosing specific weaknesses in the classifier, such as systematic misclassification between certain classes, and is critical for comprehensive model evaluation in multi-class problems.

3.2.2 Principal Component Analysis (PCA)

Principal Component Analysis (PCA) is a widely used technique for reducing the dimensionality of high-dimensional datasets. It transforms a set of potentially correlated input variables into a smaller set of linearly uncorrelated variables known as *principal components*. Each principal component represents a linear combination of the original variables and is constructed to successively capture the maximum possible variance in the data. By retaining only the leading components that account for the majority of the variance, PCA effectively compresses the data while preserving its essential structure. This dimensionality reduction helps to alleviate issues such as multicollinearity and overfitting in downstream machine learning models, especially when working with hyperspectral data characterized by hundreds or thousands of spectral bands [w7].

3.2.3 First Approach : Support Vector Machines (SVM) + PCA

1. **SVM Definition**
   A Support Vector Machine (SVM) is a supervised learning algorithm used for classification and regression tasks. It operates by finding the hyperplane in a high-dimensional feature space that maximally separates data points of different classes. The optimal hyperplane is defined by support vectors—data points closest to the decision boundary—such that the margin between classes is maximized, improving generalization on unseen data [w8].
2. **SVM workflow**

Figure 3.9 – Our SVM Workflow

3. **Architecture**
   In this first approach, the hyperspectral data from the "Cubo_Cut" cube are processed using a standard *PCA+SVM* pipeline. Each pixel of the hyperspectral cube, represented by its spectral signature, is reshaped into a feature vector. After any necessary normalization or preprocessing, these vectors are assembled into a matrix $\mathbf{X} \in \mathbb{R}^{NB}$, where N is the total number of sampled pixels and B is the number of spectral bands.

   Next, **Principal Component Analysis (PCA)** is applied to this matrix to reduce dimensionality. Concretely, one computes the eigenvectors and eigenvalues of the covariance matrix of the spectra, then projects each spectral vector onto the first k eigenvectors (the *principal components*) that retain most of the data's variance. In practice, for hyperspectral data, it is common that the first 30–40 components capture over 95% of the total variance.

   Finally, the vectors obtained for each pixel are used as inputs to a **Support Vector Machine (SVM)** classifier. The idea is that, thanks to PCA, the principal components are uncorrelated and form a lower-dimensional space in which training an SVM is more efficient and less prone to overfitting. The overall pipeline can thus be described as follows :

   1. **Spectrum extraction and preprocessing :** Load the "Cubo_Cut" cube, extract each pixel as a spectral vector, and normalize or apply any specific preprocessing (filtering, masking of irrelevant pixels, etc.).
   2. **Dimensionality reduction via PCA :** Fit PCA on the training spectra, select the first k components (those covering the majority of variance), and project all vectors into this reduced k-dimensional space.
   3. **SVM classification :** Train an SVM classifier (typically with an RBF kernel for hyperspectral data) on the projected vectors. The classifier then learns to map each principal-component vector to the correct plastic class label.

   This configuration compresses the spectral information while preserving most variance, then performs classification in a reduced and better-separated space, improving robustness and reducing the SVM's computational complexity.
4. **Training** The training phase follows a standard supervised-learning procedure, strictly separating training and test data :

   1. **Data preparation**
      — From the "Cubo_Cut" hyperspectral cube, each pixel is converted into a vector.
      — Labeled pixels (generated from annotations or associated masks) are gathered into a matrix and their corresponding labels into a vector $\mathbf{y} \in \{1, \ldots, C\}^N$, where C is the number of classes (plastic types).
      — A random split (using `train_test_split` with `test_size=0.2` and a fixed `random_state`) yields a training subset $(\mathbf{X}_{\text{train}}, \mathbf{y}_{\text{train}})$ and a test subset $(\mathbf{X}_{\text{test}}, \mathbf{y}_{\text{test}})$.
   2. **Dimensionality reduction (PCA) on the training set**
      — PCA is fitted only on $\mathbf{X}_{\text{train}}$ to avoid any information leakage.
      — The number of components k is chosen so that the sum of variances explained by these k components exceeds a predetermined threshold (over 95%).
   3. **SVM training**
      — Initialize an SVM classifier (often with an RBF kernel). The SVM's hyperparameters (regularization parameter C, kernel parameter γ, etc.) can be tuned by cross-validation on the training set or set to default values for demonstration.
      — Train the SVM on $\{\mathbf{X}_{\text{train}}, \mathbf{y}_{\text{train}}\}$. The optimization seeks to maximize the margin between classes in the k-dimensional principal-component space.
   4. **Evaluation on the test set**
      — In this step, we compute predictions, overall accuracy, and plot the confusion matrix.

   In summary, the training procedure is :

   1. Extract and randomly split spectra into training and test sets.
   2. Fit PCA on the training set and project all spectral vectors onto the k principal components.
   3. Train the SVM on the reduced features, with optional hyperparameter tuning.
   4. Evaluate the model using overall accuracy and the confusion matrix on the test set.

   This methodology ensures that the SVM has no access to test information and provides quantitative metrics (accuracy, confusion matrix) to assess how well the model segments and classifies different plastics from hyperspectral images.

3.2.4 Second Approach : Random Forest + PCA

1. **Random Forest Definition**

   Random Forest is an ensemble learning algorithm that aggregates multiple decision trees to perform classification. In computer science terms, it constructs an array of tree-based classifiers, each of which is a binary tree data structure built on a random subset of the training samples and a random subset of input features. During inference, each tree independently traverses its decision nodes—implemented as nested `if-else` checks on feature values—to reach a leaf node containing a predicted class. The final output for each pixel is obtained by majority voting over all individual tree predictions [10].

   Internally, each decision tree is represented by node objects (or arrays of node attributes) storing :

   — `feature_index` : index of the spectral band used for splitting,
   — `threshold` : numeric value of the band's reflectance at which the split occurs,
   — `left_child` and `right_child` : indices (or pointers) to the child nodes,
   — `is_leaf` : boolean flag indicating whether the node is terminal,
   — `class_counts` or `class_label` : for leaf nodes, the distribution of training samples or the majority class label.

   These nodes are stored in contiguous arrays (NumPy arrays) for memory efficiency. By using array-based representations and integer indices, the implementation can leverage vectorized operations when possible, minimizing Python-level overhead.

2. **Random Forest Workflow**

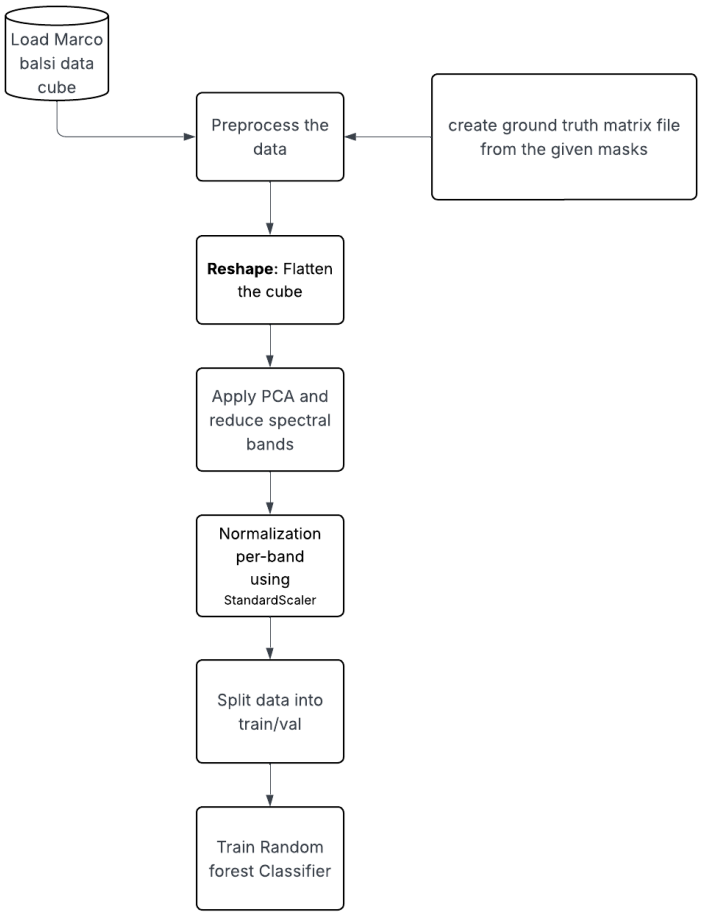

Figure 3.10 – Our Random Forest Workflow

3. **Architecture**
   In this approach, we used the Balsi–Moroni–Bouchelaghem Hyperspectral Plastic Detection Dataset. The Random Forest module is integrated at the end of the hyperspectral imaging (HSI) processing pipeline. The pipeline stages and the corresponding data structures are :

   1. **HSI Data Loading :** A hyperspectral scene is loaded as a three-dimensional array `cube`, with shape `(H, W, B)`, where `H` and `W` are spatial dimensions and `B` is the number of spectral bands. Internally, `cube` is stored as a contiguous NumPy array, enabling fast slicing and reshaping operations.
   2. **Flattening to Feature Matrix :** We reshape `cube` into `X`, a two-dimensional array of shape `(N, B)`, where `N = H * W`.
      Here, each row `X[i]` corresponds to the full spectral signature (length `B`) of the i-th pixel. The data type remains `float32` for efficient memory usage.
   3. **Label Array Creation :** Segmentation masks are loaded as a two-dimensional integer array `mask` with shape `(H, W)`, where each element is a class index (e.g., 0 for background, 1 for PET, etc.). We flatten `mask` into a one-dimensional label vector `y` of length `N`.

      This label vector aligns with `X` so that `X[i]` has label `y[i]`.
   4. **Preprocessing :** Depending on data quality, we may apply :

      Dimensionality Reduction : Principal Component Analysis (PCA) can be applied to `X` using matrix operations. PCA transforms `X` into `X_pca` of shape `(N, K)`, where `K < B`.

      This step converts `X` into a form suitable for robust tree splitting, changing the shape from `(N, B)` to `(N, K)`.

4. **Training**
    1. **Classifier Initialization :** The Random Forest classifier is instantiated using `scikit-learn`'s `RandomForestClassifier`. The parameters used are shown in Table 3.8 :

Parameter	Default	Purpose
`n_estimators`	100	Total number of decision trees in the forest
`max_depth`	20	Maximum depth of each tree
`n_jobs`	-1	Number of parallel jobs for tree construction (-1 = use all processors)

TABLE 3.8 – Random Forest hyperparameters and their roles in the algorithm

    2. **Evaluation on the test set :** The trained Random Forest is used to predict labels on $\mathbf{X}_{\text{test}}$. Overall accuracy is then computed by comparing the predicted labels $\hat{\mathbf{y}}_{\text{test}}$ with the true labels $\mathbf{y}_{\text{test}}$. Finally, a confusion matrix is generated to visualize the per-class performance and identify any misclassification patterns.

In summary, the training procedure is :

— Extract and randomly split spectra (flattened HSI pixels) into training and test sets.
— Apply any preprocessing steps (e.g., spectral filtering, normalization, PCA) to both $\mathbf{X}_{\text{train}}$ and $\mathbf{X}_{\text{test}}$, ensuring that PCA (if used) is fit only on $\mathbf{X}_{\text{train}}$.
— Initialize the `RandomForestClassifier` with chosen hyperparameters.
— Train the Random Forest on $\{\mathbf{X}_{\text{train}}, \mathbf{y}_{\text{train}}\}$.
— Predict on $\mathbf{X}_{\text{test}}$ and evaluate using overall accuracy and the confusion matrix.

3.2.5 Third Approach : Single-Phase 2D U-Net Approach (No PCA)

This section details the practical implementation of a single-phase hyperspectral plastic detection and classification method using a 2D U-Net. The focus is on data preparation, model architecture, and training workflow, all implemented in PyTorch, without the use of PCA.

1. **U-Net Overview**

    U-Net is a convolutional neural network designed for pixel-wise semantic segmentation. It features a symmetric encoder–decoder structure :

    — The encoder path extracts contextual features via successive convolutions and downsampling.
    — The decoder path upsamples feature maps to recover spatial resolution.
    — Skip connections directly link encoder and decoder layers at the same resolution, allowing precise localization.

    Originally proposed for biomedical segmentation [29], U-Net is effective for tasks with limited labeled data and complex spatial structures.
2. **Workflow Overview**
3. **Data and Preprocessing**

    — **Dataset :** Balsi–Moroni–Bouchelaghem Hyperspectral Plastic Detection Dataset.
    — **Hyperspectral Cubes :** Each sample $\mathcal{C} \in \mathbb{R}^{HWB}$, where B denotes the number of spectral bands. Ground-truth masks $\mathcal{M} \in \{0, 1, \ldots, C-1\}^{HW}$ provide pixel-wise class labels.
    — **Patch Extraction :** Cubes are split into overlapping patches of size $PP = 3232$ with stride $s = 16$. Each patch is (B, P, P), with corresponding label mask (P, P). Background-only patches are discarded or downsampled.
    — **Spectral Standardization :** Per-band mean and standard deviation are computed across training patches. Each band is normalized individually. No PCA is applied—preserving all B spectral bands as input channels.
    — **DataLoader :** A custom `Hyperspec2DDataset` class loads patches from NumPy/HDF5 :

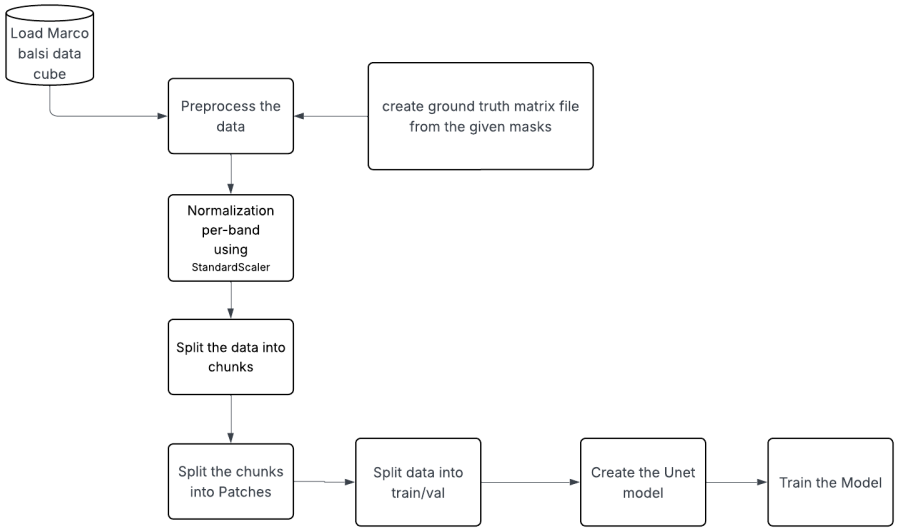

Figure 3.11 – Overview of the U-Net 2D segmentation workflow.

— Applies normalization
— Returns $\mathbf{X} \in \mathbb{R}^{BPP}$ and $\mathbf{Y} \in \{0, \ldots, C-1\}^{PP}$

Wrapped by a PyTorch `DataLoader` with `batch_size`=2.

4. **Model Architecture**

— **Input :** (B, P, P) tensor.
— **Encoder :**
  — **Level 1 :** Two 33 `Conv2d` layers with BatchNorm + ReLU, followed by 22 MaxPooling. Output : `base_ch` channels.
  — **Level 2 :** Same as above, with channels doubled.
— **Bottleneck :** Two 33 conv layers with BatchNorm + ReLU. Output channels : 4`base_ch`.
— **Decoder :**
  — **Up Block 1 :** `ConvTranspose2d` upsample, concatenate with `enc2` output, followed by two 33 conv layers.
  — **Up Block 2 :** Similar to above, concatenated with `enc1`.
— **Output Layer :** `Conv2d(base_ch, out_channels, 1)`. Final activation : `Softmax` (multi-class) or `Sigmoid` (binary).
— **Notes :**
  — All conv layers use padding=1 to maintain spatial dimensions.
  — Upsampling uses `ConvTranspose2d(2, stride=2)`.
  — Default : `base_ch` = 16. Increase to 32 if B is large (e.g., 128).
  — 10M trainable parameters for typical settings.

