Application of Machine Learning Techniques for Identifying Marine Species in Maltese Waters

Benjamin Mifsud Scicluna

A dissertation submitted to the Faculty of Science, University of Malta in (part) fulfilment of the degree of Master of Science in Applied Oceanography

Supervisor: Dr Adam Gauci
Co-supervisor: Prof. Alan Deidun

September 2024

ACKNOWLEDGEMENTS

I wish to express my appreciation to my supervisor, Dr Adam Gauci, and co-supervisor, Prof. Alan Deidun, for their invaluable professional guidance, constant patience, and unwavering support throughout the completion of this dissertation.

I would also like to sincerely thank my family and close friends, especially Emma, Jean, and Jonathan, whose assistance and support throughout this thesis have been indispensable.

ABSTRACT

The challenge of accurately identifying fish species was tackled in this research by employing machine learning and image classification techniques. The main aim was to develop an innovative algorithm capable of dynamically recognising the most common invasive Mediterranean fish species in Maltese coastal waters, based on available images. The target species for this study were identified as *Fistularia commersonii, Lobotes surinamensis, Pomadasys incisus, Siganus luridus,* and *Stephanolepis diaspros*. Machine learning models and transfer learning were utilised to facilitate precise, real-time species identification. The methodology involved the collection and organisation of images, followed by the training of models using consistent datasets to ensure comparable results. Among the models tested, ResNet18 was found to be the most accurate and reliable, with YOLO v8 being demonstrated as similarly effective but less consistent. These findings were highlighted to show the potential of the developed algorithm to significantly contribute to marine biology research, support citizen science initiatives, and enhance environmental management through accurate fish species identification.

Keywords:

image classification; machine learning; convolution neural networks; citizen science; Mediterranean basin; invasive alien species

Table of Contents

LIST OF FIGURES

LIST OF TABLES

LIST OF EQUATIONS

LIST OF ABBREVIATIONS

AI	Artificial Intelligence
ANN	Artificial Neural Networks
API	Application Programming Interface
CIMPAL	Cumulative IMPacts of invasive ALien species
CIS	Cumulative Impact Score
CNN	Convolutional Neural Network
EDM	Euclidean Distance Method
FN	False Negative
FP	False Positive
IAS	Invasive Alien Species
ILSVRC	ImageNet Large Scale Visual Recognition Challenge
IoU	Intersection over Union
IPBES	Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services
IUCN	International Union for Conservation of Nature
mAP	mean Average Precision
ML	Machine Learning
MPAs	Marine Protected Areas
NIS	Non-Indigenous Species
PR	Precision-Recall
PUS	Public Understanding of Science
ReLU	Rectified Linear Unit
SVM	Support Vector Machines
TEU	Twenty-foot Equivalent Units
TL	Transfer Learning
TN	True Negative
TP	True Positive
YOLO	You Only Look Once

1. INTRODUCTION

Biodiversity across various ecosystems is facing growing challenges due to global change. The spread of non-indigenous species (NIS) and invasive alien species (IAS) accompanied by global warming is particularly noteworthy as they often work together to impact ecosystems synergistically (Fleuré et al., 2024). The Mediterranean coastal areas serve as a perfect illustration of these biodiversity shifts. The launch and recent expansion of the Suez Canal has facilitated the entry of over 150 non-indigenous fish species from the Red Sea into the Mediterranean since 1869 (Azzurro et al., 2022). The rapid warming of the Mediterranean surface waters, exceeding 1°C since 1980, has further fuelled the northward and westward expansion of these tropical non-native species, altering both their geographic range and abundance (Shaltout & Omstedt, 2014). This marine phenomenon, known as tropicalisation, occurs when tropical species expand into temperate regions, displacing native species and causing significant ecological and evolutionary shifts (Zarzyczny et al., 2023).

Over the last several years, the utilisation of underwater image processing in species identification has gained significant attention, particularly given the challenge of these newly introduced IAS. Accurate species classification has been a crucial stepping stone in various categories, addressing fields such as fisheries, marine biology, and aquaculture. The identification of fish species serves multiple purposes in fisheries management. First and foremost, it supports after-catch inspection, particularly in countries with regulations prohibiting the fishing of protected species and imposing quotas on fishing vessels. Moreover, accurate species identification is vital for successfully managing fish harvests, and facilitating the establishment of sustainable and profitable fisheries (Ovalle et al., 2022).

Fish species identification is a multifaceted task that used to be done via manual classification methods, these are resource-demanding and liable to human errors, necessitating efficient automatic methods (Barbedo, 2022). The application of artificial intelligence (AI), specifically Convolutional Neural Networks (CNNs), are used for image classification and object classification, and have shown promise (high accuracy scores) in automating this process, reducing reliance on manual intervention (Hassoon, 2022). Furthermore, technological advancements continue to progress, enabling more precise identification of fish species despite challenges such as limited datasets, segmentation errors, image distortion, and object overlap, which collectively undermine the efficacy of existing methods (Rum et al., 2021). Additionally, distinguishing species with similar characteristics is complicated by variations in body colouration, which are further influenced by factors such as light absorption at varying depths (Barbedo, 2022). Consequently, addressing concerns related to image quality is essential for enhancing the accuracy and reliability of underwater image analysis methods (Azzurro et al., 2012).

1.1 Research Questions

This research focuses on developing an algorithm designed to accurately identify common invasive alien marine species in the Mediterranean region, specifically the Maltese islands. This will be done while tackling the challenges posed by underwater image quality, species resemblance, and varying environmental conditions. Accurate identification is crucial for effectively monitoring and controlling the spread of these species, making it essential to create innovative and dependable identification techniques. In this context, the study is guided by the following research questions:

1. How can an algorithm be developed to accurately identify invasive alien marine species in the Mediterranean region?
2. What are the challenges and limitations in accurately identifying invasive species using image-based techniques, especially in underwater settings?
3. How can the integration of citizen science and technological tools improve the tracking and mitigation of invasive marine species?

1.2 Aim and Objectives

The primary aim of this research is to develop an algorithm that accurately identifies invasive alien marine species in the Mediterranean region, addressing challenges such as underwater image quality, species similarity, and environmental variables. To achieve this, the algorithm will be integrated into a website that allows individuals to submit by-catch for identification and documentation. A comprehensive database will be created to track species findings and continuously improve the accuracy of the algorithm over time. This initiative will start locally in the Maltese Islands and gradually expand to cover the entire Mediterranean basin, engaging citizen scientists in the process.

The key objectives of this study include:

1. To develop a robust algorithm that can accurately identify invasive alien marine species in the Mediterranean region, addressing the complexities of underwater image quality and species differentiation.
2. The algorithm will be seamlessly integrated into a user-friendly website, where individuals can submit by-catch for precise identification and receive detailed species descriptions.
3. The research will establish a comprehensive and continuously updated database to document invasive species findings, contributing to a deeper understanding of species spread across the region.
4. This database will also serve as a critical tool for providing early warnings to countries, helping them prepare for potential extreme invasions and protect their local marine ecosystems.

1.3 Structure of Report

This dissertation is organised into five key chapters, each building upon the previous to provide a comprehensive understanding of the research undertaken. Following this introductory chapter, the subsequent sections detail the reviewed literature, methodology, findings, analysis, and implications. Chapter 2 offers a review of the existing literature, addressing critical themes related to the impact of invasive alien species, convolutional neural networks for species identification, and related Mediterranean case studies. In Chapter 3, the research methodology is thoroughly explained, outlining the processes and practical approaches that guided the development and implementation of the models. Chapter 4 presents the results of the study in the form of confusion matrices and graphs showcasing the learning and validation curves. Chapter 5 consists of a detailed analysis of the data collected and the performance metrics of each model, which are then compared to results from previous literature. The final chapter, Chapter 6, concludes the dissertation by summarising the key findings, discussing the limitations of the research, and proposing areas for future investigation.

2. LITERATURE REVIEW

This chapter reviews the literature relevant to this research, focusing on key areas that contribute to the understanding and application of CNNs for the identification of IAS. The review begins by exploring the impact of IAS, addressing their effects on biodiversity and the economy. It then examines the situation in the Mediterranean region, with particular attention to the Maltese Islands, followed by an analysis of the vectors responsible for the spread of IAS. Regional case studies are included to highlight specific examples of IAS introduction and their consequences. Subsequently, the discussion moves to the role of CNNs in species identification, providing an overview of feature extraction methods, the Keras library, transfer learning techniques, and key neural network architectures such as You Only Look Once (YOLO) and ResNet. The chapter also discusses case studies from the Mediterranean, with a focus on the Maltese Islands, that utilise CNNs for species identification. The role of citizen science campaigns in supporting these efforts is examined, and finally, the challenges and knowledge gaps in image classification for IAS identification are addressed.

2.1 The Impact of Invasive Alien Species

2.1.1 Impacts on Biodiversity

Alien species refer to organisms or taxa that have been introduced to regions beyond their natural habitats, often facilitated by human activities. These species possess the ability to survive, reproduce, and propagate outside their original range, exceeding their natural dispersal limitations (Evans et al., 2015). The term 'alien' encompasses various synonyms such as non-native, non-indigenous, allochthonous, foreign, exotic, immigrant, imported, transported, or adventive, as classified within the Colautti & MacIsaac (2004) system (Occhipinti-Ambrogi & Galil, 2010). Within the context of biodiversity conservation, invasive alien species (IAS) pose a significant threat, as their introduction and proliferation can have adverse effects on native ecosystems and species diversity. In addition, the recently published report by the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services (IPBES) highlighted that IAS have contributed, either independently or in combination with other factors, to 60% of documented global extinctions. Moreover, they are solely responsible for 16% of recorded extinctions among animal and plant species worldwide. One major negative impact of IAS is biotic homogenisation, a process that makes biological communities across the globe increasingly similar, causing significant consequences for the structure and functioning of ecosystems (Roy et al., 2023).

Among the many invasive alien species causing ecological disruptions, the Indo-Pacific cornetfish (*Fistularia commersonii*) serves as a notable example. Kalogirou et al. (2007) examined the diet composition of this common IAS in the southeastern Aegean Sea, analysing the stomach contents of 245 specimens caught between September 2004 and March 2005. They compared the diet across different seasons and fish sizes, finding a significant correlation between the length of the predator and the length of its prey. The diet of the blue cornetfish was found to consist of 96% fish by number and 99.95% by weight. Analysis of prey size and habitat (benthic, supra-benthic, and pelagic) indicated that as the cornetfish grows, it shifts to larger prey and adopts a more generalist feeding strategy. *Spicara smaris, Boops boops*, and species from the Mullidae family were the most significant prey by weight, while small benthic fish (particularly gobiids) and newly hatched fish made up the majority of prey items. The Atlantic tripletail, *Lobotes surinamensis*, is a widely distributed migratory species that lives in tropical and temperate marine environments, occupying both benthic and pelagic zones. While this species has been documented in various regions, including the Ligurian Sea, southern Tyrrhenian Sea, and Central Mediterranean Sea, it remains rare in the eastern Mediterranean, with only sporadic sightings (Cooperativa et al., n.d.). Another species that is defined as subtropical, the *Pomadasys incisus*, which entered the Mediterranean Sea naturally through the Straits of Gibraltar in the early nineteenth century, is currently expanding its range across the Mediterranean coast, with the exception of the Adriatic Sea (Bodilis, et al., 2013). To better understand its distribution, the spread and establishment of the bastard grunt was examined as a Mediterranean species in relation to prevailing currents and potential climate changes. In terms of biological characteristics of *P. incisus*, a study by Pajuelo et al. (2003) showcases that this species is gonochoric, exhibiting similar traits in both males and females, with an adult sex ratio close to 1:1 and comparable average sizes between the sexes. It reaches sexual maturity at the end of its second year, at around 183 mm in length, and spawns throughout the year. Growth studies (Nespereira et al., 2003) show that *P. incisus* is fast- growing but relatively short-lived, with a lifespan of about seven years. Fishing has a direct impact on the species, leading to an 80% reduction in abundance from unexploited levels (AL- Nahdi, 2021). Recruitment into commercial fisheries occurs before the fish reach sexual maturity, making all spawners vulnerable to capture under the current fishing practices. The length at first capture (168 mm) is less than the length at maturity, with 45% of the total catch consisting of fish smaller than the mature size, indicating a risk of recruitment overfishing (Pajuelo et al., 2003).

Azzurro et al. (2013) identified that the rocky-reef fish populations around Malta and Lampedusa are of significant biogeographical importance and experience varying levels of protection and human impact. Through the use of underwater visual census methods, 192 surveys were conducted in May–June and September–October 2007, employing a hierarchical spatial design across four depth layers. In total, 23 families and 61 different taxa were identified. Amongst them, two highly invasive species, *Siganus luridus* and *Fistularia commersonii,* were observed, though in relatively low numbers. Native species from the Labridae and Sparidae families were dominant in shaping the assemblage structures of both islands, with thermophilic species like *Sparisoma cretense* and *Thalassoma pavo* being particularly abundant. While the fish assemblages of Malta and Lampedusa showed similarities in species composition, richness, and overall abundance, multivariate analysis revealed significant differences, largely due to disparities in the abundance of Labridae species. Additionally, significant variations in the size distribution of the most common species were observed between the islands, with these differences showing consistent patterns over time. The spatio-temporal variability of the entire assemblage structure closely mirrored that of nekto-benthic fish, emphasising the importance of this group as indicators in future monitoring efforts.

As mentioned in the study of Akyol et al. (2018), the reticulated leatherjacket, *Stephanolepis diaspros*, inhabits inshore sandy and rocky areas with vegetation, typically at depths of up to 20 meters. It primarily feeds on small invertebrates that it extracts from rocks (Zouari-Ktari et al., 2008). The species commonly measures between 7-15 cm, with a recorded maximum size of 20-25 cm. Originally distributed in the Red Sea and the Arabian Gulf, *S. diaspros* entered the eastern and central Mediterranean via the Suez Canal, eventually reaching regions like Tunisia and southern Italy (Rim et al., 2011). As one of the earliest Lessepsian migrants, it was first documented in the Mediterranean along the Palestinian coast. Since then, it has been reported in Iskenderun Bay in the 1950s and is now well-established along the Turkish coasts, particularly in the eastern Mediterranean (Akyol et al., 2018). The species was first recorded in the Aegean Sea in 1943 and has since spread throughout the region, extending its range to the Sea of Marmara, the Adriatic Sea, the Gulf of Palermo in Sicily, Tunisia, and Maltese waters (Rim et al., 2011; Akyol et al., 2018). However, it is important to note that not all alien species are harmful, some may have positive impacts on society and the economy, particularly in sectors like agriculture or aquaculture (Evans et al., 2015; Riley, 2009).

2.1.2 Impacts on Economy

While less common, alien marine species can also have some positive effects, such as enhancing aesthetic values, creating new economic opportunities such as in fisheries and aquaculture, and increasing employment through management projects and programs (Bax et al., 2003). Furthermore, the knowledge gained about ecosystem processes and resource dynamics can be considered as a beneficial outcome.

Invasive alien marine species primarily cause adverse economic and social effects. This includes threats to human health and reduced productivity in marine-dependent sectors like fisheries, aquaculture, tourism, and marine infrastructure (Kishore et al., 2018). These impacts can lead to social consequences, with decreased employment in affected industries and a decline in the well-being of individuals due to the deterioration of their surrounding environment. Economies and societies face opportunity costs as financial resources, labour, and scientific expertise are diverted to managing these species, resulting in lost potential benefits (García-Llorente et al., 2008). The overall societal cost of invasive alien marine species is calculated by subtracting any inter- temporal benefits from the total costs, including the negative impacts and additional expenses for prevention, control, and management (Vilà et al., 2011). Economic impacts are generally measured by the change, usually a decrease, in net social benefits due to the effects of these species on the resource base and the increased management costs (Haubrock et al., 2021).

Impacts on human health are assessed by considering the reduction in working time, leading to lost income, and the additional costs of medical treatment (Mazza et al., 2014; Roy et al., 2023). When human mortality is involved, valuing the impact becomes more complex due to moral and ethical considerations (Vilà et al., 2011). Invasive species are defined as those that have successfully established themselves in a new habitat and rapidly increased in population, posing a threat to native species diversity and abundance, disrupting the ecological balance, and potentially endangering economic activities and human health (Evans et al., 2015). This poses a significant concern for biodiversity conservation, as highlighted by the International Union for Conservation of Nature (IUCN), which correlates their impact to habitat loss (Riley, 2009). Biological invasions in marine ecosystems pose significant threats to both conservation efforts and economic interests. In 2019, the global annual costs associated with these invasions were estimated to exceed US$423 billion (Roy et al., 2023). Given the magnitude of these impacts, it is crucial for national governments and international organizations to promote awareness campaigns (Occhipinti-Ambrogi & Galil, 2010).