5. **Training Setup**

— **Hyperparameters :**
  — Epochs : 200
  — Batch size : 2
  — Learning rate : 110^{-4}
  — Optimizer : Adam
  — Loss function : Cross-entropy
— **Environment :**

- GPU : NVIDIA Quadro RTX 6000 (24 GB)
- PyTorch 1.10+, CUDA 11.3, cuDNN 8.2
- Training time : 83 minutes (200 epochs)

— **Monitoring :**

- Loss and accuracy plotted with `matplotlib`
- Metrics : Validation accuracy, F1-score
- No early stopping

6. **Evaluation and Visualization**

— **Metrics :** Overall Accuracy (OA), F1-score, Intersection over Union (IoU).

— **Visualization :**

- Color-coded segmentation maps overlaid on RGB composites
- Confusion matrices and per-class IoU plots

7. **Implementation Highlights**

— **Memory Optimization :**

- Reduce `base_ch` or batch size if memory is limited.
- Optionally divide cubes into smaller chunks.

— **No PCA Justification :**

- Full-band input preserves spectral richness and enables better end-to-end learning.
- Outperformed PCA-based input in validation accuracy.
- Tradeoff : higher GPU memory usage, longer training.

8. **Summary**

The 2D U-Net (No PCA) pipeline is a robust and modular approach for hyperspectral plastic segmentation :

— **Input :** (B, P, P) hyperspectral patches.
— **Model :** Standard U-Net (2 encoders, bottleneck, 2 decoders).
— **Training :** 200 epochs, Adam optimizer, batch size 2.
— **Evaluation :** Accuracy, F1, IoU ; segmentation maps and confusion matrices.

3.2.6 Fourth Approach : Two-Phase Segmentation Pipeline

Our approach adopts a two-phase pipeline designed to balance efficiency and accuracy in plastic detection. In the first phase, we utilize a U-Net model trained on RGB imagery to perform coarse segmentation of plastic regions. This step takes advantage of the fast acquisition and high spatial resolution of RGB images, allowing rapid identification of candidate plastic areas. However, due to the limited spectral information of RGB, this phase often introduces false positives.

To address this, the second phase introduces a refinement step based on hyperspectral data. Specifically, we extract the spectral signatures of the regions identified in Phase 1 and classify them using a Support Vector Machine (SVM) trained to distinguish plastic from background based on hyperspectral reflectance profiles. This spectral verification step significantly improves precision by filtering out misclassified regions from the initial segmentation.

By combining the speed of RGB-based segmentation with the discriminative power of hyperspectral analysis, our approach achieves both efficient processing and accurate plastic detection. The two phases are described below :

1. **Phase 1 — U-Net Segmentation :**
   This phase employs a U-Net architecture trained on RGB images to generate an initial segmentation mask that separates plastic from background. The use of RGB imagery offers rapid processing due to its lower dimensionality and widespread availability. Visually salient plastic regions—often characterized by bright colors or distinct shapes—are effectively captured, leading to a relatively high recall. However, this phase is susceptible to reduced precision, as visually similar non-plastic materials may also be detected, leading to false positives.
2. **Phase 2 — Spectral SVM Classification :**
   To refine the initial mask, we use hyperspectral data in a second phase. Spectral signatures are extracted from the pixels or patches marked as plastic in the RGB-based mask. These signatures are then classified using a Support Vector Machine (SVM) trained to distinguish plastic from background based on hyperspectral

reflectance. The SVM reevaluates each candidate region : only those that exhibit characteristic plastic spectra are retained. This spectral filtering step significantly enhances segmentation precision by eliminating visually similar but spectrally different false positives. Although it's possible to reprocess the full image to recover missed regions (false negatives), our current implementation focuses on improving precision by refining only the detected regions.

This two-phase strategy offers an effective compromise between speed and accuracy in plastic segmentation, leveraging the strengths of both RGB and hyperspectral imaging modalities.

3.2.7 Two-Phase Workflow

To visually summarize our approach, Figure 3.12 illustrates the overall two-phase segmentation pipeline adopted in this work. This workflow combines the efficiency of RGB-based segmentation with the spectral precision of hyperspectral analysis, resulting in a more accurate and robust plastic detection system.

In the first phase, an RGB image is fed into a U-Net model that performs an initial segmentation to detect potential plastic regions. This step leverages the rapid acquisition and high resolution of RGB imagery to quickly identify visually distinguishable plastic candidates. However, due to limited spectral fidelity, this phase may introduce false positives.

In the second phase, the candidate regions detected in Phase 1 are further analyzed using hyperspectral data. Spectral features are extracted from the corresponding regions, and a Support Vector Machine (SVM) classifier is used to validate whether each region truly corresponds to plastic. This spectral verification step filters out false positives, thus increasing the overall precision of the detection.

This modular workflow not only improves detection performance but also allows for flexibility in processing. The RGB phase can be deployed in real-time or on resource-constrained

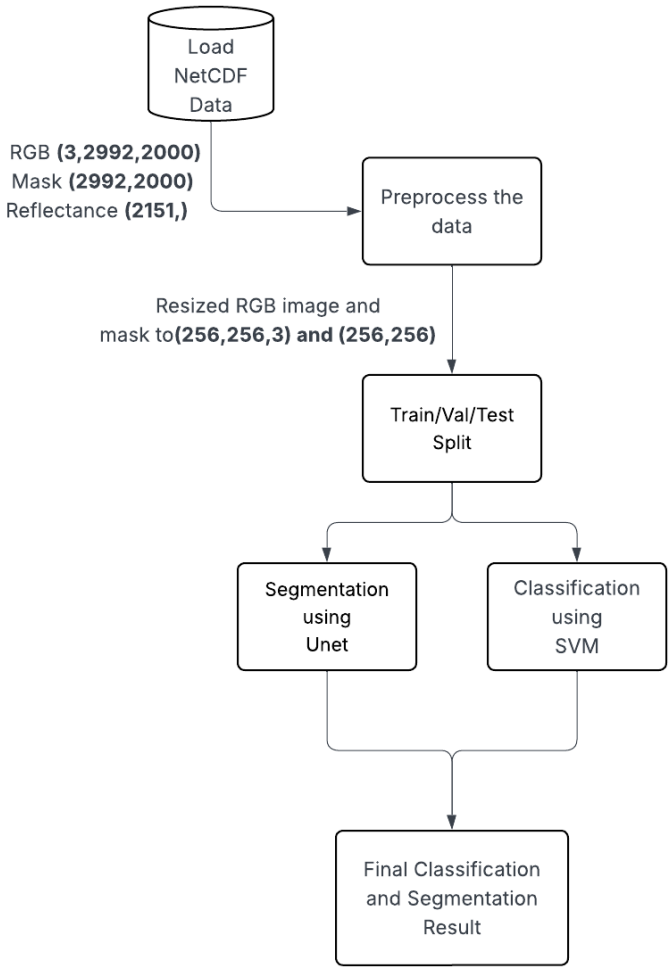

Figure 3.12 – Overview of our Two-Phase Workflow for plastic segmentation. Phase 1 uses a U-Net model on RGB images to detect candidate plastic regions. Phase 2 refines these predictions using hyperspectral classification with an SVM.

3.2.8 Architecture

Segmentation using U-Net U-Net is employed for plastic debris segmentation due to its powerful encoder–decoder architecture, which is well-suited for semantic segmentation tasks where precise, pixel-level predictions are required. Originally designed for biomedical image segmentation, U-Net has demonstrated strong performance across a wide range of domains, including environmental monitoring and remote sensing. Its strength lies in the symmetric structure that captures both low-level spatial features and high-level contextual information. The encoder progressively reduces spatial dimensions while increasing feature abstraction, whereas the decoder restores spatial resolution to produce detailed segmentation masks. Crucially, U-Net incorporates skip connections between corresponding encoder and decoder layers, enabling the network to retain fine-grained information that is often lost during downsampling. This combination of deep feature learning and spatial precision makes U-Net particularly effective for identifying irregularly shaped and distributed plastic debris in RGB imagery. The network is composed of the following components :

1. **Encoder Path :** The downsampling section comprises four convolutional blocks. Each block includes two convolutional layers with 33 kernels, followed by batch normalization and ReLU activation. A max-pooling layer at the end of each block progressively reduces the spatial resolution while preserving essential features for accurate segmentation.
2. **Bottleneck Layer :** Situated at the deepest point of the network, the bottleneck captures high-level abstract

features through multiple feature maps. Dropout regularization is applied to prevent overfitting and enhance the model's generalization ability.

3. **Decoder Path :** The upsampling section reconstructs spatial resolution using transposed convolutional layers. At each level, the upsampled output is concatenated with the corresponding feature map from the encoder path via skip connections, enabling the recovery of fine-grained spatial details.
4. **Output Layer :** A final 11 convolutional layer, followed by a sigmoid activation function, produces a pixel-wise probability map that highlights regions likely to contain plastic debris.

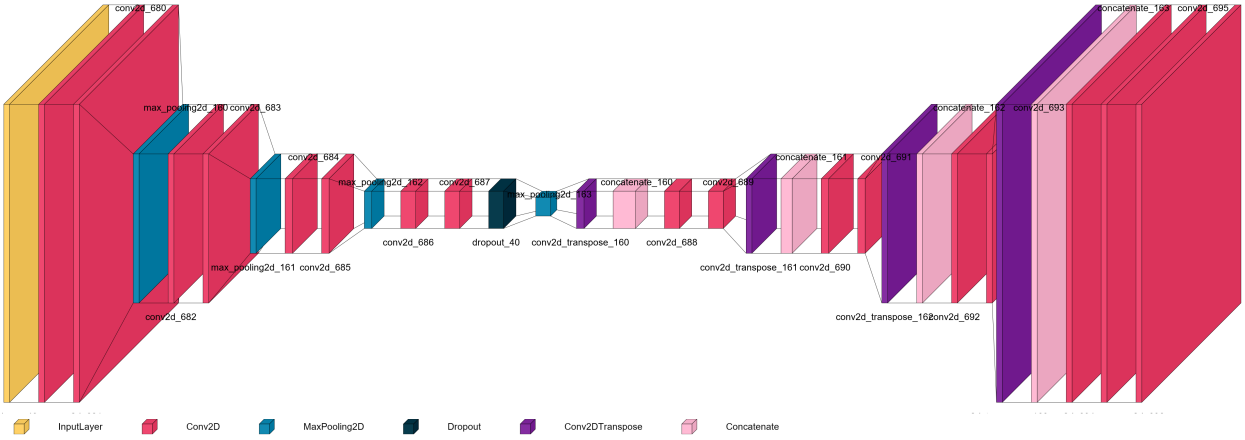

Figure 3.13 – Architecture of the U-Net model used in Phase 1 for coarse plastic segmentation with layer-wise dimensions. The network consists of an encoder-decoder structure with symmetric skip connections. Feature maps from the encoder (left) are concatenated with corresponding decoder layers (right), as indicated by the pink boxes labeled "Concat." These skip connections help preserve spatial information and enhance segmentation accuracy. Each block represents a convolutional layer or a composite operation (e.g., convolution + ReLU + pooling or upsampling).

As illustrated in figure 3.13, the U-Net segmentation network follows a symmetric encoder-decoder architecture for coarse plastic segmentation. The left side shows the contracting path (encoder), and the right side shows the expanding path (decoder), connected via skip connections. These skip connections *visualized as pink boxes labeled "Concat"* enable precise localization by combining high-resolution encoder features with decoder outputs. Each encoder block applies two 3×3 convolutions (ReLU) followed by 2×2 max-pooling, doubling the number of feature channels at each step. The bottleneck consists of two 3×3 convolutional layers (256 filters) with dropout in between. In the decoder, each upsampling block uses a 2×2 transposed convolution to halve the number of channels, concatenates the corresponding encoder features via skip connections, and applies two 3×3 convolutions (ReLU). A final 1×1 convolution produces the single-channel output mask using a sigmoid activation. The numbers and labels in the diagram indicate the dimensions and layer names, helping clarify the flow of information throughout the network.

Layer-by-Layer Architecture Description The algorithm of the U-Net model used for plastic debris segmentation is constructed as follows :

Algorithm 1 U-Net Architecture for Image Segmentation

Require: Input image $I \in \mathbb{R}^{2562563}$
Ensure: Output segmentation map $O \in \mathbb{R}^{2562561}$
Encoder Path :
$E_1 \leftarrow \text{ReLU}(\text{Conv2D}(I, 16, 33))$
$E_1 \leftarrow \text{ReLU}(\text{Conv2D}(E_1, 16, 33))$
$S_1 \leftarrow E_1$
$E_1 \leftarrow \text{MaxPool}(E_1, 22)$
$E_2 \leftarrow \text{ReLU}(\text{Conv2D}(E_1, 32, 33))$
$E_2 \leftarrow \text{ReLU}(\text{Conv2D}(E_2, 32, 33))$
$S_2 \leftarrow E_2$
$E_2 \leftarrow \text{MaxPool}(E_2, 22)$
$E_3 \leftarrow \text{ReLU}(\text{Conv2D}(E_2, 64, 33))$
$E_3 \leftarrow \text{ReLU}(\text{Conv2D}(E_3, 64, 33))$
$S_3 \leftarrow E_3$
$E_3 \leftarrow \text{MaxPool}(E_3, 22)$
$E_4 \leftarrow \text{ReLU}(\text{Conv2D}(E_3, 128, 33))$
$E_4 \leftarrow \text{ReLU}(\text{Conv2D}(E_4, 128, 33))$
$S_4 \leftarrow E_4$
$E_4 \leftarrow \text{MaxPool}(E_4, 22)$
Bottleneck :
$B \leftarrow \text{ReLU}(\text{Conv2D}(E_4, 256, 33))$
$B \leftarrow \text{Dropout}(B, 0.3)$
$B \leftarrow \text{ReLU}(\text{Conv2D}(B, 256, 33))$
Decoder Path :
$D_4 \leftarrow \text{Conv2DTranspose}(B, 128, 22, \text{stride} = 2)$
$D_4 \leftarrow \text{Concat}(D_4, S_4)$
$D_4 \leftarrow \text{ReLU}(\text{Conv2D}(D_4, 128, 33))$
$D_4 \leftarrow \text{ReLU}(\text{Conv2D}(D_4, 128, 33))$
$D_3 \leftarrow \text{Conv2DTranspose}(D_4, 64, 22, \text{stride} = 2)$
$D_3 \leftarrow \text{Concat}(D_3, S_3)$
$D_3 \leftarrow \text{ReLU}(\text{Conv2D}(D_3, 64, 33))$
$D_3 \leftarrow \text{ReLU}(\text{Conv2D}(D_3, 64, 33))$
$D_2 \leftarrow \text{Conv2DTranspose}(D_3, 32, 22, \text{stride} = 2)$
$D_2 \leftarrow \text{Concat}(D_2, S_2)$
$D_2 \leftarrow \text{ReLU}(\text{Conv2D}(D_2, 32, 33))$
$D_2 \leftarrow \text{ReLU}(\text{Conv2D}(D_2, 32, 33))$
$D_1 \leftarrow \text{Conv2DTranspose}(D_2, 16, 22, \text{stride} = 2)$
$D_1 \leftarrow \text{Concat}(D_1, S_1)$
$D_1 \leftarrow \text{ReLU}(\text{Conv2D}(D_1, 16, 33))$
$D_1 \leftarrow \text{ReLU}(\text{Conv2D}(D_1, 16, 33))$
Output Layer :
$O \leftarrow \text{Sigmoid}(\text{Conv2D}(D_1, 1, 11))$
return O

Each layer in the architecture is specified by its type, input and output dimensions, activation function, kernel and filter parameters, as well as its specific functional role within the network. This configuration adheres to the canonical U-Net framework, wherein skip connections facilitate the preservation and integration of spatial features between the encoding and decoding stages, thereby enhancing localization accuracy.