The Mediterranean Sea has recorded over 500 alien species that have led to notable impacts such as sudden declines in native species abundance and local extinctions (Galil, 2007). While immediate extinctions may not always occur, the changes induced by invasive species contribute to genetic diversity reduction, loss of ecological functions, and habitat alteration, escalating the risk of further declines and homogenisation of ecosystems (Galil, 2007). The complex interactions between native and invasive species in the Mediterranean Sea remain poorly understood, making it difficult to directly assess or test how competition may lead to niche restriction, displacement, or local extinction (Katsanevakis et al., 2023). In some cases, native Mediterranean species seem to be entirely outcompeted or partially displaced by invasive aliens, which may reflect broader human impacts on the marine ecosystem, such as habitat destruction, pollution, and rising sea temperatures (Otero et al., 2013). Despite these complexities, certain invasive alien species in the Mediterranean have received significant attention from researchers and management agencies due to their prominent effects on native ecosystems. Notably, two coenocytic chlorophytes being the *Caulerpa taxifolia*, known as "the killer alga," and *Caulerpa racemosa var. cylindracea* are among the most infamous and extensively studied (Katsanevakis et al., 2020; Galil, 2007). Further research has examined the effects of other invasive species that have entered the Mediterranean from the Red Sea via the Suez Canal (Galil, 2007).

2.2 IAS in the Mediterranean and the Maltese Islands

The Mediterranean Sea can be regarded as a significant hotspot for marine bioinvasions, given the rapidly increasing rate of species introductions and the extensive volume of shipping traffic (Mannino et al., 2017). This has led to the settlement of various invasive species in the region, including two notable examples: *Siganus luridus* and *S. rivulatus*, which were introduced in 1956 and 1927, respectively (Bariche et al., 2004). While *S. luridus* has established significant populations in both the eastern and western regions, the presence of *S. rivulatus* extends as far west as Corsica (Galanidi et al., 2018). Their impact on algal forests through grazing and herbivory has been documented, resulting in the creation of barren areas lacking macroalgae, particularly affecting species like *Cystoseira spp.*. These herbivorous fish form large schools and thrive in diverse habitats. In the Levant Sea, they are particularly abundant, with juvenile fish extensively feeding on the summer algal cover over rocky substrates (Galil, 2007). In Lebanon, siganids account for 80% of the herbivorous fish in shallow coastal areas and make up one-third of the fish biomass in rocky coastal regions of Israel (Bariche et al., 2004).

A significant outcome from these siganids is that they have displaced native herbivorous fish. An example of which is how the *S. rivulatus* has outcompeted *Boops boops* along the Libyan coast, reducing its numbers due to both species feeding on algae (Kalogirou et al., 2012). This displacement trend is also likely occurring in the southeastern Aegean Sea. Initially, it was uncertain if *S. rivulatus* was replacing *Sarpa salpa*, a species that was once common in trawl catches in the early 20th century, but it has now effectively replaced the native *S. salpa* (Shakman, 2008; Bariche et al., 2004).

Prior to the introduction of these siganids, herbivorous fish and invertebrates in the Mediterranean were minimal, and their role in the Levantine rocky habitats was negligible (Occhipinti-Ambrogi & Galil, 2010). The arrival of siganids caused a significant shift by rapidly recycling large amounts of algal material, thus accelerating the energy transfer from producers to consumers. This transfer happens within hours via the digestive systems of these fish, compared to the weeks or months typically needed for decomposition (Galil, 2007). Moreover, siganids have become a major prey item (up to 70%) for larger predators such as groupers (Goren & Galil, 2005). Their grazing pressure on intertidal rocky algae may have also facilitated the spread of an alien Erythrean mussel by providing a suitable substrate for its settlement (Occhipinti-Ambrogi & Galil, 2010).

Another invasive species, the blue-spotted cornetfish (*Fistularia commersonii*), native to the Indian and Pacific Oceans, has emerged as one of the most successful invaders in the Mediterranean and European waters, first documented along the

Mediterranean coasts of Israel in January 2000 (Azzurro et al., 2012). Similarly, the Atlantic tripletail (*Lobotes surinamensis*) which is defined as a cosmopolitan species, has been found in Maltese coastal waters. However, most of these introduced species, exemplified by the reticulated leatherjacket (*Stephanolepis diaspros)*, are often discarded by local fishermen (Deidun et al., 2010; Pešić et al., 2020). The bastard grunt (*Pomadasys incisus)*, a subtropical species, has naturally migrated into the Mediterranean Sea via the Straits of Gibraltar, along with other species of the Haemulidae family, such as the non-native *P. stridens* (Bodilis et al., 2013). An updated review by Evans et al., (2015) expands upon the previous inventory of marine alien species recorded around the Maltese Islands. With 31 new species added and six removed, the total now stands at 73, with a majority considered aliens or putative aliens. According to this study, the main known factor that introduces an alien is through lessepsian migration (through the Suez Canal), shipping, and aquaculture. Among the taxonomic groups, Mollusca and Actinopterygii are most prevalent. Notably, eight species, including *Fistularia commersonii*, *Siganus luridus* and *Percnon gibbesi*, are all classified as invasive. Shipping remains the primary introduction pathway for these species, followed by secondary dispersal within the Mediterranean Sea (Katsanevakis et al., 2013). An increasing trend in alien species reporting is evident, with a distinctive peak observed during the last decade. The impact of the warming trend in Mediterranean waters is accelerating the spread of thermophilic alien species from the Eastern to the Central Mediterranean, and the outreach of tropical and subtropical Eastern Atlantic species towards the east (Templado, 2014; Evans et al., 2015).

2.3 Vectors of Transport

Many invasive alien species have been deliberately introduced beyond their natural habitats worldwide, often due to their perceived benefits, without adequate consideration of potential negative impacts. Additionally, numerous species have been unintentionally introduced, often as contaminants in traded goods or as stowaways in shipments (Roy et al. 2023).

Shipping is responsible for over 80% of global trade, handling approximately 12 billion tonnes of ballast water annually (Le et al., 2021). Over the past 30 years, world seaborne trade has more than doubled, rising from 2,490 million tonnes in 1970 to 5,330 million tonnes in 2000 (Bax et al., 2003). The largest ships currently in operation can carry 7,500 twenty-foot equivalent units (TEU), but this will soon be exceeded by vessels designed to hold 9,200 TEU, with plans already in place for 12,500 TEU ships (Bax et al., 2003). The registered merchant fleet now includes over 45,000 vessels, with new construction contracts over the last five years adding another 6,000 ships of 300 gross registered tonnes and above (Bax et al., 2003). As the fleet grows, the frequency of ship visits increases, providing more opportunities for species to invade, even into remote, undisturbed, and protected areas (Roy et al., 2023). There are numerous vectors for transporting alien marine species. International shipping and ocean-going recreational vessels facilitate the movement of species through hull fouling, sea chests, ballast water, and other compartments (Minchin, 2006). Although anti-fouling paints and the higher speeds of modern vessels have reduced hull fouling as a vector, it remains significant, especially for smaller vessels (Arndt et al., 2021). In the two years following the 1999 elimination of the black-striped mussel from three Darwin marinas, four non-native taxa were found on yachts, commercial fishing vessels, and unauthorised vessels attempting to enter Darwin harbour (Kuris, 2002). Species that are prone to being transported as hull fouling organisms are often also moved in mariculture shipments, and many of their larvae can travel in ballast water or as juveniles and adults in sea chests or on hulls (Minchin, 2006). Marine organisms can be transported from shallow coastal waters to comparable habitats beyond their native range through multiple broad categories of vectors (Minchin et al., 2005). Historically, these vectors comprised of hull fouling, dry and semi-dry ballast, ballast water, and deliberate introductions for mariculture purposes (Kolar et al., 2010; Bax et al., 2003). More recent vectors encompass the aquarium trade, aquaculture, recreational water users, and the oil, gas, and construction industries (Fisher, 2005). The ongoing evolution nature of national and international shipping is affecting the range and momentum of potential vectors, raising the likelihood that established invasive species will be transported, and that previously species not previously transported will find new pathways for their dispersal.

2.4 Regional Case Studies

In a study conducted by Katsanevakis et al. (2016), the main pathways for the introduction of alien species in the Mediterranean region were mapped (Figure 2.1), while the Cumulative Impact Score (CIS) of each pathway, including shipping, aquaculture, and the Suez Canal, was assessed. These pathways were considered in the context of the Cumulative IMPacts of invasive ALien species (CIMPAL) on marine ecosystems which gave rise to the development of conservative additive model. This model utilised the ranges of invasive species and ecosystems, as well as the reported amount and quality of ecological effect evidence, to predict cumulative impact ratings. Assessments were conducted for each combination of 60 invasive species and 13 habitats, across each 10 x 10 km cell within the Mediterranean basin, some of which include the *Siganus luridus, Fistularia commersonii, Portunus segnis,* and *Caulerpa taxifolia*. The grading of invasive species was determined based on their contributions to the overall effect score of the Mediterranean.

The spatial heterogeneity was seen in the CIMPAL index. The way that the invasive species were first introduced into the Mediterranean Sea caused differences in the spatial patterns. In comparison to species introduced by aquaculture and through the Suez Canal, those introduced by shipping had the highest impact scores and affected a far wider area (Katsanevakis et al., 2016). From all the taxonomic groups, invasive macroalgae had the greatest overall impact. These findings provide the most accurate estimate of the regional heterogeneity in the effects of alien invasive species on Mediterranean Sea ecosystems to date.

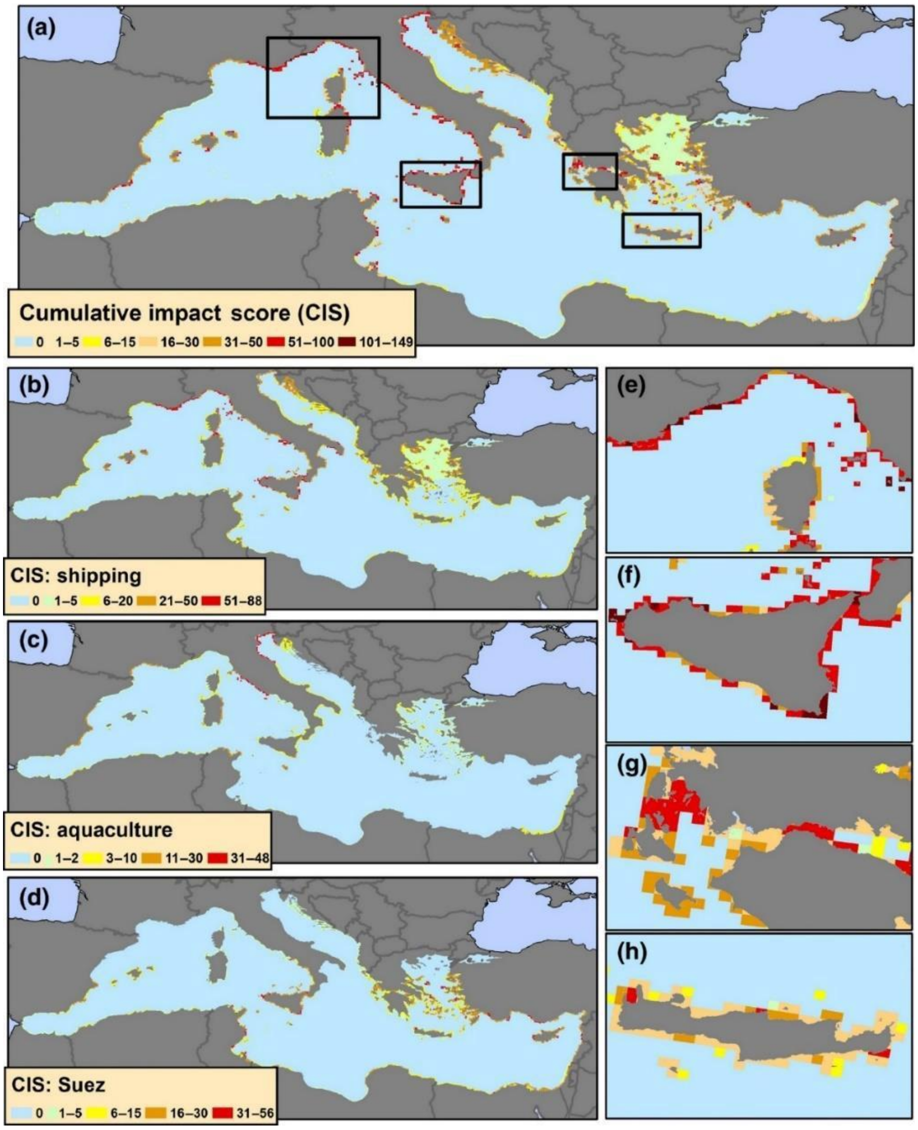

Figure 2.1 - Map (a) of the Mediterranean Sea displaying the cumulative impact (CIMPAL) score of 60 invasive alien species across 13 marine habitats, using an uncertainty-averse approach. Separate maps highlight the CIMPAL scores for marine habitats impacted by species likely introduced via shipping (b), aquaculture (c), and the Suez Canal (d). Detailed views are provided for the Ligurian Sea and Corsica (e), Sicily (f), the Greek Ionian Archipelagos and nearby gulfs (g), and Crete (h) (Katsanevakis et al., 2016).

A complementary study conducted by Kourantidou et al., (2021) estimated the total economic cost of invasive alien species in the Mediterranean basin from 1990 to 2017, which amounted to around $27.31 billion in 2017 US dollars. Most of these recorded costs appeared in publications after the mid-2000s, as evidenced by the orange line in Figure 2.2, with a marked rise in the number of annual cost reports, especially post-2006, as shown by the red line in the same figure. A substantial majority of the costs (87%) were based on predictions or estimates (Potential costs, $23.73 billion), rather than direct empirical observations (observed costs, $3.59 billion). Moreover, nearly 98% of the cost entries for the Mediterranean basin, totalling to $25.89 billion, were considered highly reliable based on their estimation methods.

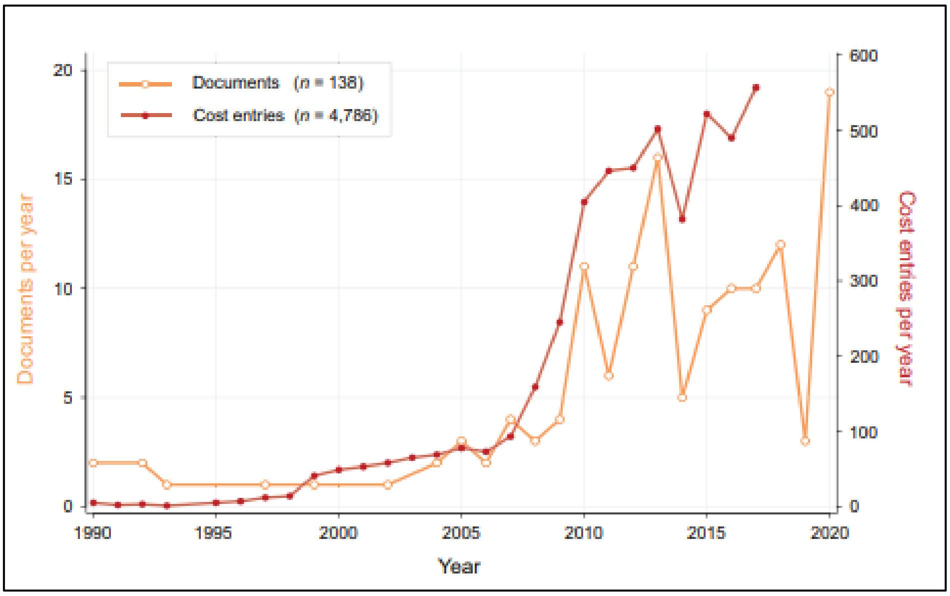

Figure 2.2 - Temporal trends reporting costs in numerous documents and their respective cost entries for IAS within the Mediterranean (Kourantidou et al., 2021).

2.5 Convolutional Neural Networks for Species Identification and Related Studies

Machine learning is considered as a subset of AI and it involves teaching machines to perform tasks without explicit programming by creating models based on training data (El Naqa & Murphy, 2015). Deep learning which is a more advanced branch of machine learning, relies on artificial neural networks with hierarchical learning representations (Sarker, 2021). This includes supervised, semi-supervised, unsupervised, and reinforcement learning approaches. As explained by Morales &

Escalante (2022), the main difference between these learning approaches is the use of labelled data: supervised uses labelled data (as used for this current research study), semi-supervised combines labelled and unlabelled, unsupervised uses only unlabelled data, and reinforcement learning learns through interaction with an environment and feedback.