The model training employs the Adam optimization algorithm, with the learning rate determined through systematic hyperparameter tuning to optimize convergence. Given the binary nature of the segmentation task "*distinguishing plastic debris from background*" the binary cross-entropy loss function is utilized. A moderate batch size is selected to achieve a balance between computational efficiency and generalization capability. Following training, the model demonstrates robust segmentation performance, quantitatively evaluated using accuracy and Intersection over Union (IoU) metrics. Comprehensive details regarding the training configuration and resulting performance are presented in Table 3.9.

Classification Using Support Vector Machine (SVM) Following the initial segmentation of plastic debris, a subsequent classification stage is implemented to refine the identification of plastic regions. This stage employs a Support Vector Machine (SVM) classifier to leverage the high-dimensional spectral information provided by hyperspectral imaging.

The classification pipeline comprises the following key steps :

1. **Feature Extraction :** Hyperspectral reflectance values are extracted from the segmented regions identified as potential plastic debris. These spectral features capture detailed material-specific signatures across multiple wavelength bands.
2. **Feature Normalization :** The extracted features are standardized via z-score normalization, ensuring zero mean and unit variance for each spectral band. This normalization process mitigates the influence of varying scales among spectral features and enhances classifier stability.
3. **SVM Training :** An SVM with a radial basis function (RBF) kernel is trained on the normalized hyperspectral features. The RBF kernel facilitates nonlinear decision boundaries, allowing effective discrimination between plastic and non-plastic classes in the complex spectral feature space.

To optimize classification performance, hyperparameters including the regularization parameter C and the kernel bandwidth gamma (γ) are fine-tuned using a grid search strategy combined with cross-validation. These adjustments balance the trade-off between model complexity and generalization capability.

The principal hyperparameter settings and the resulting classification performance metrics are detailed in Table 3.9.

U-Net Segmentation Model		
Parameter	**Value**	**Context**
Convolution Kernel Size	3x3	Kernel size for all convolutional layers to capture spatial features at fine scale
Max Pooling Size	2x2	Downsampling operation in encoder blocks to reduce spatial dimensions while preserving salient features
Dropout Rate	0.3	Applied in bottleneck layer to prevent overfitting by randomly disabling neurons during training
Output Activation	Sigmoid	Used in final layer to produce pixel-wise probabilities for binary segmentation mask generation
Optimizer	Adam	Adaptive gradient-based optimizer employed for efficient and stable convergence during training
Learning Rate	10^{-4}	Learning rate empirically determined to balance convergence speed and training stability
Loss Function	Binary Cross-Entropy	Suitable for binary classification tasks, measures difference between predicted probabilities and ground truth
Batch Size	8	Mini-batch size selected to optimize GPU memory usage and generalization performance
SVM Classification Model		
Parameter	**Value**	**Context**
Feature Scaler	StandardScaler	Standardizes features to zero mean and unit variance to improve SVM training stability and performance
Kernel Type	RBF	Nonlinear kernel enabling separation of complex spectral data by mapping into higher-dimensional space
Regularization Parameter (C)	10	Controls the trade-off between maximizing margin and minimizing classification error
Gamma (γ)	scale	Kernel coefficient determining influence range of single training examples, set adaptively based on data variance

TABLE 3.9 – Summary of training parameters and results for U-Net segmentation and SVM classification.

4 Challenges and Lessons Learned

— **High Dimensionality and Computational Load**

— Hyperspectral cubes contain dozens or hundreds of spectral bands per pixel, which makes both memory usage and processing time significant. Training U-Net on full spectral inputs (without PCA) requires careful GPU memory management (e.g., reducing batch size or dividing the data into chunks and chunks into patches). Likewise, applying PCA on very large NB matrices (where $N = HW$) can become a bottleneck if not batch-processed or if the covariance matrix is computed naively.

— **Lesson Learned :** Applying PCA or band selection as a preprocessing step (when using SVM or Random Forest) drastically reduces training time and memory footprint. For the U-Net (no PCA) variant, patch-based loading and on-the-fly standardization helped mitigate GPU exhaustion but required careful tuning of patch size and overlap stride.

— **Class Imbalance and Rare Plastic Types**

— In our datasets ("Cubo_Cut" and the seven Balsi–Moroni–Bouchelaghem cubes), certain plastic classes (e.g., PVC or PS) occupy very few pixels compared to the background or mixed plastic. This imbalance caused poor recall on underrepresented classes when training SVMs or Random Forests on flattened pixel spectra, and led U-Net to predict "background" for small plastic regions.

— **Lesson Learned :** Implementing patch sampling strategies that oversample patches containing rare classes (for U-Net), or using class-balanced SVM sampling, improved per-class performance.

— **Spectral Variability Across Cubes and Sensors**

— Each hyperspectral cube (e.g., "cube21Feb24a.mat" versus the seven Balsi–Moroni–Bouchelaghem cubes) was acquired under different lighting and environmental conditions, leading to shifts in reflectance distributions. Even after per-band standardization, classifiers trained on one cube often generalized poorly to another, especially in the case of SVM and Random Forest.

— **Lesson Learned :** Leave-One-Cube-Out (LOCO) cross-validation appears to be a promising solution to this problem. It highlighted the need for robust normalization and careful memory management, and will be implemented in future work.

5 Web Interface and API Deployment

To enhance the usability and real-world accessibility of our proposed two-phase detection framework, we developed a web-based interface connected to a back-end API. This application allows users to upload hyperspectral images and receive segmentation and classification outputs along with other relevant metadata. The back-end is powered by our trained hybrid model (U-Net + SVM), which handles both spatial segmentation and refined polymer-level classification. The API was built using Python's `Flask` framework for speed and scalability, and the front-end features a user-friendly interface to simplify interaction.

This deployment demonstrates the feasibility of integrating advanced hyperspectral analysis into practical systems such as environmental monitoring platforms, UAVs, or smart recycling stations.

6 Conclusion

The experiments conducted in this chapter demonstrate that combining multiple machine learning strategies—including classical methods (SVM, Random Forest) and deep learning (U-Net)—yields complementary strengths for hyperspectral plastic segmentation and classification. The PCA + SVM approach offers fast training and interpretable decision boundaries when dimensionality is appropriately managed. Similarly, Random Forest provides robustness to noisy labels and can exploit intra-band correlations if properly tuned. The single-phase 2D U-Net (without PCA) excels at capturing both spatial and spectral features in an end-to-end manner, but it requires careful attention to GPU memory and patch sampling to address class imbalance.

The two-phase pipeline—RGB U-Net followed by spectral SVM refinement—achieves high precision by combining the speed of RGB-based segmentation with the discriminative power of hyperspectral verification.

Nevertheless, each method presents trade-offs : classical classifiers rely on engineered preprocessing steps (e.g., PCA), while U-Net demands large annotated datasets and significant hardware resources. The two-phase pipeline reduces false positives but introduces additional complexity. Overall, the lessons learned indicate that—when data variability across cubes is high—it is beneficial to adopt normalization strategies tailored to each cube. Balancing class

distribution through targeted sampling remains critical across all approaches, especially for underrepresented plastic types.

CHAPTER 04

Implementation and Results

(Detailed implementation steps and evaluation outcomes)

1 Introduction

This chapter presents and discusses the experimental results obtained from the proposed methodologies for plastic segmentation and classification using hyperspectral imaging. The evaluation includes traditional machine learning models such as Support Vector Machine (SVM) and Random Forest (RF), a deep learning-based U-Net model, and a hybrid two-phase approach combining U-Net for segmentation followed by SVM for classification. The experiments were conducted on different hyperspectral datasets, including 'cubo_cut.mat', the BalsiMoroniBouchelaghem dataset (composed of 7 cubes), and NetCDF formatted data ('.nc4'). Each method's performance is analyzed in terms of accuracy, segmentation quality, and classification precision.

2 Hardware Environment

2.1 Local Computational Setup

A significant portion of the model training and evaluation was performed on a high-performance HP Z8 G4 Workstation, which provided the necessary computational resources to efficiently process and analyze high-dimensional hyperspectral data. The specifications of the workstation are as follows :

— **Processor :** Intel® Xeon® Gold 6138 with 40 cores and 80 threads
— **RAM :** 128 GB DDR4
— **GPU :** NVIDIA Quadro RTX 6000 with 24 GB GDDR6 VRAM

This powerful setup allowed for the smooth execution of data preprocessing pipelines, including spectral normalization, PCA-based dimensionality reduction, and patch extraction. Moreover, it significantly accelerated the training of deep learning models like U-Net by enabling batch-based GPU computation, thus reducing the overall training time and avoiding memory bottlenecks when working with large hyperspectral cubes.

2.2 Cloud-Based Computational Setup (Google Colab)

Definition : Google Colab Colaboratory (or "Colab") is a cloud-based interactive notebook environment that allows you to write and execute Python code in your browser. It requires no setup, provides free access to GPUs, and is easy to share. Colab is a great tool for students, data scientists, and AI researchers.

In addition to the local workstation, some deep learning experiments—particularly those involving large-scale models or extensive hyperparameter tuning—were conducted on Google Colaboratory.

For this study, the Colab Pro+ tier was utilized, which granted access to the NVIDIA A100 Tensor Core GPU, this GPU includes :

— 40 GB of high-bandwidth HBM2 memory
— 6,912 CUDA cores and 432 Tensor cores

Using Google Colab provided a cost-effective and scalable solution for training deep models on large hyperspectral datasets, especially when the local workstation was already occupied or when remote collaboration and sharing of Jupyter notebooks were required. The Colab environment also facilitated rapid experimentation through seamless integration with Google Drive and popular Python libraries such as PyTorch, TensorFlow, and scikit-learn.

3 Software Environment

The processing pipelines and models were implemented in Python, leveraging a range of libraries for data I/O, preprocessing, machine learning, deep learning, visualization, and utility functions. Below is a detailed list of all software components used in the provided code, along with brief definitions of each.

3.1 Programming Language

3.1.1 Python

Python is a high-level, interpreted programming language that supports object-oriented design and dynamic typing. Its built-in data structures and flexible binding make it ideal for rapid application development as well as for scripting

or integrating existing components. Python's clean, easy-to-learn syntax promotes code readability and simplifies maintenance. By offering a modular design through packages and modules, it encourages reuse and organization. The core interpreter and comprehensive standard library are freely available both as source code and binaries across all major operating systems.

3.2 Used libraries

3.2.1 Pytorch

PyTorch is an open-source, Python-based deep learning framework that provides native support for GPU acceleration and dynamic computation graphs via reverse-mode automatic differentiation. Its flexible architecture makes it straightforward for practitioners to construct, train, and debug complex neural networks, which is particularly valuable for rapid experimentation in domains such as computer vision and natural language processing[w9].

3.2.2 NumPy

NumPy (Numerical Python) is a free, open-source library at the heart of scientific computing in Python. It provides a highly optimized implementation of large, multi-dimensional arrays (also known as matrices or tensors) alongside a rich suite of mathematical routines—ranging from basic linear algebra and statistical functions to Fourier transforms and random number generation. Building on the legacy of the earlier Numeric and Numarray projects, NumPy delivers fast, vectorized operations on arrays and serves as the foundational layer for many other libraries, including Pandas, SciPy, and scikit-learn. NumPy is developed by a broad community of contributors and supported by NumFOCUS[w10].

3.2.3 h5py

The h5py library provides a Pythonic interface to the HDF5 file format, enabling efficient storage and manipulation of very large numerical datasets directly from NumPy. It allows users to treat multi-terabyte arrays on disk as if they were in-memory NumPy arrays—supporting slicing, indexing, and metadata inspection (e.g. via .shape and .dtype)—while organizing thousands of datasets within a single file using familiar Python constructs such as dictionaries and array syntax[w11].

3.2.4 xarray

formerly known as xray is an open-source Python library designed to simplify and streamline working with labeled multi-dimensional arrays. By layering dimensions, coordinates, and attributes atop NumPy-style arrays, it offers a more intuitive, concise, and robust development experience. The package provides a growing suite of domain-agnostic functions for advanced analytics and visualization, was inspired by pandas' approach to labeled tabular data, and is optimized for handling NetCDF files and scaling analyses via tight integration with Dask[w11].

3.2.5 pandas

pandas is an open-source Python library that delivers high-performance, easy-to-use data structures and tools for handling tabular data. It extends Python with intuitively designed objects—such as DataFrames and Series—that simplify common tasks like loading, aligning, merging, and transforming datasets. The name "pandas" is derived from "panel data," reflecting its origin in econometrics and its strength in managing time-series and multi-dimensional data. Under the hood, pandas relies on optimized C and Cython routines for compute-heavy operations, though its single-threaded nature can limit scalability on multi-core systems. Nevertheless, pandas remains a cornerstone of the Python data-analysis ecosystem, seamlessly integrating with libraries like NumPy, Matplotlib, and scikit-learn to turn messy or raw data into clean, analysis-ready formats[w12].

3.2.6 scikit-learn

scikit-learn is an open-source Python library that streamlines machine learning workflows with a consistent, easy-to-use API. Built on NumPy, SciPy, and Matplotlib, it supports tasks like preprocessing, regression, classification, clustering, and dimensionality reduction with efficient implementations of popular algorithms (e.g., SVMs, random forests, etc). Performance-critical code is written in Cython/C++, while Python handles the high-level interface. Compatible across Linux, macOS, and Windows, scikit-learn integrates seamlessly with pandas and Matplotlib and is maintained by a large community on GitHub[w12].

3.2.7 joblib

Joblib is a Python library designed to accelerate heavy computations by enabling easy parallelization of tasks. It offers utilities for distributing workloads across multiple CPU cores and for caching the outputs of expensive function calls. This makes Joblib particularly valuable in machine learning workflows, where it can persist intermediate results and allow experiments to be paused and resumed—potentially on different machines—without re-running the entire computation[w13].

3.2.8 TensorFlow

TensorFlow is a widely adopted open-source machine-learning platform that represents computations as data-flow graphs : each node performs a mathematical operation, and each edge carries multidimensional arrays (tensors) between them. This graph-based design lets you define, train, and run models on CPUs, GPUs, or TPUs—whether on mobile devices, desktops, or large-scale servers—without changing your code. Originally developed by Google's Brain Team for deep neural-network research, TensorFlow's flexible architecture and unified toolsets make it accessible to data scientists, software engineers, and educators across diverse fields, streamlining collaboration and accelerating development[w14].

3.2.9 OpenCV

OpenCV is a free, Apache 2-licensed library with 2,500+ optimized computer-vision and machine-learning algorithms for tasks like recognition, tracking, 3D reconstruction, and image stitching. It offers C++, Python, Java, and MATLAB APIs, runs on all major OSes, and is optimized for real-time use with CPU vector instructions and GPU support[w15].

3.2.10 matplotlib

Matplotlib is a widely used open-source Python library for creating high-quality visualizations. Originally created by neurobiologist John D. Hunter in 2002 to replicate MATLAB's plotting capabilities for visualizing epileptic brain signals, it has since been maintained and enhanced by an active community of contributors. Matplotlib enables users to generate a broad range of plots—line graphs, histograms, bar charts, scatterplots, and more—with just a few lines of code. It integrates smoothly with NumPy and can be embedded into web servers, Python shells, standalone scripts, or graphical user interfaces via its versatile APIs[w16].

3.2.11 NetCDF

NetCDF (Network Common Data Form) is both a suite of software libraries and a platform-independent data format designed for the creation, access, and sharing of multi-dimensional array-based scientific datasets. As an established community standard, NetCDF facilitates interoperability across applications and institutions. The Unidata Program Center provides and maintains NetCDF APIs for C, C++, Java, and Fortran, with additional language bindings available for Python, IDL, MATLAB, R, Ruby, and Perl[w17].

4 Results and discussion

4.1 SVM Approach

4.1.1 Training Performance and Visualization

An SVM was established using PCA-reduced hyperspectral features, consisting of 15 components, and utilizing an 80/20 train–test split. Upon fitting the model, it attained a training accuracy of roughly 96.8%. Detailed performance metrics for each polymer class are provided in Table 4.1.