Deep learning finds application in areas requiring extensive expertise, most notably in medical decision-making, speech recognition, and natural language understanding, as well as in scenarios involving evolving problem solutions or large-scale problems (Hassoon, 2022).

The convolutional neural network stands out as a deep, multilayered neural network tailored for visual pattern detection with minimal preprocessing of image pixels (Rum & Az, 2021). Additionally, there are also advanced deep learning algorithms known as deep CNNs which have proven to be quite effective and beneficial, especially when it comes to learning visual representations (Alotaibi & Alotaibi, 2020). Composed of convolutional and pooling layers, this architecture enables CNNs to effectively capture various aspects of visual information in nearly limitless ways (Li et al., 2021). The use of convolution kernels or filters, CNNs automatically extract relevant features from input images, making them highly efficient for picture classification while requiring minimal preprocessing and accommodating diverse data formats (Hassoon, 2022). Numerous CNN architectures have emerged in advancing AI such as LeNet, VGGNet, AlexNet, ZFNet, ResNet, and GoogLeNet. VGG16 rose to popularity in 2014 when it prevailed in the ImageNet Large Scale Visual Recognition Challenge Competition (ILSVRC) (Rum & Az, 2021).

2.6 Feature Extraction

The process of feature extraction (Figure 2.3) turns unprocessed data into interpretable representations for a particular classification problem (Kumar & Bhatia, 2014). Millions of pixels, each with corresponding colour information, make up an image. By computing abstract features, or a quantified representation of the image that retains pertinent information for the classification task (such as shape, texture, or colour information) and omits unnecessary information, the high dimensionality of these images is decreased (Wäldchen & Mäder, 2018). Traditionally, domain experts would create the features to be retrieved through labour-intensive and subjective manual procedures, for instance, it has been shown that edge detection in images is important to humans (Lei et al., 2016). This pattern is followed by several well-known computer vision algorithms, such as the scale-invariant feature transform (SIFT), which uses edge or gradient-based features (Wäldchen & Mäder, 2018). A popular method for object detection and picture comparison, SIFT effectively finds and defines scale-invariant and characteristic key points in images, offering a significant advancement over previous techniques (Lowe, 2004).

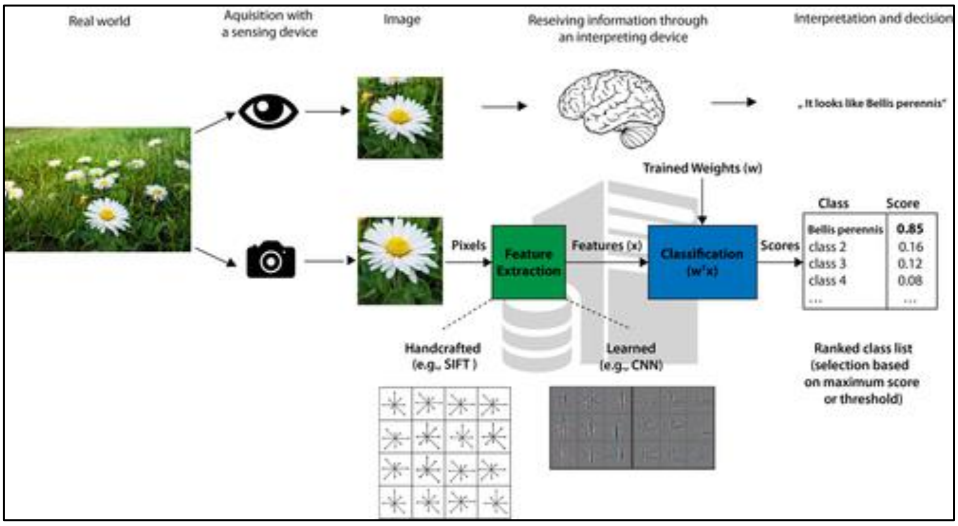

Figure 2.3 - Typical human and computer vision pipeline for species identification. The machine learning platform takes in an image and outputs the confidence scores for a predefined set of classes (Wäldchen & Mäder, 2018).

2.7 Keras Library

The open-source Keras deep learning framework is well known for having an easy-to-use interface. In addition, it removes a lot of the complexity involved in deep learning by offering a high-level Application Programming Interface (API) that makes neural network construction and training easier (Manaswi & Manaswi, 2018). Keras makes it simple for users to create and modify neural network topologies, stack various layer types, define optimisation algorithms and loss functions, and assess model performance (Gulli & Pal, 2017). Nevertheless, Keras has come to be associated with TensorFlow and is now the go-to option for creating and refining neural networks in the TensorFlow environment (Joseph et al., 2021). With its variety of pre-trained models and comprehensive documentation, it is a great choice for both novice and seasoned machine learning practitioners looking for a practical yet effective solution for deep learning applications (Hamzaoui et al., 2023).

2.8 Transfer Learning

Transfer Learning (TL) in Machine Learning (ML) involves employing a model developed for one specific task as the basis for creating a new model tailored to a different task that is still related. (Weiss et al., 2016). This technique is particularly prevalent in computer vision, where a model pre-trained on a large dataset serves as the foundation for addressing a particular problem (Ma et al., 2018). Utilising a pre-trained model, which has already acquired general features from a related task, can significantly reduce the time required compared to developing and training a model from scratch (Rum & Az, 2021; Hamzaoui et al., 2023).

Transfer methods in ML are often tailored to the specific algorithms used for learning tasks and can be seen as extensions of those algorithms (Iman et al., 2023). In the sector of TL, there are primarily two contexts where these methods are applied, inductive learning and reinforcement learning. In inductive learning, the focus is on extending established classification and inference algorithms, such as neural networks, Bayesian networks, and Markov Logic Networks, whilst reinforcement learning emphasises on adapting algorithms like Q-learning and policy search (Richardson & Domingos, 2006).

The central objective of TL is to enhance learning in a target task by utilising knowledge acquired from a source task. Transfer learning can be assessed through three key measures (Torrey & Shavlik, 2010). The first measure is the initial performance of the target task with only the transferred knowledge, compared to the performance of an agent that starts with no prior knowledge. The second measure is the time required to fully learn the target task with the benefit of transferred knowledge versus the time needed to learn the task from scratch. The third measure evaluates the final performance level achieved in the target task with transfer compared to the final level without transfer.

Furthermore, the conduction of extensive studies on the transferability of features pre-learned from the ImageNet dataset, employing various fine-tuning strategies on other datasets. They found that the effectiveness of feature transfer diminishes as the distance between the base task and the target task increases (Han et al., 2018). Donahue et al. (2014) further demonstrated the transferability of parameters from each layer of a pre-trained AlexNet by fine-tuning the network layer by layer. This method has proven highly effective in practice. Nonetheless, several questions remain when applying this method to specific problems: how to set hyper-parameters (e.g., learning rate) for fine-tuning, how to choose the fine-tuning strategy, and whether other methods can further enhance network performance.

A potential drawback of transfer learning is negative transfer, where the application of knowledge from one task impairs performance on another (Zhang et al., 2022). A significant challenge in developing transfer methods is to stimulate positive transfer between tasks that are suitably related while avoiding negative transfer between tasks that are less related (Weiss et al., 2016).

When applying knowledge from one task to another, it is often necessary to establish correspondences between the characteristics of the two tasks (Torrey & Shavlik, 2010). While many transfer learning approaches rely on human-defined mappings for this purpose, there are methods for automatic task mapping.

Furthermore, it is important to differentiate between transfer learning and multi-task learning. Multi-task learning involves learning multiple tasks simultaneously without a designated source or target task, instead, the learning agent receives information about several tasks at once (Caruana, 1997). In contrast, transfer learning, by our definition, involves an agent that initially has no knowledge of the target task or even that there will be a target task while learning the source task (Caruana, 1997). Although transfer learning methods can sometimes be applied to multi-task learning problems, the reverse is not true. This distinction is valuable because, in practical settings, agents are more likely to encounter transfer scenarios than multi-task scenarios.

2.9 TF.Keras

The TF.Keras Sequential model is a high-level API for building neural networks, traditionally used with TensorFlow as the backend (Sarang, 2020). Furthermore, it serves as a fundamental instrument to construct machine learning models, particularly those characterised by a simple stack of layers. With each layer having precisely one input tensor and one output tensor, this model type excels in straightforward architectures (Kapoor et al., 2022). The Sequential Model API enables the systematic layer-by-layer development of models, facilitating a clear progression in complexity (Chicho et al., 2021). However, it is essential to recognise that the Sequential model is not suitable for creating models with multiple inputs or outputs. Instead, it is tailored to efficiently handle scenarios where a basic stack of layers suffices to meet the modelling requirements (Chicho et al., 2021).

2.10 You Only Look Once (YOLO)

Recent advancements in marine resource and environmental research have underscored the critical ecological role of coral reef ecosystems, which flourish in warm, shallow ocean waters (Modasshir & Rekleitis, 2020). The increasing effects of global climate change and the intensifying use of marine resources highlight the urgent need for dynamic and regular monitoring of these ecosystems. Marine animals are essential for maintaining the balance of coral reefs, and their presence often indicates a healthy environment, hence tracking marine animal populations is crucial for effective ecosystem management (Goreau & Hilbertz, 2005).

To meet this need, the rise of autonomous underwater vehicles, which are used for tasks such as capturing marine life, conducting ecological surveys, and monitoring biodiversity (Kumar et al., 2018), has created a growing demand for real-time environmental monitoring technologies. These technologies support more efficient ecosystem management by enabling continuous observation of marine life in their natural habitats.

As explained by Zhong et al. (2022), the objective in practical applications is to create marine animal detectors that are both accurate and efficient. Ideally, these detectors should be lightweight and durable. Traditionally, shallow learning architectures were used for fish detection but struggled with generalisation and failed to extract high-level features from complex, changing environments. However, advances in deep learning have led to the application of deep neural networks in fish-related tasks, yielding improved results. CNNs are particularly effective due to their ability to automatically learn features and detect objects with high precision (Salman et al., 2020 & Zhong et al., 2022).

Collecting underwater images remains difficult and expensive, and annotating these images is a costly process. The dataset for underwater object detection is relatively limited, making data augmentation essential to improve model focus on semantic information (Rizzini et al., 2015). Additionally, achieving a balance between precision and speed in detection models is increasingly challenging.

Nevertheless, a computer vision technology called object detection is used to find and identify items in pictures or movies. A key step in this process is image localisation, which involves using bounding boxes (rectangular shapes that indicate the locations of objects) to accurately identify item locations. This method is not the same as image recognition or classification, where the objective is to identify the class or category of an image as a whole or individual items in it. The ground-breaking real-time object detection technique known as You Only Look Once (YOLO) was first introduced by Redmon et al. in 2016 as a deep learning model designed for object detection with a focus on high computational speed, enabling true real - time detection. Unlike previous models, YOLO operates as a region-based convolutional neural network that merges the region proposal and classification stages into a single network. This allows YOLO to directly predict bounding boxes and their associated class probabilities with a single feed-forward pass, significantly reducing the time needed for region proposals compared to earlier models (Al Muksit et al., 2022).

YOLO-based methods have emerged as a leading solution, offering strong performance in both precision and speed, and are currently a major focus of research. In computer vision, YOLO is a popular technique for object recognition. It is a machine learning system that can identify and categorise objects in photos or videos in real-time, frequently producing findings that are more accurate than those obtained using more conventional techniques (Park et al., 2021). YOLO is described as having high detection accuracy, with good generalisation, apart from being open-source (Sanchez et al., 2020).

YOLO has two fully connected layers, four max-pooling layers, and 24 convolutional layers total, as shown below (Figure 2.4).

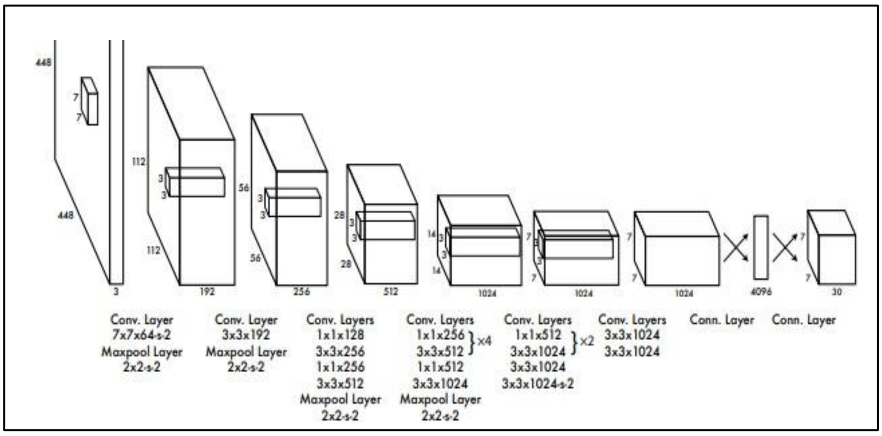

Figure 2.4 – The architecture used by YOLO (Redmon et al., 2016).

The architecture functions by resizing the image to 448 by 448 before passing it through the convolutional network. This process involves applying a 3x3 convolution after a 1x1 convolution to reduce the number of channels and produce a cuboidal output. Throughout all layers, a linear activation function is employed, except for the last layer, which utilises a Rectified Linear Unit (ReLU) as the activation function. Additionally, techniques such as batch normalization and dropout are incorporated to regularise the model and prevent overfitting (Sanchez et al., 2020).

2.11 ResNet

ResNet, or residual networks, pioneered by Microsoft researchers in 2016, placed first in the ILSVRC (Figure 2.5) with an impressive 96.4% accuracy rate (Alotaibi & Alotaibi, 2020). This innovation revolutionised convolutional neural networks (CNNs), achieving state-of-the-art performance on the ILSVRC 2015 classification task and facilitating the training of networks exceeding 1000 layers (Deng et al., 2009). ResNets, similar to highway networks, employ identity shortcut connections, ensuring seamless information flow across layers and mitigating attenuation resulting from stacked non-linear transformations (Targ et al., 2016).

He et al. (2016a) introduced the deep residual network (ResNet), a formidable deep learning model surpassing its CNN counterparts. The effectiveness of ResNet is evident in various domains, evidenced by its enhanced performance in well-established hyperspectral image classification datasets (Zhong et al., 2017). Notably, the architecture of ResNet consists of 152 layers, featuring innovative residual blocks that address the challenge of training deep architectures with identity skip links. These residual blocks ensure the continuity of information flow by copying inputs from preceding layers, thus bypassing the issue of vanishing gradients (Alotaibi & Alotaibi, 2020).

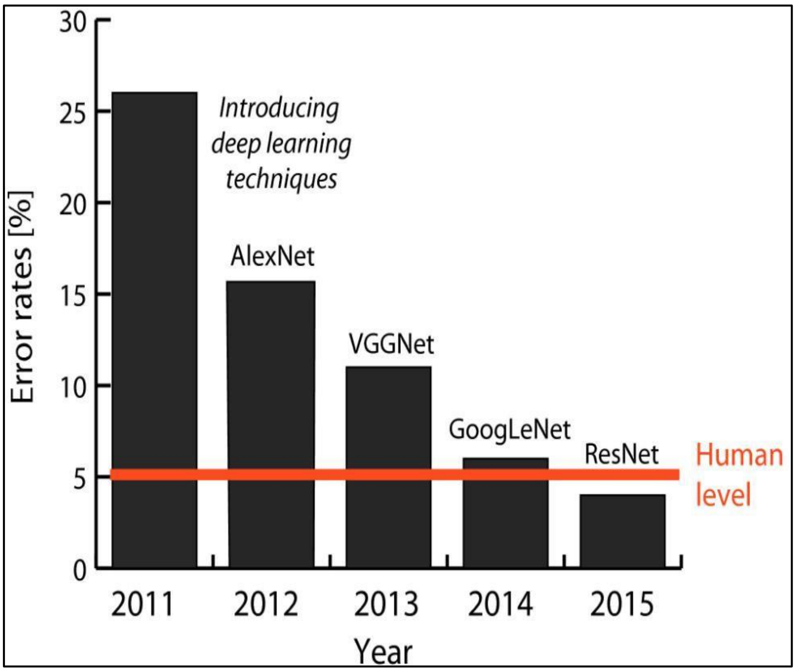

Figure 2.5 - Top-5 ImageNet Visual Recognition Challenge Classification Error Rates Challenge (Wäldchen & Mäder, 2018).