Polymer Type	Precision	Recall	F1-score
PET	0.79	1.00	0.88
PE	0.75	1.00	0.86
PP	0.82	1.00	0.90
PS	0.92	1.00	0.96
PVC	0.84	1.00	0.91
Mean	**0.82**	**1.00**	**0.90**

TABLE 4.1 – SVM Classification Performance by Polymer Type

4.1.2 Results

The SVM result is shown in Figure 4.1

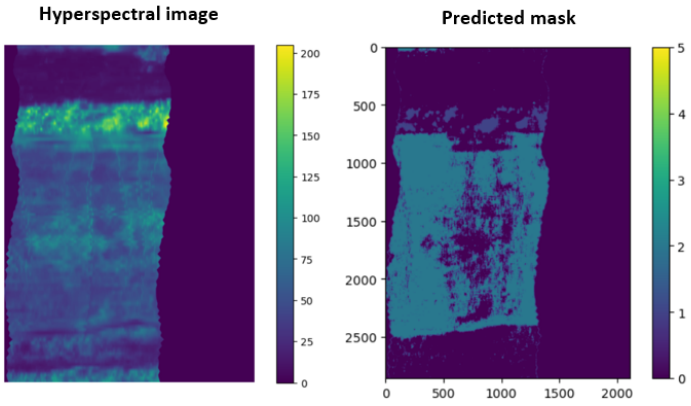

FIGURE 4.1 – Accuracy and loss during training epochs.

4.2 Random Forest Approach

4.2.1 Training Performance and Visualization

The Random Forest pipeline was set up on the same PCA-compressed hyperspectral features (15 components) with a 80/20 train–test split and 100 trees in the ensemble. After fitting, the model achieved a training accuracy of approximately 98%, detailed performance metrics for each polymer class presented in Table 4.2.

Polymer Type	Precision	Recall	F1-score
PET	0.96	0.97	0.96
PE	0.96	0.98	0.97
PVC	0.99	0.97	0.98
PP	0.98	0.95	0.97
PS	0.98	0.96	0.97
Mean	**0.97**	**0.97**	**0.97**

TABLE 4.2 – Random Forest Classification Performance by Polymer Type

The confusion matrix (Figure 4.2) reaffirms this strong performance—each polymer class sits firmly on the diagonal, with the only appreciable confusions occurring between PE and PP, which share adjacent spectral bands.

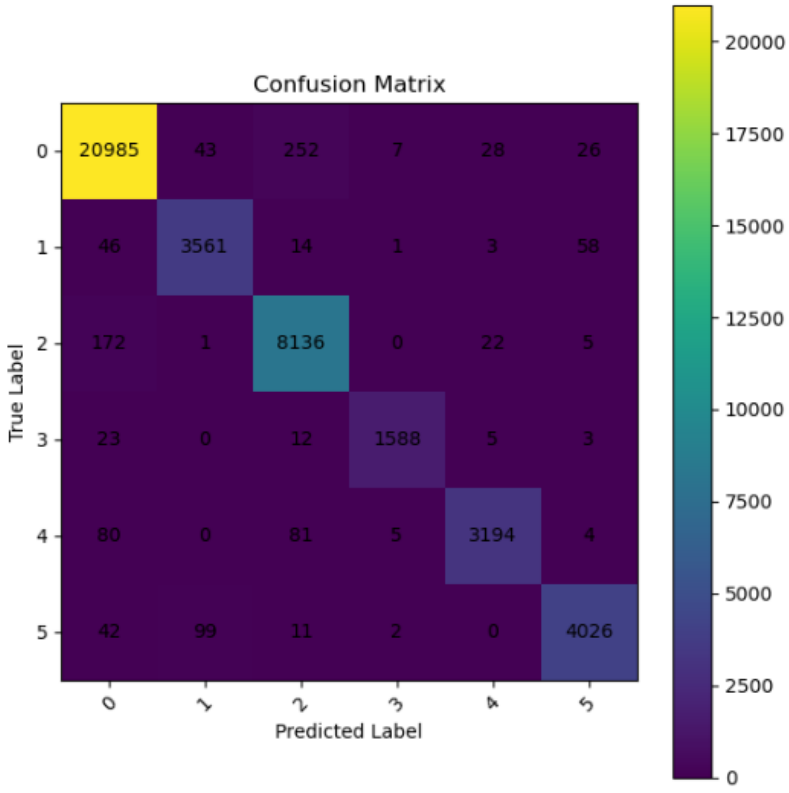

Figure 4.2 – Accuracy and loss during training epochs.

4.2.2 Results

To demonstrate the practical effectiveness of random forest, the ultimate result is shown in Figure 4.3. In this figure, every pixel is categorized by the material class it predicts.

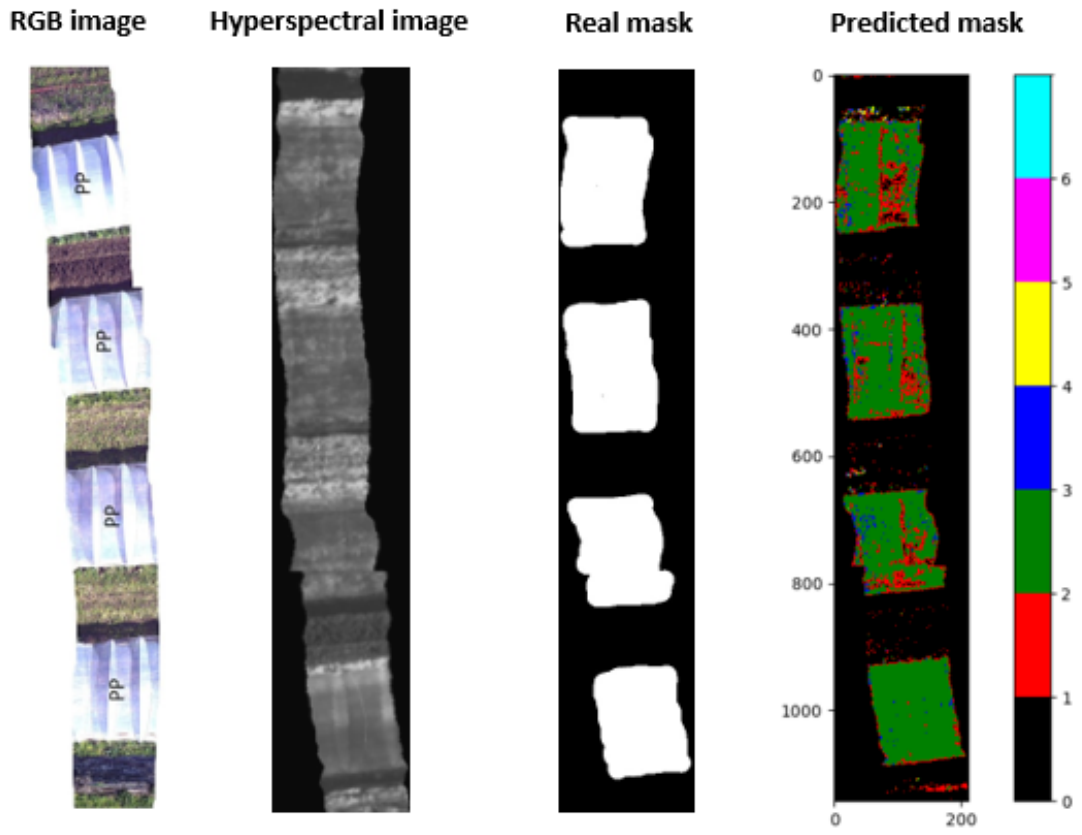

Figure 4.3 – Final Random forest output

4.3 Unet Approach

4.3.1 Training Performance and Visualization

The U-Net2D variant was subjected to an extensive training regimen to probe its convergence behavior and segmentation stability under prolonged optimization. When trained for 200 epochs, the model achieved a peak training accuracy of 99.03%, with a corresponding validation accuracy of 97.24%. Encouraged by this strong baseline, we further extended training to 250 epochs, which yielded marginal yet consistent gains : training accuracy rose to 99.08%, while validation accuracy improved to 97.59%. These results suggest that, although the bulk of feature learning and boundary refinement occurs within the first 200 epochs, an additional 50 epochs can still fine-tune the network's weights to better generalize on unseen samples.

A detailed view of this progression is presented in Figure 4.4, which charts the evolution of both training and validation accuracies alongside their respective losses over the full 250-epoch span. Evaluation metrics summarized in Table 4.3 corroborate these trends, showing diminishing returns beyond 200 epochs—validation loss plateaus even as accuracy ticks upward only slightly. Timing benchmarks, also included in Table 4.3, confirm that the extended 250-epoch run remains compatible with near real-time inference requirements, striking a practical balance between marginal performance gains and computational overhead.

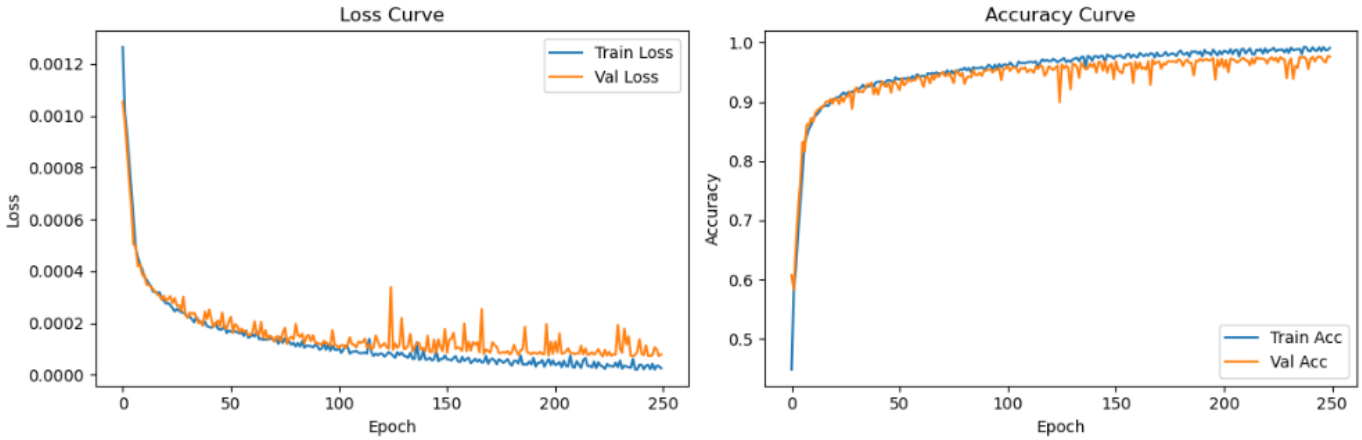

Figure 4.4 – Accuracy and loss during training epochs.

Metric	Achieved Value
IoU (Intersection over Union)	95.48%
Overall Accuracy	97.59%
F1-score	97.68%
mAP	95.84%

Table 4.3 – Segmentation Performance Metrics and Achieved Results

The confusion matrix in Figure 4.5 reveals an overwhelmingly strong diagonal—each class is correctly identified in well over 95% of cases, underscoring the U-Net2D's ability to distinguish among the six material categories. Altogether, this matrix confirms that the network has learned robust, discriminative features for each polymer type, with most errors clustered among those whose hyperspectral signatures inherently overlap (especially in the class 0 where some background patches contains some plastic debries).

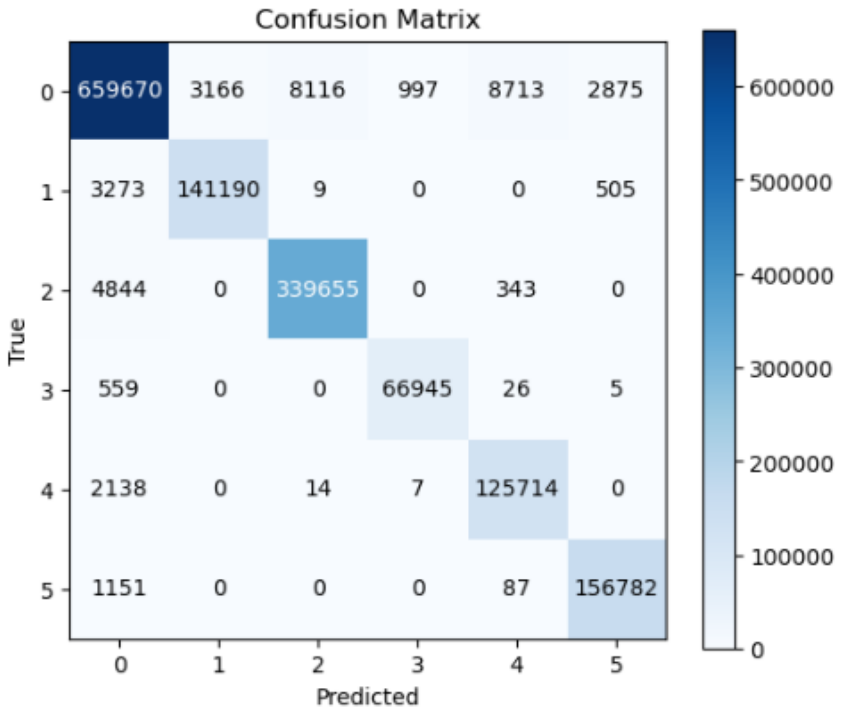

Figure 4.5 – Confusion matrix.

4.3.2 Results

To illustrate the U-Net2D's practical performance, the final output is shown in Figure 4.6. Here, each pixel has been labeled according to its predicted material class,

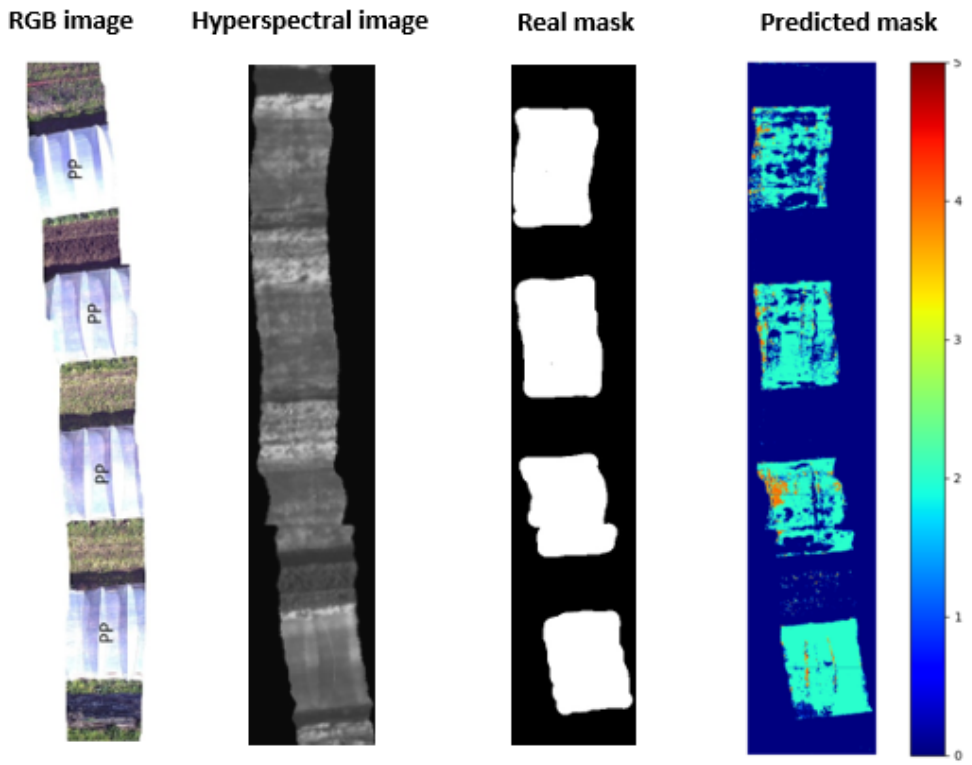

Figure 4.6 – Final U-net2D output.