A study by Mahmood, et al. (2020) demonstrates that a well-optimised, high-performing deep network trained on ImageNet will yield more effective and generalisable image representations. One such network is the deep residual network, ResNet. Compared to other CNN architectures like VGGnet, ResNets are easier to train. For instance, a 152-layer ResNet, which is eight times deeper than VGGnet, remains less complex and trains more efficiently (Ikechukwu et al., 2021). Furthermore, a 34-layer ResNet performs 3.6 billion multiply-add operations, whereas a 19-layer VGGnet requires 19.6 billion multiply-add operations, representing a 20% reduction in computational complexity (Mahmood et al., 2020). While very deep networks can lead to overfitting and accuracy saturation, ResNets address these issues through residual learning and identity mappings, allowing them to excel in image detection, localisation, and segmentation tasks (He et al., 2016 b). Moreover, ResNet is available in various configurations with different numbers of layers, including 34, 50, 101, and 152 layers. Among these, ResNet-50 is particularly well-known and comprises 50 layers, consisting of 49 convolutional layers followed by a final fully connected classification layer. The ResNet architecture (Figure 2.6) is built upon multiple basic residual blocks, which include ReLU activation functions and convolutional layers (Muslim & Zulkifli, 2021).

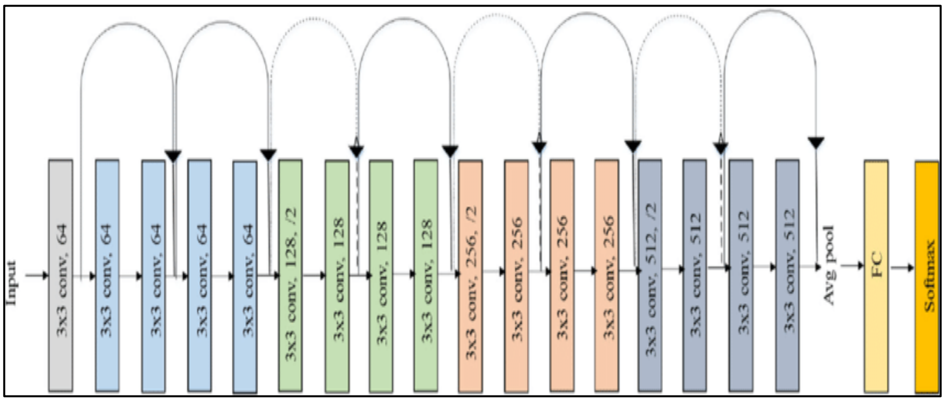

Figure 2.6 - A visualisation of the full ResNet18 architecture (Ramzan et al., 2019).

2.12 Mediterranean Case studies utilising CNNs

Catalán et al. (2023) focused on improving fish identification and classification in underwater images by evaluating the performance of two leading object detectors, YOLOv5m and Faster RCNN. The authors introduced a sizable labelled dataset encompassing 18,400 Mediterranean fish instances from 20 species across 1,600 images with varied backgrounds. The findings indicated that YOLOv5m outperformed Faster RCNN, and as evidenced by its mAP@0.5 values, which exceeded 0.8 across most scenarios.

Further analysis revealed the dominance of YOLO in diverse scenarios, showcasing promising outcomes with reduced labelled sets, diverse background incorporation during training, and high-quality labelling. Its performance consistency across various test scenarios, especially when trained with balanced datasets (E3 and E5 scenarios, Figure 2.7). Remarkably, even with fewer training instances, YOLO exhibited competitive results (Figure 2.8), especially when excluding small fish and employing balanced backgrounds.

Training Scenario	Model	P	R	mAP@0.5
E0	Faster RCNN	0.80	0.48	0.60
	YOLO	0.83	0.78	0.84
E1	Faster RCNN	0.66	0.45	0.37
	YOLO	0.75	0.75	0.77
E2	Faster RCNN	0.76	0.52	0.83
	YOLO	0.84	0.73	0.80
E3	Faster RCNN	0.81	0.45	0.70
	YOLO	0.82	0.73	0.80
E4	Faster RCNN	0.71	0.53	0.42
	YOLO	0.81	0.79	0.84
E5	Faster RCNN	0.78	0.53	0.83
	YOLO	0.88	0.71	0.83

See Table 4 for test sets. Noticeably, E5 yielded relatively good results with a low number of training objects (by eliminating fish that are only dots or very difficult to recognize at the species level).

Figure 2.7 - Performance metrics over the training datasets (Catalán et al., 2023).

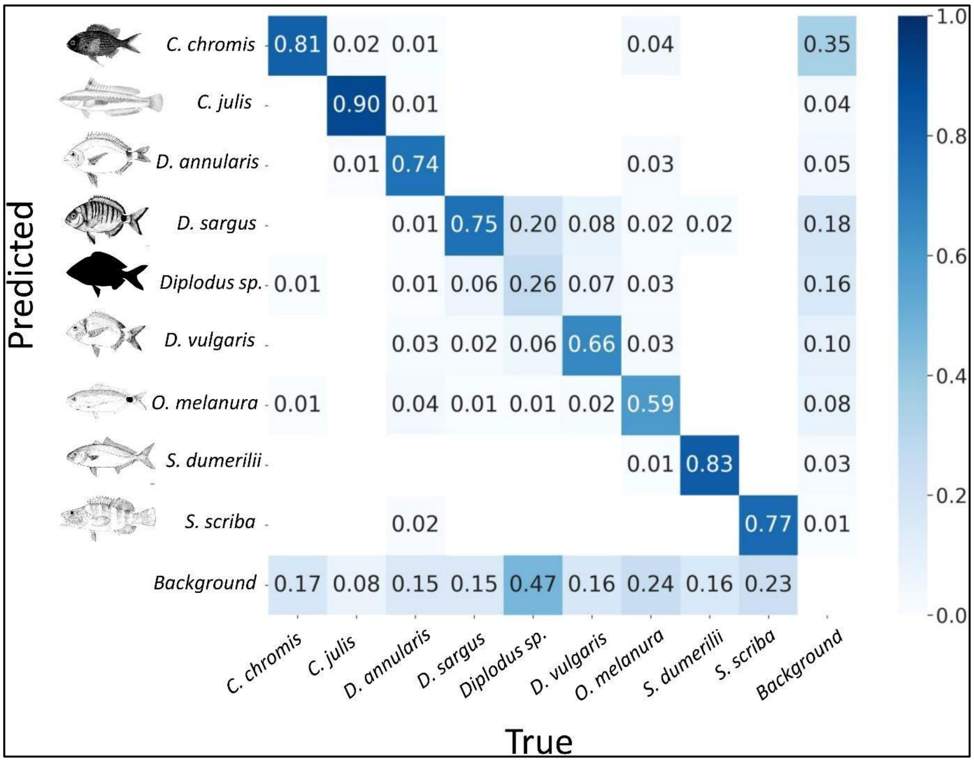

Figure 2.8 - Confusion matrix based on YOLO results for eight species (Catalán et al., 2023).

A corresponding study by Fleuré et al. (2024) examines the challenge of detecting and quantifying invasive rabbitfish species in Mediterranean coastal ecosystems using underwater image analysis. The coastal regions of the Mediterranean Sea harbour a diverse array of NIS, primarily introduced from the Red Sea via the Suez Canal. The herbivorous rabbitfishes *Siganus rivulatus* and *Siganus luridus*, are a point of interest as they have become invasive in the southeastern Mediterranean, leading to ecological damage through algae overgrazing.

To address this challenge, the study compiled a dataset of 31,285 images, including *Siganus spp*. and six common native Mediterranean fish species, sourced from 40 underwater videos across three reef habitats. Subsequently, a deep learning algorithm was trained to identify *Siganus spp*. within images containing these eight species. The performance metrics of the model were assessed using an independent dataset of 2024 images.

The results (Figure 2.9) showed that the model achieved a recall of 0.92 for the *Siganus* genus, which improved to 0.98 after confidence-based post-processing, with only 4 out of 272 *Siganus spp.* images misclassified. While the accuracy reached 0.61, requiring expert verification to discard false positives, images of native species not included in the training data exhibited similar false positive rates. Overall, the study demonstrated that adopting the model for automatic image processing followed by expert verification reduced processing efforts by up to five times compared to full manual processing.

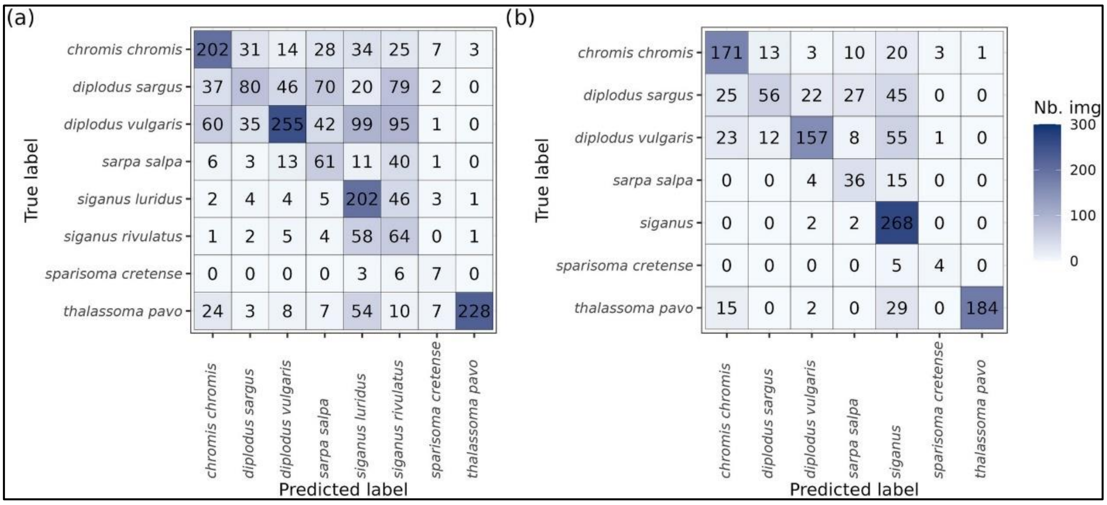

Figure 2.9 - Performance of the algorithm to identify underwater images of the *Siganus spp*. (Fleuré et al. 2024).

2.12.1 Maltese Case Study

The study by Gauci et al. (2020) focuses on leveraging citizen science initiatives for the classification of jellyfish species in Maltese waters using CNNs. With the increasing popularity of citizen science campaigns, the collection of data, especially in marine domains, has become more widespread. However, validating submitted reports, particularly regarding the taxonomic identification of jellyfish, remains a challenge due to the volume of reports and limited trained staff.

The photos used for their dataset were from the "Spot the Jellyfish" campaign to train region-based CNNs to classify the five most commonly reported jellyfish species in Maltese waters: *Pelagia noctiluca, Cotylorhiza tuberculata, Carybdea marsupialis, Velella velella,* and salps. They evaluated the reliability of their models using precision, recall, $f1$ score, and κ score metrics. Additionally, they explored the benefits of data augmentation and transfer learning techniques.

The results (Figure 2.10) showed promising outcomes, suggesting the potential for embedding automated classification methods, possibly in smartphone apps, to reduce or eliminate the need for human validation of citizen science reports.

Confusion matrices were employed to assess classification accuracy, and precision, recall, and $f1$ scores were computed for each class. The study found that models utilising the GoogLeNet feature extraction network performed best, with an increase in accuracy observed with more anchors.

Model	Species	Precision	Recall	f_1 Score	*kappa*
1	*Carybdea marsupialis*	0.92	0.90	0.91	0.92
	Cotylorhiza tuberculata	1.00	0.95	0.97	
	Pelagia noctiluca	0.90	0.87	0.88	
	Salps	0.91	0.98	0.94	
	Velella velella	1.00	0.98	0.99	
2	*Carybdea marsupialis*	0.97	0.93	0.95	0.95
	Cotylorhiza tuberculata	1.00	0.95	0.97	
	Pelagia noctiluca	0.92	0.90	0.91	
	Salps	0.93	1.00	0.96	
	Velella velella	0.98	1.00	0.99	
3	*Carybdea marsupialis*	0.98	1.00	0.99	0.96
	Cotylorhiza tuberculata	0.95	0.98	0.96	
	Pelagia noctiluca	0.91	0.98	0.94	
	Salps	1.00	0.98	0.99	
	Velella velella	1.00	0.90	0.95	
4	*Carybdea marsupialis*	0.95	0.95	0.95	0.92
	Cotylorhiza tuberculata	1.00	0.85	0.92	
	Pelagia noctiluca	0.81	0.88	0.84	
	Salps	0.93	0.98	0.95	
	Velella velella	0.98	1.00	0.99	
5	*Carybdea marsupialis*	0.98	1.00	0.99	0.99
	Cotylorhiza tuberculata	1.00	1.00	1.00	
	Pelagia noctiluca	1.00	0.98	0.99	
	Salps	1.00	1.00	1.00	
	Velella velella	1.00	1.00	1.00	

Figure 2.10 - Performance of the five different models: 1) ResNet50, 2),3), & 5) GoogleNet with varying feature layers, and 4) ResNet18 (Gauci et al., 2020).

Furthermore, the study compared the performance of different CNN architectures, such as ResNet18 and ResNet50, noting that ResNet50 provided better results with only a marginal increase in processing time.

2.13 Citizen Science Campaigns

Citizen science involves a partnership between scientists and the public, where volunteers assist in collecting and analysing data on natural phenomena that are challenging to monitor with traditional scientific methods (Dosemagen & Parker, 2019). Marine citizen science has extensive potential, covering areas such as terrestrial and shoreline observations (these include reports of stranded organisms, litter ("ANDROMEDA" project[1]), and organic matter), shallow water monitoring (for instance; changes in protected benthic communities, coral, and artificial reefs), and open sea sampling (such as; data collection from opportunistic ships and ferry boxes for underwater sampling) (Deidun & Sciberras, 2017).

Campaigns involving citizen science are valuable tools that offer several advantages, such as enhancing awareness and facilitating data collection. For instance, they can significantly increase the volume of metadata, which is essential for achieving the geographic coverage necessary to identify trends and patterns in species populations. By expanding the dataset size, citizen science initiatives can reduce errors and biases, as they help eliminate erroneous and erratic reports submitted by previous individuals (Dickinson et al., 2010). However, there are also challenges associated with citizen science data. These include potential biases and errors related to sampling efforts and techniques, as well as concerns regarding data quality (Pateman et al., 2021).

For the context of the Maltese Islands, the "Spot the Alien Fish" and "Spot the Alien" citizen science campaigns were initiated by the University of Malta in 2017 and 2019, respectively. These campaigns aim to gather and consolidate citizen science reports on NIS within Maltese waters into a national database. Participants can submit their observations through a dedicated web portal, as well as via the campaigns' social

---

[1] https://www.um.edu.mt/research/oceanographymalta/projects/andromeda/

media pages and email addresses (Deidun et al., 2021). Since its launch, the campaign has led to the discovery of two previously unrecorded fish species in Maltese nearshore waters: the Guinea angelfish (*Holacanthus africanus*), with two individuals caught by recreational fishermen, and the Azure damselfish (*Chrysiptera hyanacea*), captured on video by a SCUBA diver. The campaign has also generated numerous reports of other non-indigenous species, such as *Fistularia commersonii, Siganus luridus, Seriola fasciata,* and *Cephalopholis taeniops*, aiding in the monitoring of their population trends. Additionally, the campaign facilitated the second record of the African moonfish (*Selene dorsalis*) in Maltese waters.

Interestingly enough, there are specific methods that can improve and engage the public more with these types of campaigns. A case in point was evaluated by Bowser et al. (2013), in which they explored the possibility of gamifying citizen science initiatives, with the scope of enhancing engagement amongst current volunteers and drawing in new participants. To investigate this, they evaluated Biotracker, a gamified mobile app designed to collect plant phenology data. Their main focus was on engaging Millennials, known for their enthusiasm for technology, with this app. The study assessed how gamification could benefit users and the potential advantages of the app. Findings indicate that gamification significantly attracts Millennials, with social motivations playing a major role and educational aspects having a lesser impact.

Bonney et al. (2015) investigated the evolution of citizen science from a novel idea to a concept nearly as prominent as science itself. Initially a unique approach involving public participation in scientific research, citizen science is now becoming mainstream. Each year, the reach and accessibility of citizen science projects expand, offering more engaging, enjoyable, and rewarding experiences. This increased participation often leads to a greater understanding of scientific content, the processes of science, and the nature of scientific inquiry. Additionally, emerging evidence suggests that citizen science can have a profound impact on the lives of participants. Whether citizen science can contribute meaningfully to current models of Public Understanding of Science (PUS) as deliberative, participatory, and crucial for science-society relationships, depends on several factors. Achieving a deep PUS through citizen science will necessitate a better grasp of project design to meet specific goals. Developers will need to integrate lay knowledge with scientific expertise to frame relevant questions and empower individuals in decision-making processes. Both local and scientific communities must adopt new methods of dialogue and deliberation, challenging traditional views on expert knowledge and the value of different types of knowledge within science. Additionally, addressing issues of trust, fairness, equity, and risk should be seamlessly incorporated into discussions alongside considerations of volunteer recruitment, protocols, and data quality.

2.14 Challenges & Knowledge Gaps in Image Classification

Despite advancements in technology, the accurate identification of fish species continues to present numerous challenges, such as limited datasets, segmentation errors, distortion, and overlap of objects in images, all of which hinder the effectiveness of existing methods (Rum & Az, 2021). Furthermore, distinguishing between species with similar characteristics, heightened by variations in body colouration due to factors like light absorption at different depths, presents additional obstacles (Barbedo, 2022). Another difficulty is the degradation of image quality in underwater settings which is consistently highlighted as a significant factor contributing to misestimates and misclassifications in image-based techniques. This issue underscores the importance of addressing image quality concerns to improve the accuracy and reliability of underwater image analysis methods (An et al., 2021).

As discussed by Salman et al. (2016), unconstrained underwater fish classification is particularly challenging due to the complex conditions present in natural environments. Factors such as varying lighting conditions, water turbidity, background confusion from reef structures and underwater vegetation, and intra-species variability due to the movement of the fish all contribute to the difficulty of accurate classification. Videos of fish are typically recorded with digital cameras, and since there is no control over the underwater environment, these factors can significantly complicate the task. The variability in underwater imagery can greatly impact classification performance. To address these challenges, one approach involves using hierarchical classification trees combined with Support Vector Machines (SVMs) trained on image features. This method employs Gaussian Mixture Modelling for decision-making, which has shown notable improvements over Principal Component Analysis and standard SVM classification techniques. Fish species recognition is inherently subject to challenges such as video quality, water turbidity, algal growth, background patterns of coral reefs, and fluctuations in light intensity, all of which can significantly affect the performance of classification methods.

Another study, done by Iqbal et al. (2019) explored the complex issue of identifying and recognising fish species. Researchers face numerous obstacles, including distortion, noise, segmentation errors, and occlusion. Historically, studies were confined to controlled environments, primarily focusing on terrestrial objects. However, the rapid increase in demand for identifying and recognising aquatic species has highlighted the limitations of these traditional approaches. Classifying fish in natural, unconstrained environments presents significant challenges. Variability in lighting conditions, background confusion caused by coral reefs and aquatic plants, and water turbidity, complicate the process. Additionally, the similarity in shape, texture, and colour among different fish species, further complicates accurate classification.

Efforts to identify visual images have been ongoing, yet remain unresolved due to issues such as segmentation errors, distortion, and the overlapping and occlusion of objects within coloured images (Bai et al., 2008; Kim & Hong, 2009). The core challenge in object classification involves accurately determining the occurrence of each fish species. Resolving this necessitates solutions capable of addressing issues like fish size and orientation variability, picture quality, and segmentation in an automated fish classification system (Rum & Az, 2021).

One approach to addressing the challenge of limited sample size and data diversity involves engaging individuals outside the research community through citizen science initiatives. By involving a broader range of contributors, datasets can be enriched thereby enhancing the robustness and representativeness of data used in research (Barbedo, 2022). Furthermore, existing datasets, such as Fish4Knowledge[2], often lack the characteristics relevant to marine and freshwater aquaculture, creating a notable data gap. Thus, there is a pressing need to develop more comprehensive datasets that encompass a broader range of features, thereby addressing this limitation and expanding these datasets used in research (Liu et al., 2021).

[2] https://homepages.inf.ed.ac.uk/rbf/fish4knowledge/

3. METHODOLOGY

This chapter outlines the methodology employed in this research, detailing the steps taken to achieve the objectives of this study. It begins by describing the area of study, providing context for the geographical and environmental scope of the research. Following this, the framework reinforcing the research is explained, establishing the theoretical and practical structure guiding the investigation. Furthermore, the process of generating the image dataset, which serves as the foundation for developing the classification models is elaborated upon. This is followed by an overview of the model development process, including the feature extraction and selection techniques. Finally, the training process for the models is subsequently discussed, highlighting the approaches used to optimise performance.

3.1 **Area of Study**

The Maltese Islands are located in the centre of the Mediterranean and cover approximately 316 square kilometres. They exhibit a typical Mediterranean climate with an average annual rainfall of around 530 mm, predominantly occurring between October and March. The mean monthly temperature ranges from 12°C to 26°C, and the islands are notably windy and sunny (Schembri, 1997).

As explained by Sciberras & Schembri (2007), the Maltese Islands' location at the biogeographic boundary between the western and eastern Mediterranean makes them a critical monitoring point for the entry and spread of alien marine species. Increasing marine traffic, both commercial and touristic, heightens the risk of new invasions. The islands also serve as a convergence point for alien species from the Atlantic and the Red Sea, creating unique opportunities to study the interactions between different biogeographic populations.

The invasion of native ecosystems by non-indigenous species poses a significant threat to the integrity of biological communities, economies, and even human health, a concern recognised globally (Vitousek et al., 1996). Invasive marine species are identified as one of the four greatest threats to the oceans of the world (IMO, 2000-2007). These invasive species can impact local ecosystems through predation, competition, genetic contamination, habitat modification, and the introduction of new parasites and pathogens. Human activities often exacerbate these invasions, accelerating the rate of introduction beyond natural levels (Johnston et al., 2017).

Since the early 20th century, there has been a notable increase in the number of alien marine species reported in Maltese waters. Shipping, aquaculture, and the expansion of Lessepsian immigrants (species migrating through the Suez Canal) are the primary vectors for these invasions, contributing to 20%, 11%, and 32% of the introductions, respectively (Sciberras & Schembri 2007). Additionally, the warming trend of Mediterranean waters and increasing marine traffic facilitate the spread of warm-water species from the Atlantic and Indo-Pacific regions into Maltese waters (Mannino et al., 2017).

3.2 **Framework**

As observed in Figure 3.1, the comprehensive workflow of the image classification process can be observed. It initiates with an image dataset that serves as the source of data for model training. The next step is pre-processing the raw images to enhance their quality and suitability for training. This stage includes resizing, normalisation, and noise reduction to ensure the images are optimised for the next steps. To further improve the dataset, data augmentation techniques are applied, creating modified versions of the images through rotations, flips, zooms, and other transformations. Additionally, this increases the variability and size of the dataset, making the model more robust and less prone to overfitting.

Following data augmentation, the images are fed into the model training phase, where the machine learning model learns to recognise patterns and features within the images. During this phase, transfer learning can be optionally applied, using a pre-trained model on a similar task and fine-tuning it for the specific dataset and classification task at hand. Once the training is complete, the resulting trained model is capable of extracting relevant features from new images. These features, which capture important information such as edges, textures, and shapes, are stored in a feature database. This database serves as a repository of learned features, facilitating future reference and comparison.

In the image classification stage, the trained model uses the extracted features to classify new images. It compares the features of the new images with those stored in the feature database to determine the most likely class for each image. The final step in the workflow is producing an output description, where the model provides the classification results, indicating the identified class or label for each image. This systematic approach, incorporating data augmentation, model training, and feature extraction, ensures accurate and reliable image classification results.

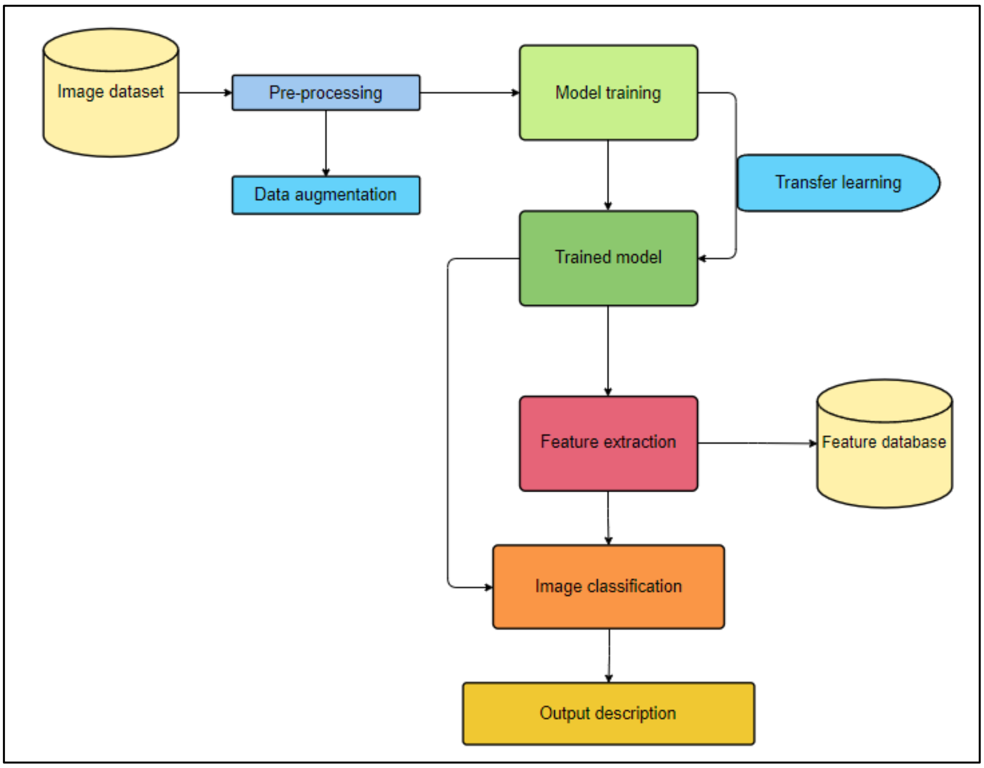

Figure 3.1 - Illustrates the proposed model training framework, which begins with the pre-processing and augmentation of an image dataset. This dataset is then utilised to train the model using transfer learning methods. In the final stage, the model extracts features from the images, saves them in a database, and produces a descriptive output (Mifsud Scicluna et al., 2024).

3.3 **Generation of the Image Dataset**

A diverse dataset was assembled by collecting images from multiple sources, such as Google, FishBase[3], iNaturalist[4], contributions from divers, and the Spot the Alien[5] citizen science initiative. The images shown in Figure 3.2 varied in terms of dimensions, resolution, angles, and colour patterns, and included both juvenile and adult specimens. This variation in the dataset allowed for comprehensive training of the algorithm, helping it to accurately learn and adjust to the complexities of identifying different fish species.

[3] https://fishbase.se/search.php

[4] https://www.inaturalist.org/

[5] https://www.aliensmalta.eu/

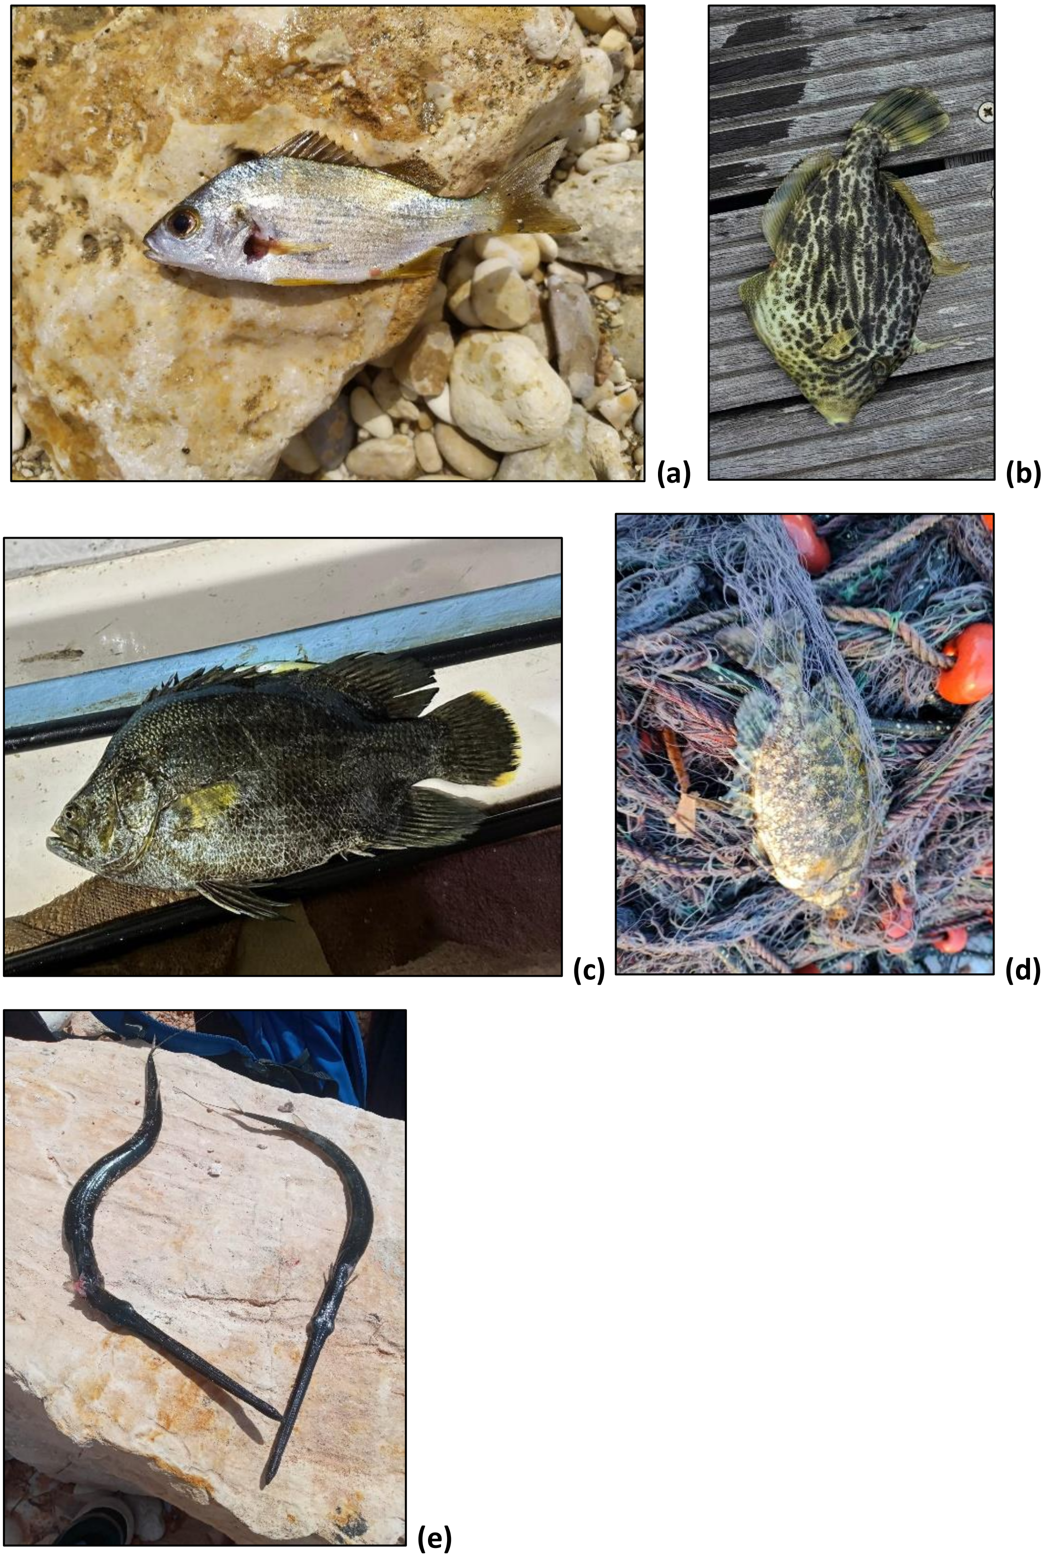

Figure 3.2 - The five invasive fish species chosen for this research study; (a) Pomadasys incisus, (b) Stephanolepis diaspros, (c) Lobotes surinamensis, (d) Siganus luridus, & (e) Fistularia commersonii (Source: Spot the Alien).

Table 1 presents the distribution of images among various categories, detailing the final dataset of five fish species used in this study, which includes a total of 1,337 images. These species were chosen for two main reasons: their common presence in Maltese waters and the availability of numerous images, which is crucial for effective model training. The dataset was divided into training, validation, and testing subsets, with proportions of 70%, 20%, and 10%, respectively, to maintain consistency during model training. All models were trained on this standardised dataset to ensure comparable results. However, it is important to note that approximately 15% of the test images for the YOLO model were excluded because they were not recognised due to format limitations imposed.