4.4 Two Phase : Unet+SVM Approach

4.4.1 4.4.1 Segmentation Results Before Classification

4.4.2 Training Performance and Visualization

To evaluate the segmentation capability of the proposed U-Net architecture, the model was trained under different epoch settings to analyze performance convergence and stability. Initially, the model was trained for 5 epochs, achieving a segmentation accuracy of approximately 99.28%. Although this initial result was promising, further experiments were conducted to assess whether longer training would lead to additional performance gains.

After extending the training to 10 epochs, the model reached an improved accuracy of 99.35%, accompanied by a notable reduction in validation loss. These results indicate better convergence, suggesting that the model effectively captured the underlying spatial features within the dataset. However, no significant improvements in accuracy or loss reduction were observed beyond 10 epochs, confirming that 10 epochs provided a good balance between performance and computational efficiency.

The training progression is illustrated in Figure 4.7, which shows trends in accuracy, loss, and evaluation metrics (Table 4.4) across epochs. Timing information is also included to highlight the framework's suitability for near real-time environmental monitoring.

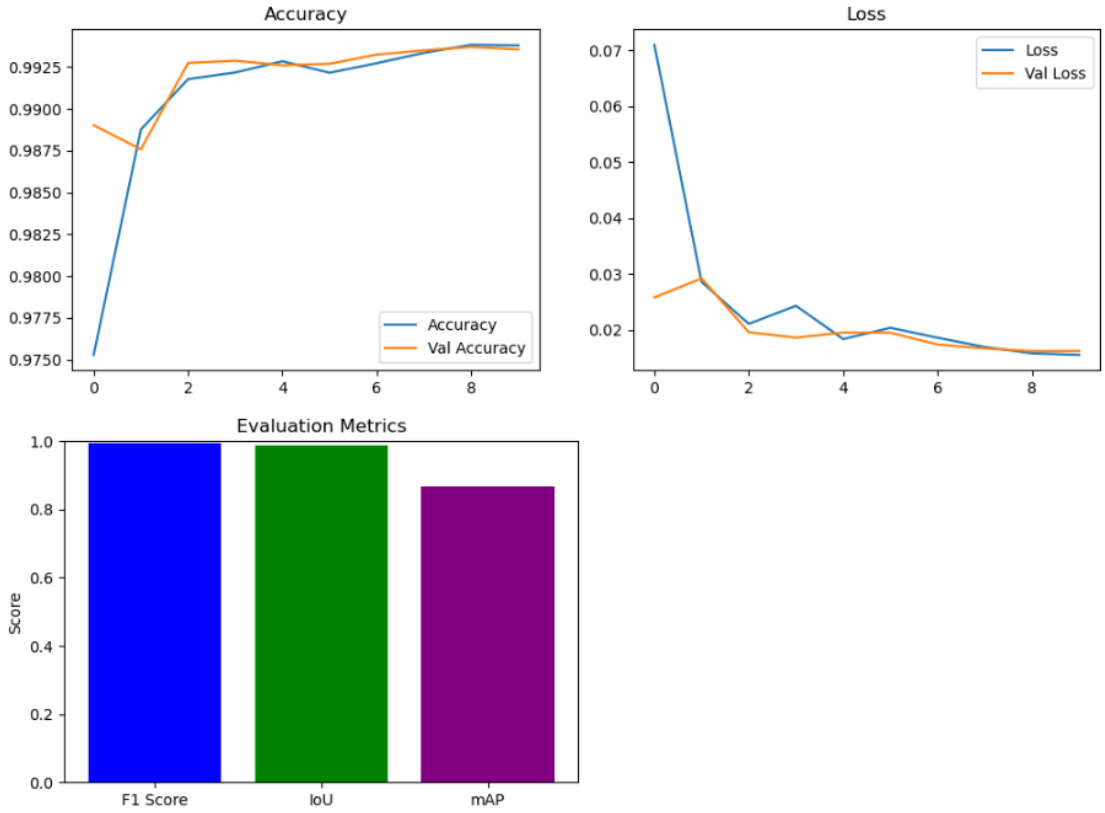

Figure 4.7 – Accuracy, loss, and evaluation metrics during training epochs.

Metric	Achieved Value
IoU (Intersection over Union)	98.75%
Overall Accuracy	99.35%
F1-score	99.42%
mAP	87.91%

Table 4.4 – Segmentation Performance Metrics and Achieved Results

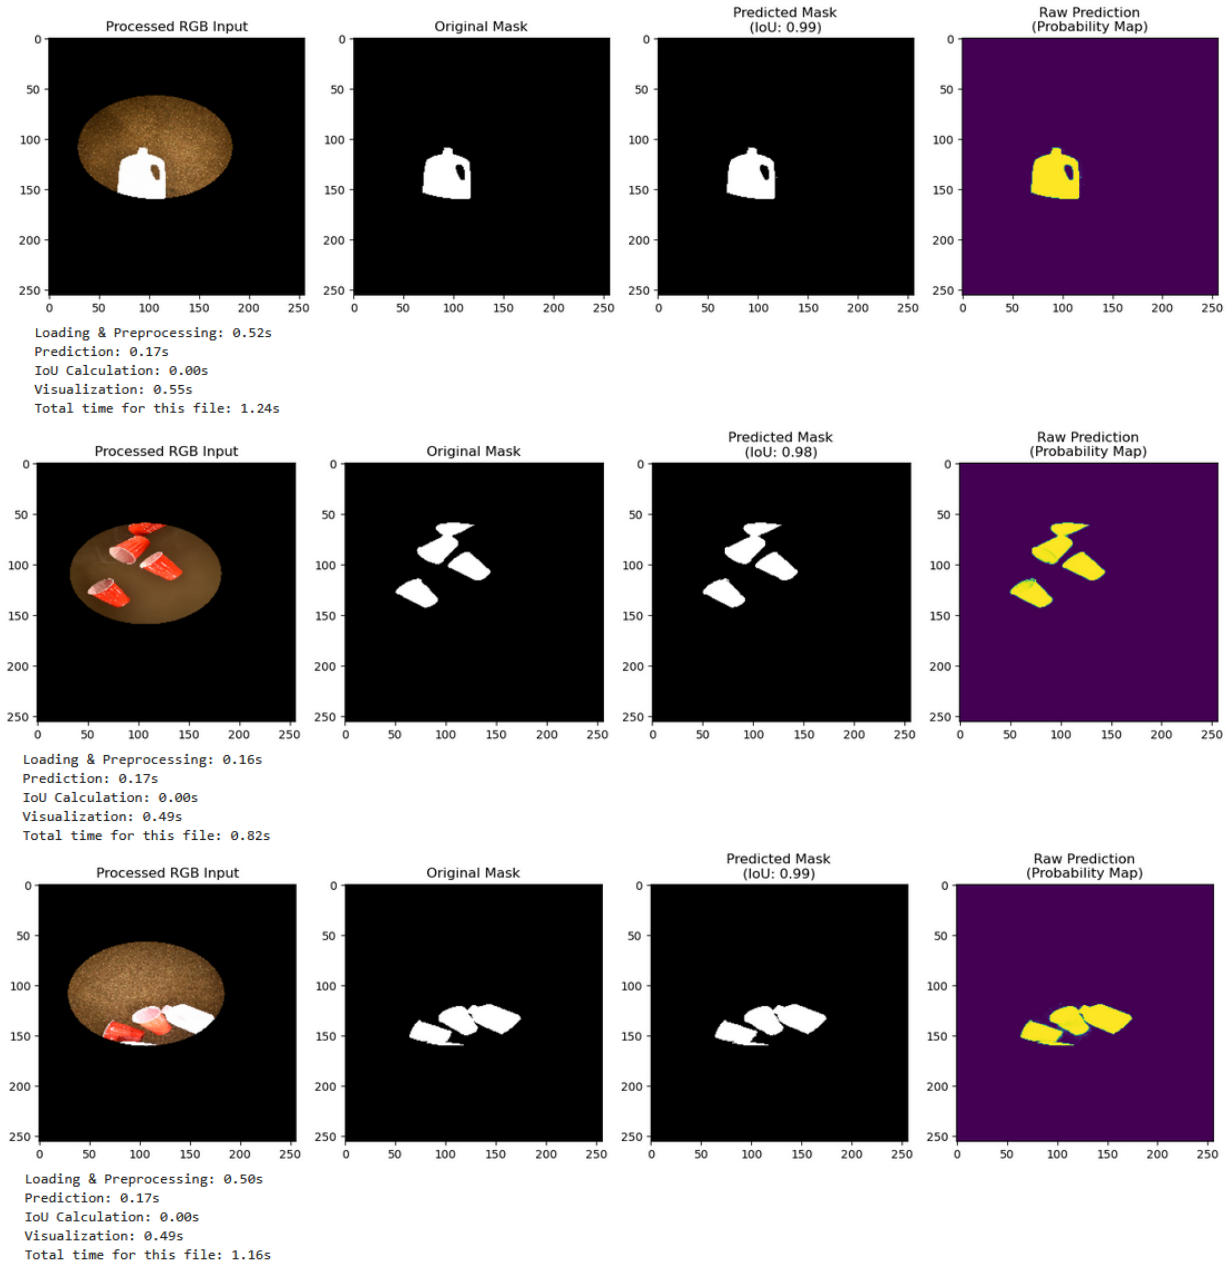

Figure 4.8 – Segmentation Results : Another Example of RGB Input, Ground Truth Mask, Predicted Mask (IoU Score), and Probability Map.

4.4.3 4.4.2 Classification Results

The confusion matrix, presented in Figure 4.9, demonstrates strong overall classification performance, with high accuracy across most classes. Notably, the majority of misclassifications occur between EPSF and PP classes, which can be attributed to their similar hyperspectral signatures. Furthermore, the probability calibration plot corroborates that the SVM model yields well-calibrated confidence scores, indicating reliable predictive uncertainty estimates for each class prediction.

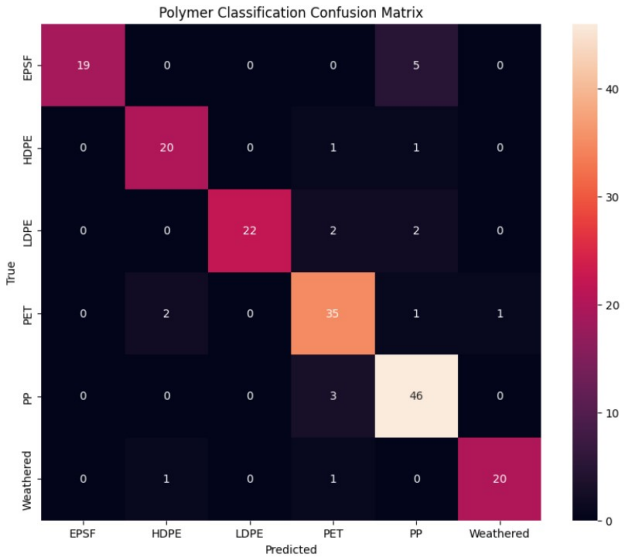

Figure 4.9 – Polymer Classification Confusion Matrix.

The classifier achieved an overall precision of 91.6% over 307 validation samples, with detailed performance metrics for each polymer class presented in Table 4.5.

Polymer Type	Precision	Recall	F1-score
EPSF	1.00	0.79	0.88
HDPE	0.87	0.91	0.89
LDPE	1.00	0.85	0.92
PET	0.83	0.90	0.86
PP	0.84	0.94	0.88
Weathered	0.95	0.91	0.93
Mean	**0.915**	**0.883**	**0.893**

Table 4.5 – Classification Performance by Polymer Type

The use of multiple evaluation metrics namely precision, recall, and F1-score—demonstrates the robustness of our classification approach beyond the scope of simple accuracy. Each of these metrics offers a distinct perspective on the classifier's effectiveness. Specifically, *precision* quantifies the proportion of correctly identified positive instances among all instances predicted as positive (i.e., the ratio of true positives to the sum of true and false positives). *Recall*, on the other hand, measures the proportion of true positives correctly identified out of all actual positive instances (i.e., the ratio of true positives to the sum of true positives and false negatives), reflecting the model's sensitivity.

The *F1-score*, defined as the harmonic mean of precision and recall, offers a balanced assessment by penalizing extreme values in either metric, making it particularly useful in contexts with class imbalance. Unlike overall accuracy, which includes both true positives and true negatives and may be misleading in datasets with skewed class distributions, these metrics focus specifically on the model's ability to detect relevant plastic types.

4.4.4 4.4.3 Final Classification and Segmentation Results

To visualize the effectiveness of the proposed pipeline, final classification outputs are illustrated for each polymer class following semantic segmentation and SVM-based classification. Figures 4.10 through 4.15 present representative examples of correctly identified plastic types, including EPSF, HDPE, LDPE, PET, PP, and weathered samples. Each visualization includes the model's confidence score for the predicted class, highlighting the reliability of the classification output.

The high-confidence predictions depicted in these figures demonstrate the system's ability to not only segment plastic

materials accurately but also to assign them to the correct polymer category, even under varying environmental conditions. Notably, consistent confidence scores across distinct polymer types affirm the robustness and generalization capacity of the integrated U-Net and SVM pipeline.

Note : In each figure from figure4.10 to figure 4.15, the first image is shown in its original aspect ratio, while the second (segmentation output) has been resized for visual comparison—this may cause slight differences in aspect ratio between the two images. Additionally, a third histogram has been added to represent the classification rates achieved by the combined two-phase approach, complementing the individual results from Phase 1 and Phase 2.

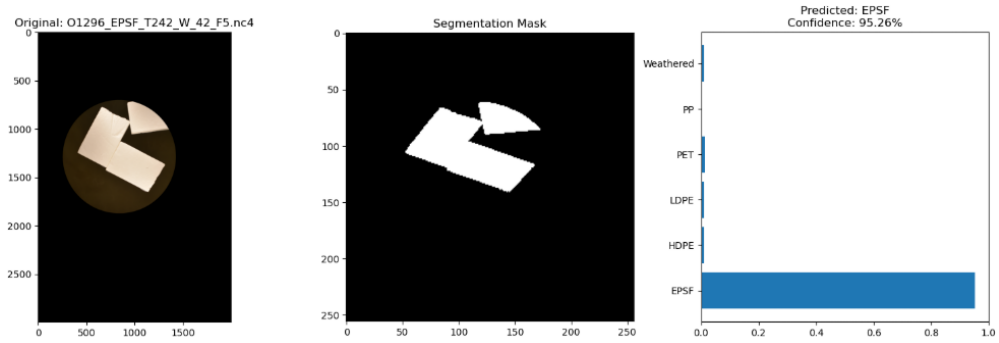

Figure 4.10 – Final Classification Output : Expanded Polystyrene Foam (EPSF) Polymer with Confidence Score.

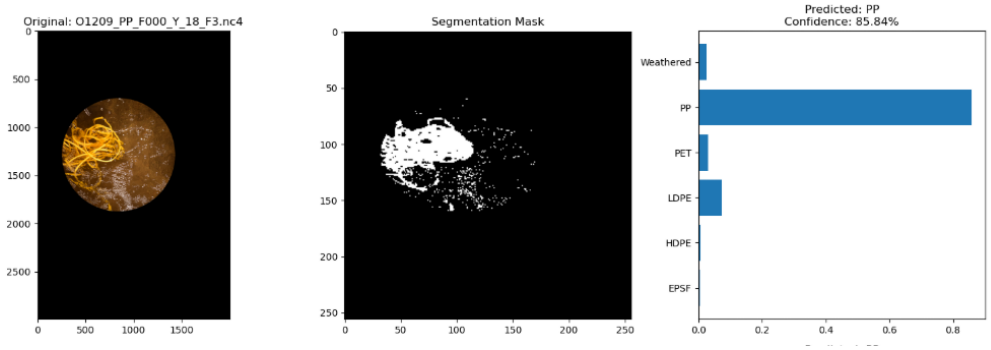

Figure 4.11 – Final Classification Output : Polypropylene (PP) Polymer with Confidence Score.