Table 1 - The number of images allocated for training, validation, and testing for each fish species.

Species	Number of training images	Number of validation images	Number of testing images
Fistularia commersonii	277	79	41
Lobotes surinamensis	203	58	30
Pomadasys incisus	123	35	19
Siganus luridus	177	50	26
Stephanolepis diaspros	153	43	23

During the pre-processing phase, images were resized to suit the requirements of various models. Irrelevant backgrounds were cropped out, and bounding boxes were added for the YOLO model to concentrate on pertinent objects. YOLO v8 does not have specific pixel size requirements, but images for ResNet18 were resized to a uniform dimension of 224 pixels, while those for the TF.Keras sequential model were resized to 180 pixels. This adjustment in image sizes ensured consistency in training and evaluation across different models. Moreover, data augmentation techniques were applied to increase the diversity and robustness of the datasets.

3.4 **Development of the Classification Models**

Given the limited number of high-resolution images, the current dataset was used until a more extensive high-resolution dataset becomes available for future testing. Nonetheless, it is crucial to include a variety of image resolutions and augmented data in the training process to improve the robustness of the algorithm (Li et al., 2000). For the YOLO model, ROBOFLOW was employed for annotating all images (see Figure 3.3) and for manually creating bounding boxes around key objects.

(a)

(b)

Figure 3.3 - Bounding boxes applied using ROBOFLOW to identify two species: (a) *Fistularia commersonii* and (b) *Lobotomes surminensis*. Each bounding box highlights the individual species within the image for clearer recognition and analysis *(Source: Spot the Alien).*

In this study, three models were adopted to distinguish between five common invasive fish species, drawing from a range of journal articles and sources. From these models, two are image classifiers (ResNet & TF.Keras), whilst the third one, YOLO, serves as an object detector. Although all these models are commonly used in species identification research, image classifiers tend to be the favoured option for this purpose. The final models used were built upon the open versions freely available on the internet, and were further optimised with the use of transfer learning to fit the main aims of this dissertation.

Transfer learning in ML entails using a model designed for one task to serve as the basis for developing a new model aimed at a different, yet related task. This approach utilises insights gained from addressing one problem and applies them to another, thereby promoting more efficient learning. In computer vision, transfer learning is particularly prevalent as it typically involves employing a pre-trained model from an extensive database as a starting point for a specific task (Ma et al., 2018). This method involves employing a pre-trained convolutional neural network that has learned to extract features from thousands of natural images across different categories. By using a model that has already developed generalised features from a related task, transfer learning considerably decreases the time and computational resources required when compared to developing and training a new model from scratch (Rum & Az, 2021).

3.5 **Feature Extraction and Selection**

Subsequently, the models were trained and optimised to extract the features from the fish images. Feature extraction is a crucial step in the process of transforming raw data into interpretable representations tailored for a specific classification task. In the context of fish images, this involves reducing the high dimensionality of the images, which consist of millions of pixels each containing colour information, by calculating abstract features. These features are quantified representations that capture relevant information such as shape, texture, or colour while discarding redundant details (Jogin et al., 2018; Wäldchen & Mäder, 2018). The extracted features are then stored in a feature database. This database serves as a repository of filtered information, which enhances the efficiency and accuracy of the subsequent classification task. By focusing on the most pertinent aspects of the data, the classification process becomes more streamlined and effective.

The next step is image classification, where the model uses the stored features to categorise the images into predefined classes. This involves comparing the abstract features of a new image with those in the feature database to determine the most likely class. The final output of this process is a classification result, which provides a label, or a description of the image based on the learned features. This output can be used for various applications, such as identifying fish species in marine research, supporting automated monitoring systems, or enhancing ecological studies.

3.6 **Training Process**

The models were trained over various epochs to ensure adequate learning and to achieve acceptable accuracies while minimising overfitting. To assess the effectiveness of the training, confusion matrices, graphs, and accuracy metrics (such as precision, recall, and the $f1$ score were generated. Based on these evaluations, the most effective model was selected and subsequently fine-tuned for final implementation.

Often, in classification tasks, the available data is not sufficient to develop models that are both accurate and robust. A common solution to this issue is data augmentation (Fawzi et al., 2016). This technique involves creating synthetic data to expand the training dataset, which helps to reduce overfitting. In this study, with only 1,337 training images available, data augmentation was employed to enhance model robustness despite the limited data.

For the ResNet18 model, data augmentation methods such as random resizing and random horizontal flipping were applied, along with normalization, to both training and validation datasets. Similarly, the TF.Keras Sequential model utilised data augmentation techniques to improve dataset quality and model performance. Lastly, the YOLO v8 also integrated automatic data augmentation methods. These practices ensured that all models benefited from augmented data during training, thereby enhancing their ability to generalise and preventing overfitting.

4. RESULTS

This chapter presents the results of the study, providing a comprehensive analysis of the performance of the developed classification models as outlined in the previous methodology chapter. It begins with the presentation of confusion matrices, which illustrate the accuracy of the models in distinguishing between different species. This is followed by a discussion of the error metrics, which quantify the types and frequency of misclassifications. The chapter then evaluates the performance of the models using various metrics, offering a detailed assessment of their effectiveness. An in-depth error analysis is also conducted to identify potential areas for improvement in the models. The training and validation processes are then discussed, outlining how the models were fine-tuned and optimised to achieve their final performance. Finally, the development of a website to showcase and apply the results of the research is detailed, demonstrating how the findings are translated into a practical tool for species identification and promoting citizen science.

4.1 Confusion Matrices

A confusion matrix is a table used to help visualise the performance metrics of a classification model by comparing the actual and predicted labels. These labels were obtained after evaluating all three classification models using the same set of images as can be seen in Table 2. In these matrices, the columns correspond to the actual species depicted in the test images, while the rows represent the classifications made by each model. Confusion matrices were utilised because they help to identify whether a classifier is behaving randomly or facing challenges in distinguishing between specific classes (Marom et al., 2010).

Table 2 - The confusion matrices displaying the classification outcomes for the five selected fish species across three models: TF.Keras, ResNet18, and YOLO v8. Each matrix provides insight into the performance and accuracy of the respective model in correctly identifying and distinguishing between the different species.

		Actual				
Model		*Fistularia commersonii*	*Lobotes surinamensis*	*Pomadasys incisus*	*Siganus luridus*	*Stephanolepis diaspros*
TF.Keras	*Fistularia commersonii*	24	2	8	4	3
	Lobotes surinamensis	4	20	3	1	2
	Pomadasys incisus	3	2	11	1	2
	Siganus luridus	4	2	3	12	5
	Stephanolepis diaspros	3	0	3	5	12
ResNet18	*Fistularia commersonii*	40	0	0	1	0
	Lobotes surinamensis	1	28	0	0	1
	Pomadasys incisus	0	0	18	1	0
	Siganus luridus	1	1	1	21	2
	Stephanolepis diaspros	1	2	0	0	20
YOLO v8*	*Fistularia commersonii*	32	0	0	0	0
	Lobotes surinamensis	2	17	0	4	1
	Pomadasys incisus	1	0	9	1	0
	Siganus luridus	1	0	6	20	0
	Stephanolepis diaspros	1	1	0	0	15

**A select number of images in the test dataset could not be used due to file format restrictions.*

4.2 **Error Metrics**

Accuracy, precision, recall, and $f1\ score$ are essential evaluation metrics used to assess the effectiveness of machine learning models. These metrics were computed using Equations 1 to 4, and they provide a comprehensive understanding of the model performance. True Positive (TP) reflects the number of instances where the model correctly predicts the positive class, in this case, these are the fish species. False Positive (FP) represents the instances where the model incorrectly predicts the positive class when it is negative. True Negative (TN) counts the instances where the model correctly identifies the negative class, while False Negative (FN) denotes the situations where the model fails to recognise the positive class, mistakenly predicting it as negative.

Precision (Equation 1) measures the ratio of TP predictions to the total number of positive predictions, indicating how accurate the positive predictions are. Recall (Equation 2) calculates the ratio of TP predictions to all actual positive instances, reflecting the ability of the model to capture all true positives. The $f1\ score$ (Equation 3) provides a balance between precision and recall by computing their harmonic mean, which is particularly useful in scenarios where a balance between the two metrics is needed, especially when either precision or recall is disproportionately high. Finally, accuracy (Equation 4) evaluates the overall correctness of the model by considering the proportion of true predictions (both positive and negative) across the entire dataset.

Equation 1 - Equation for calculating the precision

$$Precision = \frac{TP}{TP+FP}$$

Equation 2 - Equation for calculating the recall

$$Recall = \frac{TP}{TP+FN}$$

Equation 3 - Equation for calculating the $f1$ score

$$f1\ score = \frac{2\ (Recall \times Precision)}{Recall \times Precision}$$

Equation 4 - Equation for calculating the accuracy

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

4.3 **Model Performance Metrics**

The calculated metrics for all five species across each model can be observed in Table 3, offering a comprehensive comparison of their performance in species identification. This table allows for a detailed assessment of how effectively each model (TF.Keras, ResNet18, and YOLO v8) predicts and distinguishes between the fish species based on precision, recall, $f1$ score, and accuracy.

Table 3 - The precision, recall, $f1\ score$, and accuracy metrics for the classification of five fish species using three models. These metrics offer a thorough evaluation of the capability of each model to accurately identify and differentiate between the species.

Model	Species	Precision	Recall	$f1$ Score	Accuracy
TF.Keras	*Fistularia commersonii*	0.63	0.59	0.61	0.57
	Lobotes surinamensis	0.77	0.67	0.71	
	Pomadasys incisus	0.39	0.58	0.47	
	Siganus luridus	0.52	0.46	0.49	
	Stephanolepis diaspros	0.50	0.52	0.51	
ResNet18	*Fistularia commersonii*	0.98	0.93	0.95	0.91
	Lobotes surinamensis	0.93	0.90	0.92	
	Pomadasys incisus	0.95	0.95	0.94	
	Siganus luridus	0.81	0.91	0.86	
	Stephanolepis diaspros	0.87	0.87	0.87	
YOLO v8	*Fistularia commersonii*	0.86	1.00	0.93	0.84
	Lobotes surinamensis	0.94	0.71	0.81	
	Pomadasys incisus	0.60	0.82	0.69	
	Siganus luridus	0.80	0.74	0.77	
	Stephanolepis diaspros	0.94	0.88	0.91	

In Figure 4.1, the performance of the YOLO model is highlighted through its ability to correctly identify two fish species with a confidence score of 90%. This high level of accuracy and confidence underscores the robustness of this model in handling unseen data, reflecting its capacity to generalise well beyond the training set. This high confidence score further indicates that the YOLO model not only makes accurate predictions but does so with a substantial degree of certainty, which is pivotal for applications that require reliable and precise species identification.

(a)

(b)

Figure 4.1 - Results from the YOLO model, where the model successfully identifies two species in images (a) and (b) with a confidence score of 90%.

Conversely, Figure 4.2 illustrates the performance of the TF.Keras model on unseen test images, revealing both its strengths and limitations. In one instance, the model successfully identifies *Stephanolepis diaspros*, demonstrating its capability to accurately classify this species. However, in another instance, it misidentifies *Fistularia commersonii*, which highlights a key weakness in the performance of the model. This misclassification suggests that TF.Keras may struggle with distinguishing between species that have subtle visual differences, leading to errors in prediction. But in this case, there are clear morphological distinctions, hence there is a high probability that the background where the species is found may be a source of error as it alters the confidence in terms of accurate classification.

(a)

(b)

Figure 4.2 - Unseen test images presented to the TF.Keras model. In this figure, the model correctly identifies *Stephanolepis diaspros* in image (a), however, in image (b), the model incorrectly identifies the species, which is actually *Fistularia commersonii*, highlighting a misclassification error.

4.4 **Error Analysis**

4.4.1 Misclassifications

The confusion matrices provide a clear illustration of where each model struggles with misclassifications, highlighting specific areas of concern. For the TF.Keras model, misclassifications are particularly prevalent in the identification of *Pomadasys incisus*, which has a notably low precision of 0.39 and an $f1\ score$ of 0.47. This indicates a significant number of false positives, where the model incorrectly identifies other species as *Pomadasys incisus*. Similarly, *Siganus luridus* and *Stephanolepis diaspros* also present challenges for TF.Keras, with both species showing relatively low precision and recall, leading to $f1\,score$ below 0.52. Moreover, the ResNet18 model demonstrates some difficulties in distinguishing *Siganus luridus* ($f1\ score$ of 0.86) and *Stephanolepis diaspros* ($f1\ score$ of 0.87), suggesting occasional confusion between these species. Nevertheless, ResNet18 excels at correctly identifying *Fistularia commersonii* and *Pomadasys incisus* with high precision and $f1\ score$s of 0.95 and 0.94, respectively. The YOLO v8 model struggles particularly with *Lobotes surinamensis*, where the recall drops to 0.71, resulting in an $f1\ score$ of 0.81, indicating that the model fails to correctly identify a significant portion of actual instances of this species.

A closer analysis of the error patterns reveals that certain species are consistently misclassified across models, indicating potential areas for improvement. For instance, *Pomadasys incisus* emerges as a challenging species for both TF.Keras and YOLO v8, with low precision and $f1\,scores$, indicating that this species is often confused with the others. This pattern may be due to the morphological similarities between *Pomadasys incisus* and other species, such as body shape and colouration, which are not easily distinguished by the models. Additionally, *Siganus luridus* shows a tendency to be misclassified by both TF.Keras and ResNet18, highlighting challenges in recognising this species, possibly due to its variable appearance in different environments. In contrast, *Fistularia commersonii* is consistently well-identified by ResNet18 and YOLO v8, suggesting that this species has distinctive morphological features that these models capture effectively. Overall, the error patterns indicate that while some species present clear, distinguishable characteristics that models can reliably identify, others with more subtle or variable features are prone to higher misclassification rates (Rahti et al., 2017).

4.4.2 Variability and Reliability of Results

The variability and reliability of the results can be estimated by examining the consistency of the performance metrics across different models and species, as well as considering potential factors that could influence these outcomes. The ResNet18 model demonstrates high reliability, as evidenced by its consistently strong performance across all metrics for most species, with $f1\ score$ ranging from 0.86 to 0.95. This proposes that ResNet18 is a robust model, providing reliable results with minimal variability when identifying different species.

In contrast, the TF.Keras model exhibits greater variability in its results, particularly in its lower precision and $f1\ score$s for species like *P. incisus* ($f1\ score$ of 0.47) and *Siganus luridus* ($f1\ score$ of 0.49). These fluctuations indicate that the predictions of this model are less consistent and more prone to errors, suggesting lower reliability in its classification performance. The variability in the performance of YOLO is also noticeable, particularly with *Lobotes surinamensis*, where the recall drops to 0.71, affecting the overall $f1\ score$. Despite this, YOLO v8 shows relatively high precision for other species, indicating that while some predictions are variable, the model generally provides reliable outcomes.

The reliability of these results is also supported by the substantial confidence scores reported by the YOLO model, particularly in its correct classifications, which were accompanied by confidence levels of 90%. However, the observed misclassifications and the lower performance of TF.Keras indicate that environmental factors such as lighting conditions, image quality, and species similarities may introduce variability into the results. Therefore, while ResNet18 and YOLO v8 demonstrate high reliability with low variability, the performance of TF.Keras infers that further improvements are needed to enhance the consistency and reliability of this model in real-world applications.

4.5 **Model Training and Validation**

In this section, the overall process of training each model and their respective metrics such as validation and training loss, and their accuracies, will be evaluated. Accuracy measures the efficiency of a classification model, typically expressed as a percentage, and represents the proportion of correct predictions matching the actual values, making it easy to interpret. Accuracy is often visualised through graphs to illustrate the performance of the model. In contrast, a loss function quantifies how much the predicted values deviate from the true values, summing the errors for each sample in the validation set.

During model training, the accuracy and loss of validation data can fluctuate under different conditions. Ideally, as training progresses, errors should decrease with each epoch, while accuracy should improve, a total of three scenarios could occur as elaborated in the study of Rum et al., (2021). If loss increases and accuracy decreases, it indicates that the model is not learning effectively. Contrastingly, if both loss and accuracy increase, this may suggest diverse probability values or overfitting, especially when using the softmax function in the output layer. Ultimately, if loss decreases and accuracy increases, it indicates that the model is learning and performing well.

4.5.1 TF.Keras

Starting with the accuracy, Figure 4.3 (left) shows how the accuracy of the model on the training data changes throughout the given epochs, known as the training accuracy. On the other hand, the validation accuracy (Figure 4.3, right) illustrates how the accuracy of the model on the validation data changes over the same provided epoch range.

The training accuracy starts at around 0.3 (30%) and steadily increases, reaching 1.0 (100%) by the 10th epoch. This means that by the end of the 10th epoch, the model is correctly classifying all of the training data. The validation accuracy starts at 0.45

(45%) and increases to about 0.65 (65%) by the 10th epoch. Although there is an improvement, the validation accuracy lags significantly behind the training accuracy.

The training accuracy improving to 100% while validation accuracy only reaches 65% is a classic sign of overfitting. The model is performing extremely well on the training data but is not reciprocated for the new, unseen data (validation set). In a well-adapted model, the gap between training and validation accuracy should be smaller (Vabalas et al., 2019). Despite early stopping, the rapid rise in training accuracy and the slower, more modest rise in validation accuracy suggest that the model is learning to memorise the training data rather than generalise.

Unlike the accuracy graphs, the training loss is a line that represents how the loss, which by definition is a measure of how far off the predictions of the model are from the actual values, on the training data decreases over the epochs (Cui et al., 2019). The training loss starts at around 1.58 and decreases smoothly, reaching about 0.2 by the final epoch. This steady decrease indicates that the model is learning and improving its performance on the training data. In addition, the validation loss starts at 1.45, spikes at the 5th epoch to around 1.4, and then gradually decreases, ending at 1.2.

The spike at the 5th epoch is unusual and could indicate a temporary overfitting to some patterns in the training data that are not present in the validation data (Ying, 2019). Furthermore, some noisy or irrelevant features in the training data could have been the reason for the poorer performance of the validation data during that epoch. After the spike, the validation loss decreases but remains relatively high, indicating that while early stopping and data augmentation facilitated improvement, the model still struggles to be more broadly applicable. The steady decrease in training loss shows the model is grasping the training data effectively. However, the validation loss does not decrease as much, and the fact that the validation loss does not decrease as smoothly as the training loss suggests that the model is overfitting (Ying, 2019).

Overall, mitigation steps were still taken to try to overcome overfitting, one of which was stopping the model early by giving it a shorter epoch range which likely aided, but the gap between training and validation performance suggests that overfitting was already occurring by that point. Another measure implemented was applying data augmentation techniques to the provided images within the dataset to increase the diversity of training data, potentially improving the adaptability.

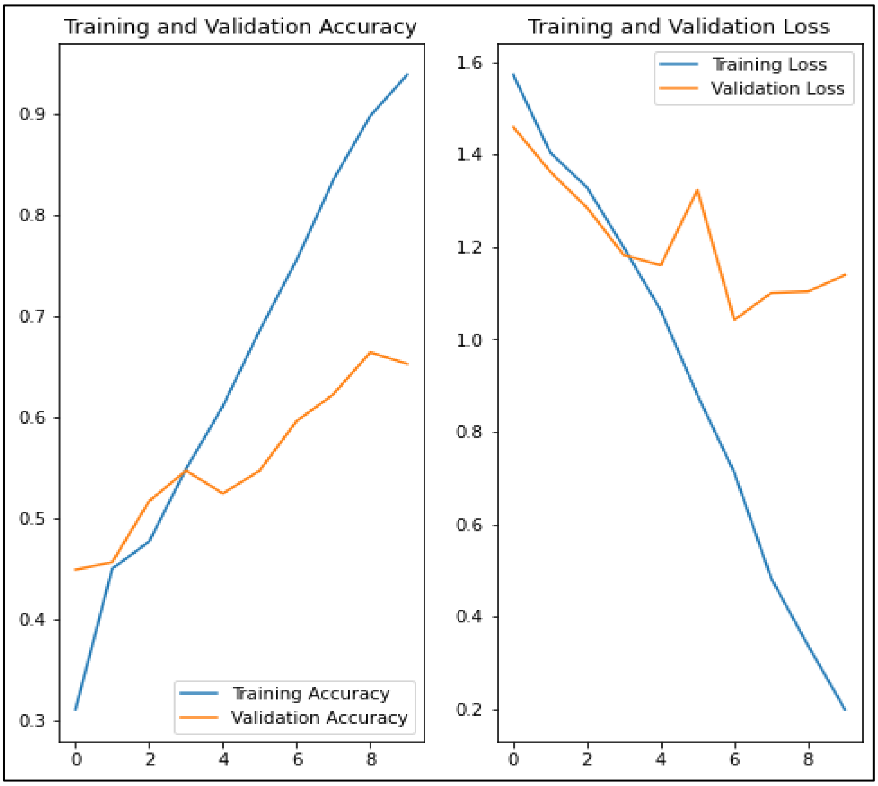

Figure 4.3 - The graphs of the TF.Keras model showcasing the training and validation accuracy (left), and their subsequent losses (right).

4.5.2 ResNet18

Following a similar graph format to that of the TF.Keras model, the training accuracy (Figure 4.4, left) of the ResNet18 model starts at around 0.3 (30%) and steadily increases, reaching 0.89 (89%) by the 50th epoch. This gradual increase suggests that the model is learning well over time, improving its ability to correctly classify the training data. In contrast, the validation accuracy (Figure 4.4, right) starts at 0.53

(53%) and increases to 0.91 (91%) by the last epoch. Notably, the validation accuracy ends higher than the training accuracy, which is unusual and suggests firm standardisation (Zhang et al., 2021).
Both the training and validation accuracies steadily increase, with validation accuracy even surpassing the training accuracy. This is a strong indicator of good generalisation, meaning the model is not only learning the training data but is also effectively applying that learning to unseen data (Moradi et al., 2020). With reference to Figure 4.4 (less), it can be observed the parallel decrease in both training and validation loss, meaning that the model is steadily learning.

Unlike the TF.Keras model, there is no significant gap between training and validation accuracy, and also the absence of any significant divergences between the training and validation loss indicates that the ResNet18 model is not overfitting.

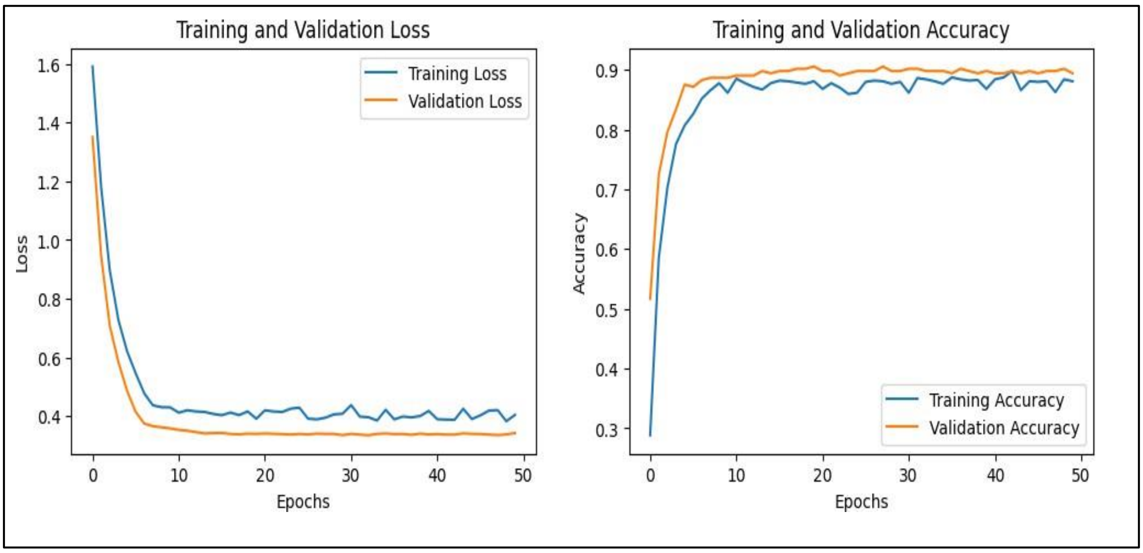

Figure 4.4 - The training and validation accuracy (left), and their subsequent losses (right) for the Resnet18 model.

4.5.3 YOLO

Figure 4.5 shows the loss curves for different components during the training and validation phases of the YOLO v8 model. It is pertinent to point out that both Figures

4.5 and 4.6 are the final run of the YOLO model, hence these are the last 169 epochs, from the total of 500 epochs and above.

The top row contains the training losses, these being the; train/box_loss graph showing the bounding box regression loss during training. It indicates how well the model is learning to predict the correct bounding box coordinates for the fish. The consistent downward trend suggests that the model is improving over time. The train/cls_loss graph shows the classification loss during training, which measures how well the model distinguishes between different classes of fish. The loss decreasing over time is a good sign, showing that the predictions of the model are becoming more accurate. Lastly, the train/dfl_loss graph indicates the distribution focal loss during training, which is related to the confidence scores assigned to each prediction, the decreasing trend indicates better confidence in predictions.

On the other hand, the bottom row consists of the validation losses. The val/box_loss is the bounding box loss on the validation set, and although it fluctuates more than the training loss, it generally decreases, indicating that the model is not overfitting and responding positively to unseen data. Moreover, the val/cls_loss is the classification loss on the validation set, the fluctuations suggest some variability in performance across different validation epochs, which could be due to the complexity of the dataset or inherent variability in the validation samples. Finally, the val/dfl_loss shows the distribution focal loss on the validation set, with similar fluctuations but a general downward trend, suggesting improving confidence in the predictions over time.

The other set of graphs (Figure 4.6) depict key performance metrics for evaluating the ability of YOLO to identify the five fish species. Starting with the metrics/precision(B), the fluctuations here indicate that precision varies from epoch to epoch, but it generally trends upward, suggesting that the model is becoming better at minimising false positives over time. Additionally, the metrics/recall(B) showcase an upward trend indicating that the model is getting better at identifying actual positives, though there is some variability in performance. The metrics/mAP50(B) represents the mean Average Precision (mAP) at an Intersection-over-Union (IoU) threshold of 0.5. This is a comprehensive metric that considers both precision and recall across all classes. The gradual increase suggests an overall improvement in the ability of the model to detect and classify the fish species. Ultimately, the metrics/mAP50-95(B) are similar to mAP50 but averaged over IoU thresholds from 0.5 to 0.95. The increasing trend, though with some fluctuations, indicates that the performance of the model is improving across a broader range of IoU thresholds.

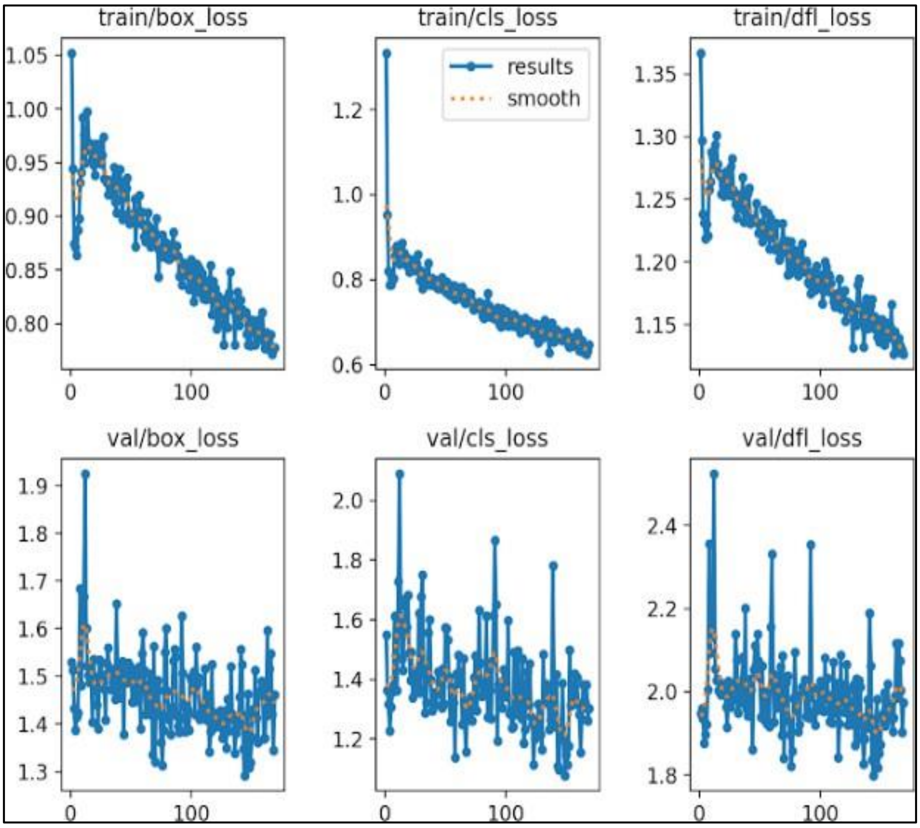

Figure 4.5 - The training and validation losses of the YOLO model.

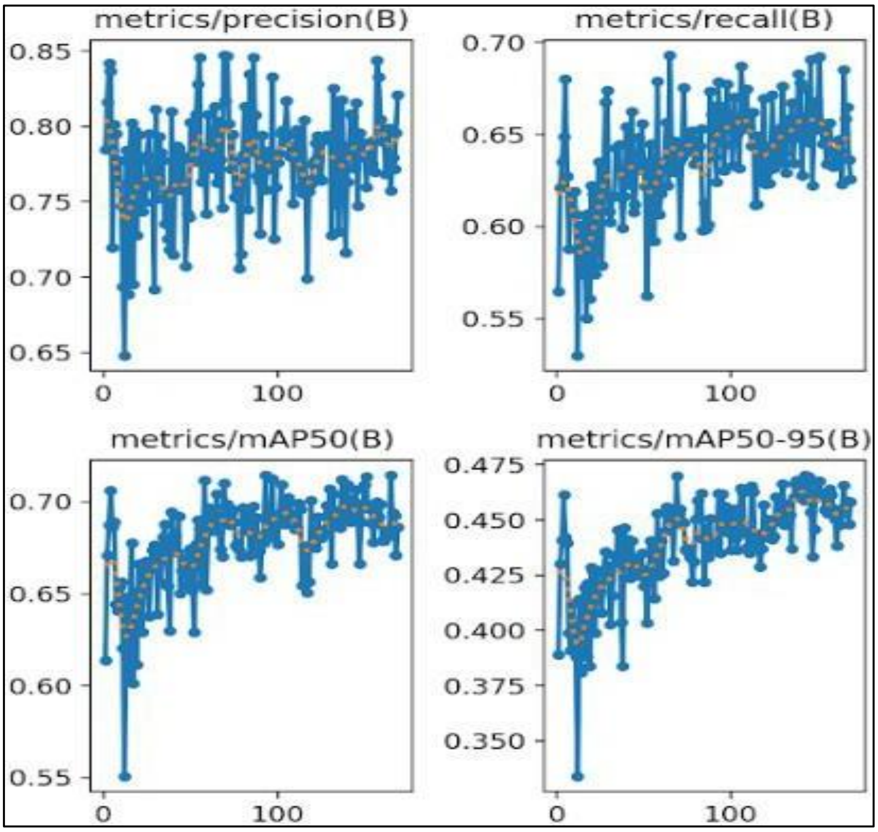

Figure 4.6 - The metrics for precision and recall, alongside the mAP50 and mAP50-95 for the YOLO model.

4.5.3.1 The Precision-Recall (PR) Curve

The precision measures the accuracy of the positive predictions made by the model. It is defined as the number of true positives divided by the number of true positives plus false positives. A higher precision means that when the model predicts a species, it is more likely to be correct. In contrast, the recall, as discussed before, measures how well the model identifies all relevant instances of a class. In addition, the mAP@0.5 (mean Average Precision at IoU threshold 0.5) is the average precision score across all classes at an IoU threshold of 0.5. It provides a summary measure of model performance across different classes. The average mAP@0.5 of 0.710 indicates the overall performance of the model in detecting and classifying fish species.

When analysing the PR graph produced by the YOLO model (Figure 4.7), the five species can be seen to experience different values when compared to those of the confusion matrix. This could be due to either the precision and recall values being calculated at a fixed point, or the use of mAP could have varied the results slightly. As mentioned previously, the produced graph is over the final 169 epochs of the model from the total, hence it could have given some variance to the results.