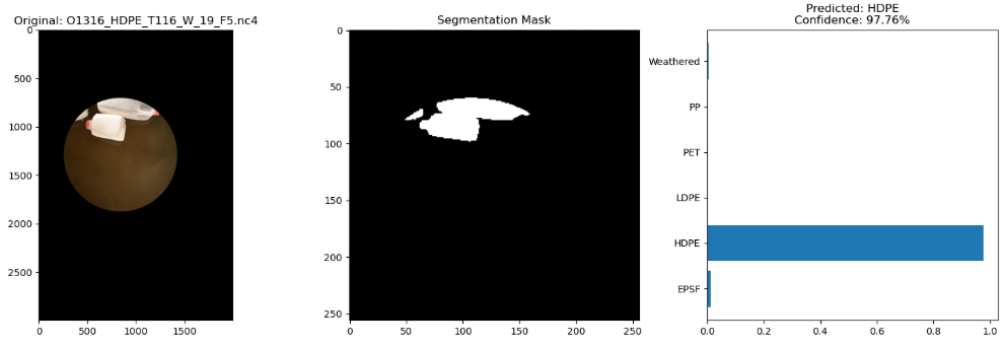

Figure 4.12 – Final Classification Output : High-Density Polyethylene (HDPE) Polymer with Confidence Score.

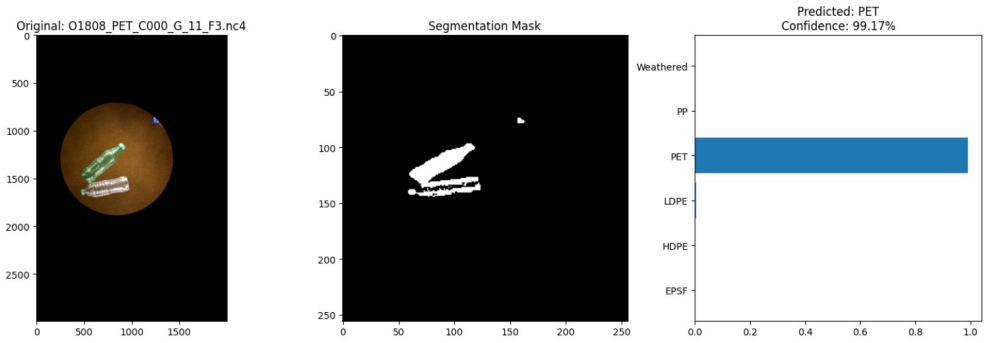

Figure 4.13 – Final Classification Output : Polyethylene Terephthalate (PET) Polymer with Confidence Score.

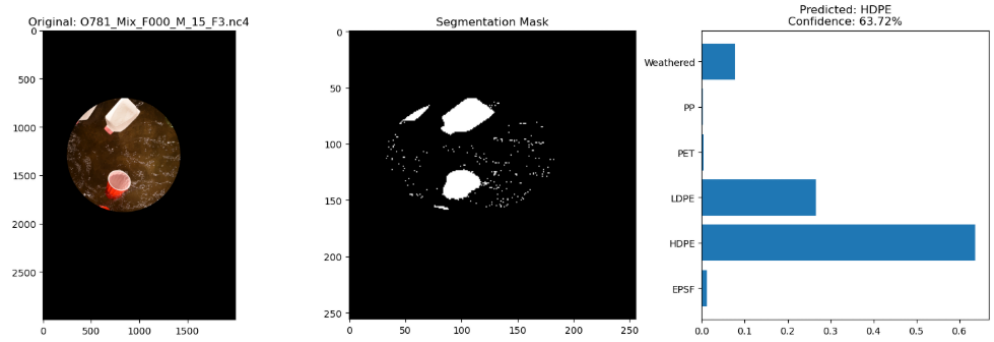

Figure 4.14 – Final Classification Output : Mixed Polymers with Confidence Score.

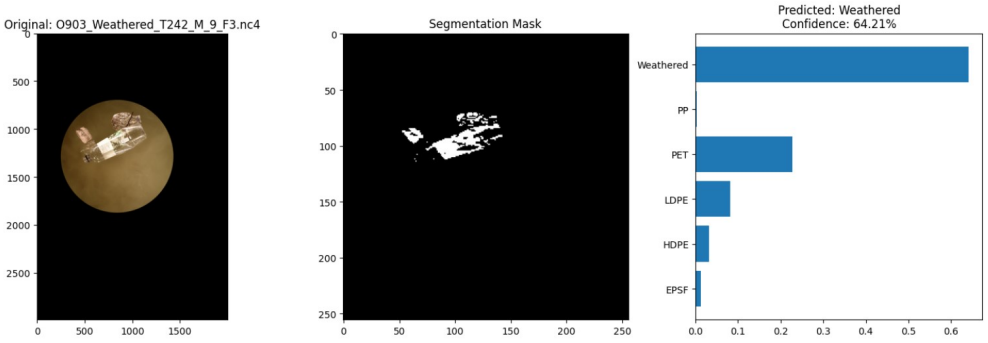

Figure 4.15 – Final Classification Output : Weathered Polymer Sample with Confidence Score.

4.5 Discussion

All four methods—SVM, Random Forest, U-Net, and the two-phase U-Net+SVM pipeline offer distinct strengths for hyperspectral plastic analysis. The SVM classifier on PCA-reduced features is lightweight and achieves a high F1-score (0.98), but its pixel-wise decisions lack spatial context, leading to occasional misclassifications among spectrally similar polymers. Random Forest reached 98% accuracy and is robust against overfitting, yet its outputs can be noisy without explicit spatial smoothing.

By contrast, the U-Net architecture excels at integrating spectral and spatial information. It delivers superior segmentation quality (IoU 95.5%, F1 97.7%) and inherently smooth predictions that respect object boundaries. Its ability to learn complex feature hierarchies makes it particularly adept at distinguishing plastics in heterogeneous backgrounds, and even modest additional training yields incremental gains without overfitting.

Building on this, the two-phase U-Net+SVM approach further elevates performance by first applying U-Net for precise localization (IoU 98.8%, F1 99.4%) and then using SVM to assign polymer labels with high confidence. This cascade combines the best of spatial segmentation and spectral classification, yielding clearer delineation and more accurate polymer identification than either method alone. Although its overall accuracy (88%) reflects some error propagation between stages, careful tuning of segmentation thresholds and classifier hyperparameters can minimize such effects.

In practical deployments, we recommend U-Net for scenarios where segmentation fidelity and spatial coherence are paramount, and the two-phase pipeline when both precise localization and reliable polymer classification are required, recognizing that the added complexity is justified by significantly improved end-to-end performance.

5 Conclusion

This chapter has highlighted the diverse capabilities of both classical and deep learning approaches in tackling the challenges of plastic detection using hyperspectral imaging. While traditional classifiers such as SVM and Random Forest offer efficiency and simplicity, their reliance on spectral data alone limits their spatial interpretability. Deep learning methods, particularly U-Net, introduce a paradigm shift by harnessing both spectral and spatial dimensions for more coherent segmentation.

Moreover, the integration of segmentation and classification into a unified two-phase framework demonstrates that combining complementary models can significantly enhance system robustness and interpretability. This modular design not only improves accuracy but also offers flexibility for adapting to complex, real-world environments.

The outcomes of this comparative analysis lay the foundation for selecting and deploying the most suitable architecture based on application-specific requirements.

General Conclusion

In this thesis, we have addressed one of the most urgent environmental challenges of our era : the identification and classification of plastic waste in complex natural environments. Leveraging the power of hyperspectral imaging (HSI) combined with advanced artificial intelligence (AI) techniques, we proposed a multi-approach framework that enhances both the accuracy and efficiency of plastic detection systems in real-world scenarios.

Our work was structured around four complementary computational methodologies : (1) A dimensionality-reduced spectral classifier using Principal Component Analysis (PCA) followed by Support Vector Machines (SVM) ; (2) A Random Forest model for robust spectral discrimination ; (3) A 2D U-Net deep convolutional network capable of direct spectral-spatial segmentation ; and (4) A hybrid two-phase pipeline combining deep learning-based RGB segmentation and hyperspectral SVM classification for precise polymer-level discrimination.

Through rigorous experimentation on heterogeneous hyperspectral datasets—ranging from aerial UAV acquisitions to controlled laboratory scenes—we validated the effectiveness of each method. Notably, the two-phase hybrid approach achieved state-of-the-art results with a segmentation performance of 99.42% F1-score, 99.35% accuracy, and 98.75% Intersection over Union (IoU). The classification phase reached an overall precision of 91.5% across six polymer types, achieving 100% precision for EPSF and LDPE.

From a computational perspective, the full framework demonstrated real-time capabilities, processing 10 hyperspectral files in just 12.7 seconds on GPU hardware. This reinforces the system's potential for deployment in embedded platforms, including drones and IoT-based environmental sensors.

This thesis contributes both methodologically and practically to the fields of computer vision, remote sensing, and environmental informatics. It bridges deep learning and classical machine learning within a hyperspectral context, tackling fundamental challenges such as spectral redundancy, spatial heterogeneity, and class imbalance. Moreover, the integration of spatial and spectral cues underlines the importance of multi-modal reasoning in environmental AI systems.

Advanced Perspectives : Building on the robust methodological foundation of this work, several advanced research and development trajectories emerge, aiming to push the boundaries of real-time plastic waste detection and environmental AI systems :

- **Self-Supervised and Few-Shot Learning for Spectral Data :** Future systems can leverage self-supervised learning paradigms to pre-train models on unlabeled hyperspectral data, drastically reducing the dependence on annotated samples. In parallel, few-shot and meta-learning approaches could enable rapid adaptation to new plastic types or environmental conditions with minimal data.
- **Synthetic Data Generation and Virtual Environments :** Employ physics-based rendering engines and generative models (e.g., GANs or NeRFs) to simulate hyperspectral plastic waste scenes under various lighting and environmental conditions. This can significantly augment training datasets and support robust generalization across unseen scenarios.
- **Federated and Privacy-Preserving Environmental AI :** To support decentralized and collaborative environmental monitoring, federated learning frameworks can be developed, allowing models to be trained across heterogeneous datasets from different regions or stakeholders—without sharing sensitive or proprietary data.
- **Cross-Modality Fusion and 3D Scene Understanding :** Extend the current 2D spectral-spatial analysis to integrate other sensing modalities such as LiDAR, thermal imaging, or 3D photogrammetry. This would enable full-scene understanding, enhancing plastic detection in cluttered or occluded environments like forests, coastlines, or landfills.
- **Neurosymbolic AI for Plastic Waste Reasoning :** Integrate symbolic reasoning systems with deep learning architectures to move from detection to contextual understanding (e.g., identifying illegal dumping patterns, source attribution, or predicting environmental impact), thereby empowering AI systems with explainable logic-based decision-making.
- **Dynamic Onboard Optimization and Continual Learning :** Design adaptive models that can dynamically reconfigure based on onboard resource constraints (GPU/TPU usage, power) or update continuously using online data streams. This enables persistent plastic surveillance on autonomous platforms (e.g., long-range drones or underwater robots).
- **AI-Driven Policy Support and Decision-Making Tools :** Develop user-friendly interfaces and policy dashboards powered by AI to assist environmental agencies and NGOs in resource allocation, pollution forecasting, and remediation prioritization—based on real-time plastic detection and accumulation analytics.

— **Contributions to Open Standards and Environmental Digital Twins :** Align with emerging international efforts (e.g., UNEP, ESA, ISO) to contribute to open standards for spectral plastic detection. This includes integrating the proposed system into larger environmental digital twin infrastructures for simulating and managing ecosystems at scale.

In conclusion, this work lays the groundwork for an intelligent, adaptive, and real-time plastic detection ecosystem, marking a significant step toward sustainable environmental monitoring using cutting-edge AI and hyperspectral technologies. The fusion of computational efficiency, spatial-spectral reasoning, and real-world applicability establishes a solid foundation upon which future research, policy development, and industrial deployment can be built—ultimately contributing to a cleaner, smarter, and more sustainable planet.

Bibliographie

[1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11) :2274–2282, 2012. Available at : https://ieeexplore.ieee.org/document/6205760.

[2] Muhammad Ahmad. *Deep Learning for Hyperspectral Image Classification*. Phd thesis, University of Messina, 2021. Doctoral Programme in Cyber Physical Systems, XXXIV Cycle. Available at https://iris.unime.it/retrieve/handle/11570/3212324/446618/Revised_Thesis.pdf.

[3] Ahed Alboody, Nicolas Vandenbroucke, Alice Porebski, Rosa Sawan, Florence Viudes, Perine Doyen, and Rachid Amara. A new remote hyperspectral imaging system embedded on an unmanned aquatic drone for the detection and identification of floating plastic litter using machine learning. *Remote Sensing*, 15(3455) :1–23, July 2023. https://doi.org/10.3390/rs15143455.

[4] María C. Alonso, José A. Malpica, and Alex Martínez de Agirre. Consequences of the hughes phenomenon on some classification techniques. In *Proceedings of the 2011 American Society for Photogrammetry and Remote Sensing (ASPRS) Annual Conference*, Milwaukee, Wisconsin, USA, May 1–5 2011. ASPRS. https://www.asprs.org/wp-content/uploads/2010/12/Alonso.pdf.

[5] Marco Balsi, Monica Moroni, and Soufyane Bouchelaghem. Plastic litter detection in the environment using hyperspectral aerial remote sensing and machine learning. *Remote Sensing*, 17(5) :938, 2025. https://www.mdpi.com/2072-4292/17/5/938.

[6] Stuart J. Barnes. Understanding plastics pollution : The role of economic development and technological research. *Environmental Pollution*, 249 :812–821, 2019. https://doi.org/10.1016/j.envpol.2019.03.108.

[7] José M. Bioucas-Dias, Antonio Plaza, Nicolas Dobigeon, Mario Parente, Qian Du, Paul Gader, and Jocelyn Chanussot. Hyperspectral unmixing overview : Geometrical, statistical, and sparse regression-based approaches. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 5(2) :354–379, 2012. https://ieeexplore.ieee.org/document/6200362.

[8] Martín C. M. Blettler, Elie Abrial, Farhan R. Khan, Nuket Sivri, and Luis A. Espinola. Freshwater plastic pollution : Recognizing research biases and identifying knowledge gaps. *Water Research*, 143 :416–424, October 2018. https://doi.org/10.1016/j.watres.2018.06.015.

[9] Giuseppe Bonifazi, Eleuterio Francesconi, Riccardo Gasbarrone, Roberta Palmieri, and Silvia Serranti. A preliminary study on the utilization of hyperspectral imaging for the on-soil recognition of plastic waste resulting from agricultural activities. *Land*, 12(10) :1934, 2023. https://www.mdpi.com/2073-445X/12/10/1934.

[10] Leo Breiman. Random forests. *Machine Learning*, 45(1) :5–32, 2001. https://link.springer.com/article/10.1023/A:1010933404324.

[11] Bruno and Edward. Automated sorting of plastics for recycling. Technical report, North Carolina Department of Environment and Natural Resources, 2000. Available at : https://p2infohouse.org/ref/09/08620.pdf or https://www.semanticscholar.org/paper/Automated-Sorting-of-Plastics-for-Recycling-Edward-Bruno/e6e5110c06f67171409bab3b38f742db6dc110fc.

[12] Zhen Cao, Lammert Kooistra, Wensheng Wang, Leifeng Guo, and João Valente. Real-time object detection based on uav remote sensing : A systematic literature review. *Drones*, 7(10) :620, 2023. https://www.mdpi.com/2504-446X/7/10/620.

[13] Turgay Celik. Change detection in satellite images using a genetic algorithm approach. *IEEE Geoscience and*

Remote Sensing Letters, 7(2) :386–390, 2010. Available at : https://ieeexplore.ieee.org/abstract/document/5395684.

[14] Bowen Chen, Liqin Liu, Zhengxia Zou, and Zhenwei Shi. Target detection in hyperspectral remote sensing image : Current status and challenges. *Remote Sensing*, 15(13), 2023. https://www.mdpi.com/2072-4292/15/13/3223.

[15] Nicholas C. Coops, Richard H. Waring, Michael A. Wulder, and Joanne C. White. Prediction and assessment of bark beetle-induced mortality of lodgepole pine using estimates of stand vigor derived from remotely sensed data. *Remote Sensing of Environment*, 113(5) :1058–1066, 2009. https://doi.org/10.1016/j.rse.2009.01.013.

[16] Sara Freitas, Hugo Silva, and Eduardo Silva. Hyperspectral imaging zero-shot learning for remote marine litter detection and classification. *Remote Sensing*, 14(21) :5516, 2022. https://www.mdpi.com/2072-4292/14/21/5516.

[17] Roland Geyer, Jenna R. Jambeck, and Kara Lavender Law. Production, use, and fate of all plastics ever made. *Science Advances*, 3(7) :e1700782, 2017. Additional data available at : https://www.science.org/doi/10.1126/sciadv.1700782 and https://ourworldindata.org/grapher/global-plastic-production-projections?time=earliest..2060.

[18] Mahdiyeh Ghaffari, Mickey C. J. Lukkien, Nematollah Omidikia, Gerjen H. Tinnevelt, Marcel C. P. van Eijk, Stanislav Podchezertsev, and Jeroen J. Jansen. Systematic reduction of hyperspectral images for high-throughput plastic characterization. *Scientific Reports*, 13 :21591, 2023. https://www.nature.com/articles/s41598-023-49051-y.

[19] Els Knaeps and Mark Bollen. Operational remote sensing mapping of estuarine suspended sediment concentrations (ormes). http://academia.edu/28171540/Operational_Remote_sensing_Mapping_of_Estuarine_suspended_Sediment_concentrations_ORMES_, Jan, 2006. 10 pages, 34 views.

[20] Winnie W. Y. Lau, Yonathan Shiran, Richard M. Bailey, Ed Cook, Martin R. Stuchtey, Julia Koskella, Costas A. Velis, Linda Godfrey, Julien Boucher, Margaret B. Murphy, Richard C. Thompson, Emilia Jankowska, Arturo Castillo Castillo, Toby D. Pilditch, Ben Dixon, Laura Koerselman, Edward Kosior, Enzo Favoino, Jutta Gutberlet, Sarah Baulch, Meera E. Atreya, David Fischer, Kevin K. He, Milan M. Petit, U. Rashid Sumaila, Emily Neil, Mark V. Bernhofen, Keith Lawrence, and James E. Palardy. Evaluating scenarios toward zero plastic pollution. *Science*, 369(6510) :1455–1461, Jul 2020. https://www.science.org/doi/10.1126/science.aba9475.

[21] Zuchuan Li, Liwei Li, Rui Zhang, and Jianwen Ma and. An improved classification method for hyperspectral data based on spectral and morphological information. *International Journal of Remote Sensing*, 32(10) :2919–2929, 2011. https://www.tandfonline.com/doi/abs/10.1080/01431161.2010.510488.

[22] John S MacDonald, Susan L Ustin, and Michael E Schaepman. The contributions of dr. alexander f.h. goetz to imaging spectrometry. *Remote Sensing of Environment*, 113 :S2–S4, 2009. https://doi.org/10.1016/j.rse.2008.10.017.

[23] Nisha Maharjan, Hiroyuki Miyazaki, Bipun Man Pati, Matthew N. Dailey, Sangam Shrestha, and Tai Nakamura. Detection of river plastic using uav sensor data and deep learning. *Remote Sensing*, 14(13), 2022. https://www.mdpi.com/2072-4292/14/13/3049.

[24] Uriel Martinez-Hernandez, Gregory West, and Tareq Assaf. Low-cost recognition of plastic waste using deep learning and a multi-spectral near-infrared sensor. *Sensors*, 24(9), 2024. https://www.mdpi.com/1424-8220/24/9/2821.

[25] Mohammadali Olyaei, Ardeshir Ebtehaj, and Christopher R. Ellis. A hyperspectral reflectance database of plastic debris with different fractional abundance in river systems. *Scientific Data*, 11 :1253, November 2024. https://www.nature.com/articles/s41597-024-03974-x.

[26] Caroline Orset, Nicolas Barret, and Aurélien Lemaire. How consumers of plastic water bottles are responding to environmental policies. *Waste Management*, 61 :13–27, March 2017. Available at : https://doi.org/10.1016/j.wasman.2016.12.034.

[27] Roberta Palmieri, Riccardo Gasbarrone, Giuseppe Bonifazi, Giorgia Piccinini, and Silvia Serranti. Hyperspectral imaging for detecting plastic debris on shoreline sands to support recycling. *Applied Sciences*, 14(23) :11437, December 2024. https://doi.org/10.3390/app142311437.

[28] Young-Je Park, Shungudzemwoyo P. Garaba, and Bruno Sainte-Rose. Detecting the great pacific garbage patch floating plastic litter using worldview-3 satellite imagery. *Optics Express*, 29(22) :35288–35298, 2021. https://doi.org/10.1364/OE.440380.

[29] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net : Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, volume 9351 of *Lecture Notes in Computer Science*, pages 234–241. Springer, October 2015. https://link.springer.com/chapter/10.1007/978-3-319-24574-4_28.

[30] Attilio Sbrana, Aline Gabriel de Almeida, André Mobaier de Oliveira, Henrique Seschin Neto, João Paulo Cesar Rimes, and Maria Cristina Belli. Plastic classification with nir hyperspectral images and deep learning. *IEEE Sensors Letters*, 7(1) :6000404, 2023. https://ieeexplore.ieee.org/document/10007047.

[31] Peg Shippert. Why use hyperspectral imagery ? *Photogrammetric Engineering & Remote Sensing*, 70(4) :377–380, 2004. https://docs.wixstatic.com/ugd/5aa63c_aecfc7d277bd49a184eade573d74d5bd.pdf.

[32] Chutimet Srinilta and Sivakorn Kanharattanachai. Municipal solid waste segregation with cnn. In *2019 5th International Conference on Engineering, Applied Sciences and Technology (ICEAST)*, pages 1–5, Luang Prabang, Laos, 2019. IEEE. https://ieeexplore.ieee.org/document/8802522.

[33] Allan T. Williams and Nelson Rangel-Buitrago. The past, present, and future of plastic pollution. *Marine Pollution Bulletin*, 176 :113429, 2022. Available at : https://doi.org/10.1016/j.marpolbul.2022.113429.

[34] Xiangrong Zhang, Qiang Song, Ruochen Liu, Wenna Wang, and Licheng Jiao. Modified co-training with spectral and spatial views for semisupervised hyperspectral image classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 7(6) :2044–2055, June 2014. https://www.researchgate.net/publication/264564193_Modified_Co-Training_With_Spectral_and_Spatial_Views_for_Semisupervised_Hyperspectral_Image_Classification.

Webography

— [w1] https://www.oecd.org/en/about/news/press-releases/2022/06/global-plastic-waste-set-to-almost-triple-by-2060.html,
Last checked : 28/01/2025

— [w2] https://ourworldindata.org/explorers/plastic-pollution?facet=none&country=USA~CHN~IND~GBR~MYS~DEU&hideControls=false&Metric=Plastic+emitted+to+ocean&Per+capita=true&Share+of+world+total=false&Source=Meijer+et+al.+%282021%29,
Full citation : Geyer et al. (2017) ; OECD (2022) – processed by Our World in Data.
Last checked : 28/01/2025

— [w3] https://oceanliteracy.unesco.org/plastic-pollution-ocean/,
Last checked : 28/01/2025

— [w4] https://enlightngo.org/language/en/post/17944,
Last checked : 03/03/2025

— [w5] https://data.mendeley.com/datasets/y8cvcs8tt5/1,
Last checked : 26/02/2025

— [w6] https://www.researchgate.net/publication/372163006/figure/fig2/AS:11431281172897338@1688699539406/Intersection-Over-Union-IoU-a-The-IoU-is-calculated-by-dividing-the-intersection-of.ppm,
Last checked : 16/05/2025

— [w7] https://www.ibm.com/think/topics/principal-component-analysis,
Last checked : 18/05/2025

— [w8] https://www.ibm.com/think/topics/support-vector-machine,
Last checked : 03/05/2025

— [w9] https://www.nvidia.com/en-us/glossary/pytorch/,
Last checked : 03/06/2025

— [w10] https://www.nvidia.com/en-us/glossary/numpy/,
Last checked : 03/06/2025

— [w11] https://www.h5py.org/#:~:text=The%20h5py%20package%20is%20a,they%20were%20real%20NumPy%20arrays.,
Last checked : 04/06/2025

— [w12] https://www.nvidia.com/en-us/glossary/pandas-python/,
Last checked : 03/06/2025

— [w13] https://www.nvidia.com/en-us/glossary/scikit-learn/,
Last checked : 03/06/2025

— [w14] https://www.analyticsvidhya.com/blog/2023/02/how-to-save-and-load-machine-learning-models-in-python-using-joblib-library/,
Last checked : 03/06/2025

— [w15] https://www.nvidia.com/en-us/glossary/tensorflow/,
Last checked : 03/06/2025

— [w16] https://opencv.org/about/,
Last checked : 03/06/2025

— [w17] https://datascientest.com/en/matplotlib-master-data-visualization-in-python,
Last checked : 03/06/2025

— [w18] https://www.unidata.ucar.edu/software/netcdf/?query_float=plastic,
Last checked : 03/06/2025

ANNEX

(Supplementary materials and additional information)

Annexe Start-up

5.1 Presenting the project

5.1.1 Project idea

Our web-based platform automates the segmentation and classification of plastic materials using hyperspectral imaging (HSI) data, providing a high-performance alternative to manual sorting methods. The system processes raw HSI files (.mat, .nc4) through integrated machine learning pipelines, deploying U-Net for pixel-level segmentation and SVM for polymer classification. It delivers interactive segmentation maps, quantitative reports (accuracy, F1-score, IoU), and material composition statistics—exportable as GeoTIFF/CSV files. By enabling adaptive learning with calibrated spectral libraries, the solution achieves 95% accuracy while drastically reducing processing time and scaling to terabyte-scale datasets. This advancement significantly enhances efficiency in recycling operations and environmental monitoring, supporting industrial sorting facilities and ecological research with auditable, metrics-driven outputs.

5.1.2 Proposed Value Proposition

— **High-Accuracy Analysis** : Unsupervised detection of over six plastic types (PET, HDPE, PVC, PP, PS, LDPE) with accuracy reaching 99% by combining deep autoencoders and spectral clustering.
— **Real-Time Processing** : Responsive web interface enabling on-the-fly processing upon hyperspectral data upload, with instant visualization of segmentation and classification results.
— **Multi-Data Support** : Full compatibility with `.mat` (MATLAB) and `.nc4` (NetCDF) , ensuring seamless integration into existing laboratory and recycling workflows.
— **Cost Reduction** : 30–50% reduction in manual sorting labor costs at recycling facilities through partial automation of pre-selection and verification steps.
— **Intuitive User Interface** : Graphical dashboard with drag-and-drop data import, and interactive visualization of spectral signatures and detected classes.
— **Robustness in Real-World Conditions** : Our algorithms demonstrate tolerance to synthetic variations in lighting and positioning through rigorous simulation-based validation. The model's architecture—designed to emulate real-world spectral distortions, has been systematically tested under controlled virtual environments replicating diverse conditions. While current validation is simulation-bound, results confirm the framework's potential for real-world adaptation.

These value propositions form the foundation of our solution : combining scientific precision, operational efficiency, and scalability to address the growing challenges of plastic waste sorting and valorization.

5.2 Project Team

The project team comprises the following members :
Yahiaoui Achraf : Expertise in artificial intelligence ; developed the hybrid Two-Phase U-Net + SVM architecture for hyperspectral data processing (.nc4 files).
Messioud Wassim : Expertise in artificial intelligence ; developed the U-Net segmentation model for hyperspectral image analysis (.mat files).
Mohammed Chaib : Expertise in web development ; implemented the website backend and model integration pipeline.
Abderrahmen Barkache : Expertise in web development ; designed the frontend interface and managed backend-frontend integration.
Hallaci Samir : Project supervisor and team lead ; expertise in AI and computer vision, his role was to oversee project execution, validate technical approaches, and maintain scientific rigor throughout development.

5.3 Project Objectives

Short-Term (1-2 years) :

- Deploy a Minimum Viable Product (MVP) web platform featuring the hybrid Two-Phase U-Net + SVM pipeline for processing hyperspectral data in `.nc4` format, enabling automated plastic segmentation and classification.
- Implement user authentication and data management systems for secure handling of industrial-scale hyperspectral datasets.
- Achieve sustained classification accuracy exceeding 95% across six major plastic polymer types (PET, HDPE, PVC, LDPE, PP, PS) through continuous model optimization.

- Expand file compatibility to include `.mat` hyperspectral data formats and integrate U-Net segmentation capabilities.

Mid-Term (3-5 years) :

- Conduct pilot testing with Algerian recycling facilities to validate platform performance under operational conditions.
- Work towards generalizing the hyperspectral market and include other fields (medical, agriculture, etc).
- Developing user training programs on hyperspectral image analysis and domain-specific implementations.

Long-Term (5-10 years) :

- Position the platform as a global reference in hyperspectral intelligence by offering end-to-end solutions across sectors such as agriculture, healthcare, environment, and manufacturing.
- Establish a large-scale annotated hyperspectral dataset repository spanning multiple domains, enabling benchmarking, collaborative research, and the development of generalized AI models for spectral analysis.
- Contribute to international research and standardization efforts, helping define best practices and protocols for spectral data acquisition, interpretation, and ethical deployment.

5.4 Project implementation schedule

Task	No.	Key Steps	Month 1	2	3	4	5
Project Research	1	Getting to know hyperspectral imaging and learning how to manipulate the data	✓	✓			
Model Development	2	Core model training (U-Net 2D, Two-phase U-Net + SVM)			✓	✓	
Web Deployment	3	Web platform development					✓

TABLE 5.1 – Project completion schedule

5.5 Innovation aspects

This project pioneers the first fully web-based HSI plastic classification service that :

— Hybrid Two-Phase Architecture : Combines deep convolutional segmentation with classical SVM classification for enhanced robustness on diverse data formats.

— Plug-and-Play Data Support : Users need no preprocessing expertise—raw cubes are automatically calibrated and normalized.

— Edge-to-Cloud Scalability : Supports both lab-scale and drone-acquired imagery, with elastic GPU scaling to accommodate high-volume uploads.

5.6 Strategic market analysis

5.6.1 Market Segments

— **Environmental Research Institutions & Universities** : Require automated tools to analyze large HSI datasets in plastic-pollution studies.

— **Recycling & Waste-Management Companies** : Seek precise polymer sorting to optimize material streams and improve quality control.

— **NGOs & Government Agencies** : Monitor coastal and inland pollution hotspots ; need rapid, repeatable analyses to inform policy.

— **Industrial Manufacturers** : Validate corporate waste-stream composition for sustainability reporting.

— **Individual Researchers & Citizen Scientists** : Access to advanced HSI tools without capital investment in hardware or software.

5.6.2 Measuring the intensity of competition

In Algeria, there are currently no domestically available hyperspectral imaging solutions. Local laboratories and recycling facilities still rely on time-consuming manual spectral analyses or basic multispectral instruments that cannot process full HSI cubes. No Algerian company offers cloud-hosted, deep-learning-based segmentation and classification services—most advanced platforms remain academic prototypes abroad or hardware-bundled systems with limited flexibility. This platform therefore occupies a unique position, delivering high-accuracy, user-friendly, scalable HSI analysis with rapid turnaround and minimal manual effort.

5.6.3 Stratégies marketing

- **Operational Efficiency :**
  - Our cloud-based platform, powered by AI and hyperspectral imaging, enables full automation of plastic detection and classification, reducing most manual sorting needs in recycling centers.
  - With a responsive web interface and real-time processing, complete spectral diagnostics are delivered in under 1 minute compared to several hours in conventional lab workflows.
- **Targeted Promotional Offers :**
  - The "Express Detection" program offers 10 free analyses per month for universities and research labs, encouraging early adoption in the scientific community.
  - A tiered subscription model (Basic, Pro, Enterprise) is available based on data volume, analysis frequency, and API access.
  - Discounts are provided to early pilot partners.
- **Specialized Content Marketing :**
  - We regularly publish case studies that demonstrate the platform's real-world impact (e.g., "40% reduction in sorting errors after one month in a recycling plant").
  - Monthly webinars are hosted with environmental and HSI experts to showcase data acquisition techniques and model interpretability.
  - Multilingual video guides are available to support drone-based plastic waste mapping and user onboarding.
- **Strategic Partnerships :**
  - We collaborate with hyperspectral sensor manufacturers to embed our API into embedded platforms, including industrial drones and fixed scanning stations.
  - Partnerships with environmental agencies, municipalities, and NGOs allow for the deployment of detection campaigns in sensitive coastal and urban zones.
  - Additionally, we are building educational alliances with universities and engineering schools to integrate our platform into training programs on environmental monitoring, AI, and remote sensing.

5.7 Production and Organization Plan

5.7.1 Production Process

The developed solution is entirely software-based and accessible through a web platform, allowing users to remotely upload their hyperspectral data files (*.mat*, *.nc4*) for automated analysis. This online platform approach offers flexible access and leverages the power of hyperspectral imaging to identify plastic polymers. The production process of our service consists of several key stages :

- **Software Development :** This involves designing and implementing the web platform and AI algorithms. This includes the development of the front-end interface allowing users to upload files and visualize results, as well as back-end services for data processing and user management. The core processing engine relies on hyperspectral imaging pipelines : a *U-Net* segmentation model isolates the materials in the image, coupled with an *SVM classifier* (Support Vector Machine) to identify the plastic type based on spectral data. The combination of deep learning and classical algorithms like SVM is a proven approach for plastic classification in hyperspectral images.
- **Testing and Validation :** Once initial modules are developed, rigorous internal testing is performed using labeled datasets, either proprietary or from academic collaborations. These tests help fine-tune model parameters (e.g., improving U-Net training on diverse plastic waste images) and ensure service reliability. After internal validation, **pilot field deployments** are scheduled in collaboration with local recycling centers. For instance, the company plans to install its detection solution in a recycling facility to assess real-world performance. This

involves acquiring hyperspectral images of waste on conveyor belts and verifying real-time segmentation and classification accuracy. This in-situ validation step is critical to confirm system effectiveness in industrial environments, particularly as recent studies highlight the potential of hyperspectral imaging for both factory-level sorting and remote plastic waste detection.

— **Support and Maintenance :** Once deployed, the company provides online technical support for users. This includes full documentation (user guides, FAQs, tutorials) and **onboarding** training resources to help new clients integrate the platform. A technical team monitors performance, proactively detects anomalies, and ensures ongoing maintenance. Regular software updates (new features, classifier improvements, security patches) are deployed with zero downtime thanks to the cloud infrastructure. A dedicated communication channel (chat or ticketing system) enables users to report issues or request help, ensuring high satisfaction. This proactive support and training system enhances user experience and long-term adoption.

— Although plastic detection is the initial application, the platform is designed to be modular and extensible, with the ambition of supporting other use cases of hyperspectral imaging in the future. These include applications in agriculture (e.g., crop health monitoring, soil analysis), food quality inspection, medical diagnostics, and environmental surveillance. By abstracting the data ingestion and analysis layers, the system is capable of integrating new models and spectral signatures specific to each domain. This long-term vision positions the product not just as a plastic detection tool, but as a general-purpose hyperspectral analysis platform.

5.8 Human Resources

While the current team consists of four core developers covering AI, front-end, and back-end components, future organizational growth will necessitate expanding the team to address broader objectives beyond plastic detection. The long-term vision is to build a multidisciplinary and scalable team capable of supporting diverse hyperspectral imaging applications across multiple domains. The expanded team will include :

— **Domain Experts :** Specialists in agronomy, environmental science, medical imaging, and materials science to help interpret spectral signatures and validate use-case-specific models.

— **Full-Stack Developers :** To scale the web platform's functionality, support data visualization for different sectors, and ensure integration with external hardware or cloud pipelines.

— **Data Engineers :** To manage the ingestion, preprocessing, and storage of massive volumes of hyperspectral data across different formats and acquisition systems.

— **Customer Success and Support Specialists :** To provide specialized onboarding and user assistance depending on the application sector.

— **R&D Partnership Manager :** To coordinate collaborations with research institutes and industrial stakeholders across fields.

This future team structure is designed to ensure that the platform remains robust, modular, and scientifically valid across a wide range of hyperspectral imaging applications. The combination of technical, scientific, and operational expertise will enable the company to become a leader in hyperspectral analytics as a service.

5.9 Supply Chain

As a software and cloud service company, the **supply chain** involves technological resources, data, and external services essential for platform production and continuous improvement. The key components are :

— **Input Hyperspectral Data :** The "fuel" of our AI engine is data. Training and testing models requires a large volume of representative hyperspectral images. Sources include public datasets, academic partners, and field-collected data. Constant access to high-quality data is essential for performance. This includes sharing agreements with universities, and manual annotation or spectroscopic verification to ensure reliable ground truth.

— **Cloud Infrastructure :** The platform relies on compute-intensive resources to process large hyperspectral images. Rather than owning hardware, the company uses cloud providers offering GPU-equipped instances. This IaaS model allows rapid scaling and provides essential services like load balancing, data backup, and CI/CD.

— **Future Hardware Integration :** Though currently SaaS-based, future end-to-end solutions may include hyperspectral imaging hardware. Partnerships with camera manufacturers (e.g., SPECIM, Headwall) will provide sensors for pilot projects. Hardware will integrate into our pipeline through standard protocols, enabling continuity from data capture to analysis. Logistics will be managed through leasing or conditional purchases depending on project success.

— **Support Services and Tools :** Daily operations depend on services like internet connectivity, project management tools, code repositories, security audits, and potentially environmental certification consultancy. Services are selected for quality, cost-efficiency, and reliability. CI/CD pipelines and secure data handling (encryption, backups) ensure compliance and continuous delivery.

5.10 Workforce Planning

To support scale-up and product enrichment, a **workforce development plan** is established to anticipate key hires and work organization over the next 12–24 months :

— **Immediate Hires :** Two technical roles are prioritized. First, a *Data Acquisition Specialist* will manage hyperspectral data collection, annotation, and storage. Second, a *Hyperspectral Imaging Expert* will enhance scientific depth and provide training to the internal team and users.
— **Organizational Evolution :** New hires will restructure the team. The data specialist offloads dataset management from AI developers. The HSI expert collaborates with all teams : refining models, adapting to sensor formats, and advising on spectral visualization. A *technical project manager* may be appointed internally to coordinate development sprints and inter-team communication.
— **Training and Skill Development :** Continuous learning will be encouraged through MOOCs (massive open online course), certifications, and internal knowledge-sharing sessions. This ensures agility and alignment with technological and market evolutions.
— **Mid-Term Recruitment :** Depending on growth, roles like *Cloud Architect*, or *Customer Support Specialist* may be added to handle scale, optimize infrastructure, and improve responsiveness. “‘

The workforce plan aims to proactively structure the team for future demands. Targeted hiring and ongoing upskilling will ensure that human capital remains an asset rather than a bottleneck.

5.11 Strategic Partnerships

Strategic partnerships are central to the startup's organizational plan. They accelerate technological development, provide critical resources, and validate the solution in the target market. Four main partnership axes are pursued :

— **Industrial Recycling Partners :** Collaborations with local plastic recycling companies provide real-world testing and potential early adopters. These partnerships ensure the platform meets operational needs and facilitate future commercial deployment.
— **Academic Partnerships :** Ties with universities and research centers provide scientific expertise, access to experimental data, and a talent pool. Collaborations include data exchange, joint research, and internships, keeping the solution state-of-the-art.
— **Institutional Partnerships :** Relationships with public agencies (e.g., the Ministry of Environment) align the project with national sustainability goals. These institutions can offer support, environmental data, and legitimacy, opening public-sector opportunities.
— **Technology Partners :** Collaborations with HSI sensor manufacturers ensure hardware compatibility and joint value propositions. Examples include co-marketing, case studies, and hardware-software demos, supporting future integrated solutions.

This partnership strategy creates a robust ecosystem around the startup. Industry partners ground the solution in operational reality, academia drives R&D, institutions boost legitimacy, and tech providers extend functionality. Each collaboration is win-win, reinforcing the startup's mission : collaborative innovation for environmental impact through hyperspectral plastic detection.

5.12 Financial Plan

5.12.1 Costs and Expenses

Identifying all essential costs and investments is crucial for operating in a competitive market and delivering reliable services to clients. The expenses are categorized into :

a) Initial Costs

Infrastructure :

— Office rental or co-working setup for administration and development.
— Installation of computing hardware, network access, and security systems.
— Cloud platform setup (GPU-enabled instances, secure storage).

Equipment :

— Hyperspectral camera rental or pilot device (optional, for demonstrations).
— Developer workstations and peripheral equipment.
— Development and integration tools (minimal license costs).

Software and AI Stack :

— In-house development of platform front-end, back-end APIs.
— U-Net segmentation model and SVM classifier implementation.
— Secure database (PostgreSQL/MongoDB) and data pipelines for hyperspectral processing.

Pilot Deployment (optional) :

— Temporary field site setup for pilot demonstration.
— Basic camera mounting system and transport logistics.

b) Monthly Operational Costs

Personnel :

— Salaries for developers (AI/Full-stack), cloud engineers, and support staff.
— Consulting agronomists, remote sensing specialists.
— Monthly training and certification (AI, cloud platforms).

Infrastructure :

— Cloud compute (GPU), storage and bandwidth.
— Licensing of API tools (map, weather, spectral DBs).
— Serverless functions for automated processing and inference.

Logistics and Office :

— Office rent, internet, maintenance.
— Subscriptions to collaborative tools (GitHub, Trello, Figma).

Marketing and Customer Support :

— Facebook/Instagram ads, SEO targeting ("plastic detection", "drone waste mapping").
— Content creation (educational videos, datasheets).
— Hotline, WhatsApp Business, email, webinar hosting.

c) Recurring Costs

— Periodic calibration or renewal of pilot devices (if applicable).
— Annual domain and platform license renewals.
— Cloud security updates, encrypted storage renewals.

d) Other Costs

— Liability and cyber-insurance.
— Compliance with GDPR-like data protection laws.
— Certification fees (environmental conformity, drone permissions).
— Security software (firewalls, CI/CD monitoring).

5.12.2 Funding Methods and Sources

a) Internal Funding :

— Founders contribute initial capital for incorporation and MVP development.
— Reinvesting early revenue from pilots or small-scale contracts.

b) External Funding :

Bank Loans :

— Startup credit to support product deployment and team expansion.
— Loan repayment aligned with revenue milestones.

Investors :

— Angel investors or early-stage venture capital (seed round).
— Strategic partnerships with environmental or agritech VC funds.

c) Public Grants and Subsidies :

— Governmental programs for AI innovation and sustainability.
— Accelerators and incubators supporting hyperspectral and recycling technologies.

d) Crowdfunding :

— Campaigns targeting scientific communities, environmental activists, and precision agriculture supporters.
— Rewards could include early platform access or co-branding on field pilots.

5.12.3 Disbursement Timeline

Date	Amount (DZD)	Use	Balance (est.)
2025-01-01	2,500,000	Initial founder injection	2,500,000
2025-01-15	800,000	Hyperspectral camera acquisition	1,700,000
2025-02-01	400,000	Data acquisition field trips	1,300,000
2025-02-15	300,000	Annotation and labeling tools	1,000,000
2025-03-01	500,000	Software MVP	500,000
2025-03-15	200,000	model deployment	300,000
2025-04-01	300,000	Presentation, publishing, and outreach	0

Table 5.2 – Planned disbursement schedule for early-stage development.

5.12.4 Monitoring and Financial Control

To ensure long-term sustainability :

— Cloud expenses and inference pipelines will be monitored monthly.
— Sensor feedback loops and edge calibration are periodically reviewed.
— Subscription prices may adapt based on usage statistics and domain expansions.

The startup ensures transparency and fiscal responsibility to attract partners and investors. Regular financial dashboards will be updated to reflect operational health and ROI of platform deployments.

Business Model Canvas

Key Partners

- Research institutions
- Annotated data providers (drone operators, waste-management agencies)
- Cloud/hosting services (local Algerian providers)
- Recycling facilities & environmental labs (pilot deployments)

Key Activities

- Data collection & preprocessing (hyperspectral cubes + masks)
- Model training & validation
- Web app integration (real-time inference pipeline)
- User support & performance monitoring
- Ongoing model retraining (new data, accuracy tuning)

Key Resources

- Labeled hyperspectral datasets (MATLAB .mat, NetCDF .nc4, JPEG/BMP masks)
- Trained model weights & code repositories (Python, TensorFlow/PyTorch)
- Web infrastructure (servers, databases, RESTful API)
- GPU compute resources (cloud instances or on-premise GPUs)

Value Proposition

- High-Accuracy Plastic Detection: Unsupervised segmentation of multiple plastic types (PET, HDPE, PP, LDPE, PVC, PS, etc.)
- User-Friendly Web Interface: Upload raw hyperspectral files, get segmentation maps and classification summaries instantly
- Cost & Time Savings: Automates manual sorting, reducing labor and error rates
- Scalable & Flexible: Works with various hyperspectral sensors (drone-based, lab-based), easily extendable to new plastic classes
- Near Real-Time Feedback: Rapid inference results for swift decision-making in field or industrial contexts
- adaptable for any other segmentation and classification tasks or fields

Customer Relationships

- Technical Support Channels: Email & chat troubleshooting for upload issues, result interpretation
- Onboarding & Training: Video tutorials, virtual workshops on data acquisition and result usage
- Regular Updates: In-app notifications or newsletters about model improvements, new features

Channels

- Web Application Portal
- Academic & Industry Conferences
- Professional Networks & Social Media
- Partnerships with Local Agencies
- Email Newsletters

Customer Segments

- Recycling Facilities & Waste-Management Companies: Automate and improve plastic sorting accuracy
- Environmental Research Institutions & Universities: Plastic-pollution studies, mapping plastic types in field data
- Municipal & Regional Agencies (Algeria): Monitoring plastic hotspots (landfills, coastlines) via drone surveys
- Manufacturing & Packaging Industries: Ensure proper sorting and recycling of corporate waste streams
- NGOs & Nonprofits (Marine Debris): Map plastic pollution in inland waterways and coastal zones
- Personal use: like agricultures etc.

Cost Structure

- Data Acquisition: Licensing annotated hyperspectral datasets; drone/sensor rental or maintenance
- Compute & Storage: GPU cloud instances for training/inference; large-scale data storage costs
- Development & Maintenance: Salaries or contractor fees (developers, data scientists, QA)
- Cloud Hosting & Bandwidth: Web-server fees, database costs, file-transfer expenses for large uploads
- Marketing & Outreach: Conference participation, digital ads, promotional materials
- Administrative Overhead: Office expenses, legal/compliance fees (data privacy, Algerian regulations)

Revenue Streams

- Subscription Plans: Tiered monthly or annual fees (number of analyses, premium features, API access)
- Pay-Per-Use Model: Fixed fee per analysis (e.g., per GB of hyperspectral data processed)
- Consulting & Custom Integration: One-time fees for on-premises deployments, tailored model retraining
- Licensing to Academia & Government: Annual licenses for unlimited academic or public-sector access (academic pricing)
- Training Workshops & Certifications: Paid sessions for environmental professionals on hyperspectral plastic detection
- Data Partnerships: Revenue share or fees from data-providing partners (drone companies, lab suppliers)

Figure 5.1 – Business Model Canvas.