The *Fistularia commersonii* had a precision of 0.677, which was slightly below the average mAP of 0.710, it indicates that while the model has a good recall (detecting all instances), its precision is somewhat lower. This means that there are more false positives relative to true positives for this species. The model might be making more incorrect positive identifications for *Fistularia commersonii* compared to the average. For *Lobotes surinamensis*¸ the precision value was of 0.788 which is above the average mAP of 0.710. It suggests that the model is relatively good at accurately identifying *Lobotes surinamensis* when it makes a prediction. There are fewer false positives compared to true positives, indicating better precision for this species. Moreover, the *Pomadasys incisus* scored a precision of 0.592 which was also below the average mAP. It implies that the predictions of YOLO for *Pomadasys incisus* are less accurate, with more false positives relative to true positives. The model might be struggling to correctly identify this species or may be making more incorrect predictions. Furthermore, *Siganus luridus*, similar to the score of *Pomadasys incisus,* had a precision score (0.595) that ranked below the average. It conveys that the precision of the model for *Siganus luridus* is also lower than the average threshold. This articulates that the model might be making more false positive predictions for this species as well. Lastly, the precision score (0.897) for *Stephanolepis diaspros* was significantly above the average mAP of 0.710. It indicates that the model is very good at accurately predicting *Stephanolepis diaspros* when it makes a positive identification. The high precision score suggests fewer false positives and a high level of confidence in its predictions for this species.

Some highlights from this graph include that *Stephanolepis diaspros* stands out with the highest precision score, hence consisting of more accurate positive predictions. Moreover, *Lobotes surinamensis* also shows above-average precision, nonetheless not as high as *Stephanolepis diaspros*. Additionally, *Pomadasys incisus* and *Siganus luridus* have lower precision scores, which may emphasise the challenges in accurately identifying these species with false positives. The fact that all recall values end at 1 demonstrates that the model detects all instances of the species, but precision varies significantly. This denotes that while the model is notable at detecting instances (recall), its ability to accurately predict the species (precision) varies among different the five fish species.

The model performs very well in detecting and accurately predicting *Stephanolepis diaspros* and *Lobotes surinamensis*, while it struggles with *Pomadasys incisus* and *Siganus luridus* in terms of precision. The high recall of YOLO across all species implies robust detection capabilities, but varying precision indicates that there is room for improvement in reducing false positives for certain species (Muksit et al., 2022).

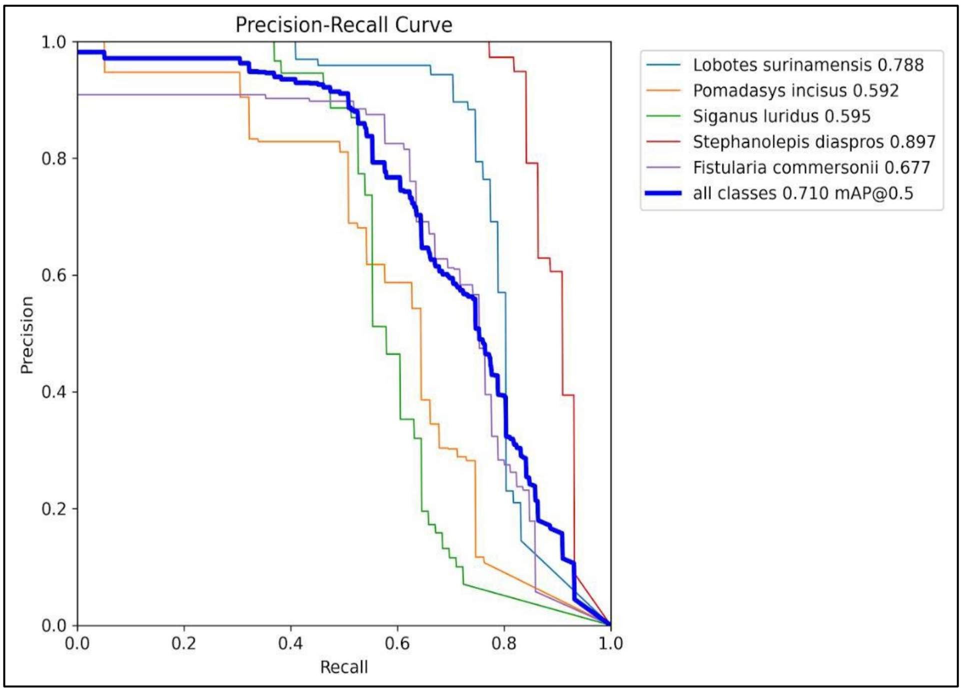

Figure 4.7 - The Precision-Recall curve for the YOLO model for the 5 fish species, alongside their coherent values at mAP IoU 0.5.

4.5.3.2 The $f1$ Confidence Curve

The $f1$ curve showcases the $f1\ score$ of the model as a function of prediction confidence for each class. In resemblance to the PR curve values, the $f1\ score$ also differs from the primary ones found in the confusion matrix (Table 3) but still withholds the same trends. This score is usually a single value that represents the performance of the model at a specific threshold of confidence. In contrast, the $f1$-Confidence curve (Figure 4.8) shows how the $f1\ score$ changes across all possible confidence thresholds (Wu & Li, 2024). As observed, the *S. diaspros* and the *L. surinamensis* were the only two species that surpassed the 0.73 threshold, whilst the other three species were found below the average (<0.73).

The peak of each curve represents the best possible $f1\ score$ that could be achieved for that class at an optimal confidence threshold. As confidence increases, the model tends to make fewer predictions, and those predictions are more likely to be correct, which initially increases the $f1\ score$. However, after a certain point, the number of predictions may become too low, which can cause the $f1\ score$ to drop. If the $f1\ score$ remains high over a wide range of confidence levels, it indicates that the confidence of the model scores is well-calibrated. Furthermore, the curve helps in selecting an optimal confidence threshold for making predictions in real-world applications. Moreover, it provides insights into the performance differences across classes. A prime example is the *F. commersonii*, this class had a significantly lower $f1\ score$ across confidence levels, indicating that the model struggles much more with this species compared to the others.

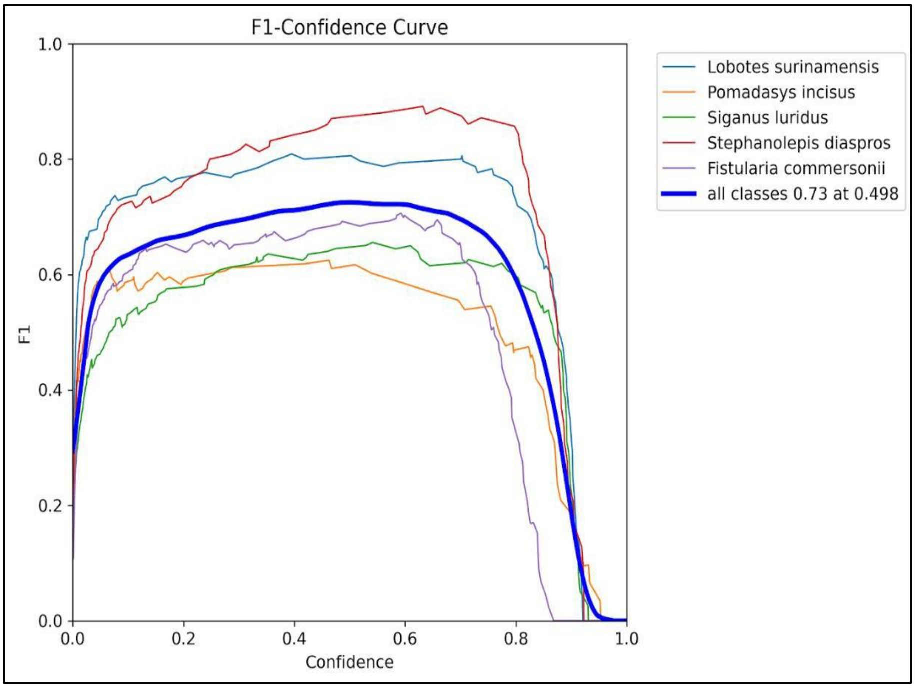

Figure 4.8 - The $f1$-Confidence curve of the YOLO model for the five fish species.

4.6 **Website Development**

The development of a unified website is underway to support the ResNet18 model, which has demonstrated the best performance among the three models evaluated. With the use of the Gradio package (Figure 4.9), the website will allow users to upload images of unclassified fish. The ResNet18 model will analyse these images and provide its highest-confidence prediction regarding the fish species.

The initial phase involves creating an online platform that integrates the ResNet18 model for use by citizen scientists. This platform will feature an interactive map and enable users to report new sightings of marine species. The long-term objective is to expand the database to include a variety of marine invasive species that were not targeted in this study, such as crabs, sea hares, and algae, across different taxa. This initiative aims to enhance the monitoring of marine biodiversity, track invasive species, and increase awareness through citizen science.

Users will be able to submit their observations, with the algorithm identifying the species and providing a brief description alongside the submitted photo. Each submission will record the location of the sighting, and after review, the image will be added to the database to continuously train and improve the algorithm.

Initially, the interactive map will focus on the Maltese Islands, engaging local citizens through the 'Spot the Alien Fish' citizen science campaign. The goal is to build a strong foundation of local participation and data collection in Malta, which will be pivotal in scaling the project to cover the entire Mediterranean region in subsequent stages.

This approach aims to create a comprehensive and dynamic tool that supports species identification and contributes to broader marine conservation and invasive species management efforts. The integration of citizen science will be crucial, enabling individuals to contribute to scientific research and conservation.

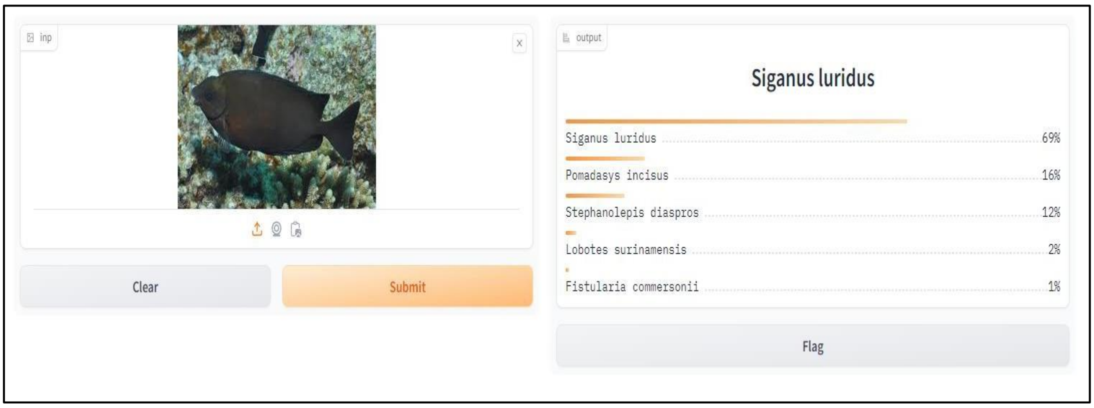

Figure 4.9 - The initial website development using a package known as Gradio.

5. DISCUSSION

This chapter provides a detailed discussion of the results obtained, offering a critical interpretation in the context of the broader research objectives. The discussion begins by analysing the confusion matrices and interpreting the patterns of correct and incorrect classifications in the model outputs. Following this, the analysis focuses on areas of poor performance and identifies the limitations of the study, exploring factors that may have contributed to suboptimal results. Moreover, it then compares the findings with existing literature, highlighting similarities and differences with prior research in the field. In addition, key challenges in fish species identification are addressed, along with improvements for tackling these issues. Finally, the discussion concludes by addressing the practical implications of the study, particularly in terms of strategies to reduce the spread of invasive IAS, and how the research can contribute to ongoing efforts in this area.

5.1 Discussion of Confusion Matrices Results

It is worth noting that a research gap is present when concerning four of the five fish species studied here (*Fistularia commersonii*, *Lobotes surinamensis*, *Pomadasys incisus*, and *Stephanolepis diaspros*), as none have been previously included in image classification studies. This gap underscores a clear opportunity for further research into the application of image classification techniques for these invasive alien species.

A recent study by Fleuré et al. (2024) focused on automated identification protocols for two invasive rabbitfish species in the Mediterranean, including *Siganus luridus*. They compiled a dataset of 31,285 images featuring *Siganus spp.* and six native Mediterranean fish species, collected from 40 underwater videos filmed across three reef habitats. The researchers utilised a deep learning algorithm based on ResNet50 to train the model for identifying *Siganus spp.* within this dataset of eight species.

The results indicated that the model achieved a recall of 0.92 for the *Siganus* genus, which improved to 0.98 after applying confidence-based post-processing. In comparison, the models used in our study reported slightly lower recall values: 0.91 (ResNet18), 0.74 (YOLO), and 0.46 (TF.Keras) as shown in Table 3. For precision, Fleuré et al. (2024) reported a precision of 0.61, while our models achieved precision values of 0.80 (YOLO) and 0.81 (ResNet18), with the TF.Keras model recording a considerably lower precision of 0.52.

The accuracy values presented in Table 3 show that the ResNet and YOLO models performed significantly well, achieving accuracies of 0.91 and 0.84, respectively. This indicates that these models generally provide robust and consistent automated fish identification. In contrast, the TF.Keras sequential model underperformed, with an accuracy of only 0.61.

Regarding precision, the fish species that were classified with the highest accuracy include *Lobotes surinamensis*, with two out of three models reporting precision values exceeding 0.93, and *Fistularia commersonii*, with two out of three models achieving precision values greater than 0.86, including an outstanding precision of 0.98 from the ResNet18 model.

Interestingly, the ResNet18 model achieved the highest precision of 0.95 for *Pomadasys incisus*, whereas the other models recorded lower precision values of 0.60 and 0.39, making *Pomadasys incisus* the most challenging species to identify with those models. Both the YOLO and ResNet models demonstrated precision values exceeding 0.80 for *Stephanolepis diaspros*, with YOLO attaining an impressive precision of 0.94 for this species.

Additionally, high recall and $f1$ score metrics were observed for two of the three classification models, further confirming their effectiveness in accurately identifying fish species from the provided images. The $f1$ score in this study ranged from 0.47 to 0.95, with the lowest scores for all species coming from the TF.Keras model, while the other two models scored between 0.69 and 0.95. Notably, the lowest $f1$ score of the ResNet model was 0.86 for *S. luridus* and 0.87 for *S. diaspros*, both of which are relatively high compared to the scores from the other models.

5.2 Analysis of Poor Performance Metrics and Study Limitations

Upon reviewing the outcomes of this study, it is crucial to acknowledge its limitations. The variability in the results of the models that yielded poor performance in some cases can be attributed to several factors. These issues can arise from both model-specific limitations and data-related challenges.

The research study by Villon et al. (2021) elaborates on how underwater image classification presents several challenges. One significant drawback of deep learning algorithm classifiers is their limitation to identifying fish species only in images where a single individual is centred. As a result, raw underwater videos or images must be cropped to isolate individual fish from background noise, such as surrounding habitat features or other species. Additionally, the effectiveness of these classifiers depends heavily on having a sufficient number of training images for each target species. This requirement often means gathering thousands of images per species, which can lead to reduced recall and precision values.

Similar ideas arise when considering variations in environmental conditions. These factors heavily impact the clarity and readability of underwater images, light intensity alterations, algae, background coral reef patterns, and image quality, are a few that affect the ability of the model to correctly identify the target species from the image (Salman, 2016). Furthermore, if the training dataset contains an unequal distribution of species, with some classes being underrepresented, models may not learn to accurately classify them (Huang et al., 2016). Moreover, models may become biased towards more frequent classes, leading to higher misclassification rates for less common species (Schaaf et al., 2021). This bias can result in poor performance, as observed with TF.Keras for *Pomadasys incisus* and *Siganus luridus*.

The architecture of simpler models like TF.Keras may not be sophisticated enough to extract and differentiate complex features, especially when compared to deeper models like ResNet18 and YOLO, resulting in lower performance for species that require more multifaceted feature extraction.

When a model overfits to the training data, it may perform effectively on the training set but poorly on unseen data, particularly for species that were underrepresented or not well captured in the training data (Ying, 2019). Overfitting can lead the model to over-specialise on the given training dataset, thus making it struggle to handle unseen images that vary slightly from the ones already seen.

Species with similar morphological characteristics are subject to misinterpretation by the model. Good examples of this are *Siganus luridus* and *Lobotes surinamensis,* as these species share visual similarities, especially when they are young (darker colourations), hence the models might struggle to differentiate between them, leading to lower precision and recall. This issue could be aggravated by the models' inability to capture subtle differences in features that are critical for correct classification (Rauf et al., 2019). High variability within a species, such as differences in size, shape, or colour, can make it more difficult for models to learn a consistent pattern for classification (Mittal et al., 2022). Contrastingly, when species are not only similar to each other but also exhibit high variability within the same class, this can make the models susceptible to misunderstanding, resulting in frequent misclassifications of the species (Siddiqui et al., 2018).

5.3 Comparison with Previous Studies

Rum et al. (2021) showcased the creation of FishDeTec, a mobile app designed to identify freshwater fish in Malaysia, utilising the VGG16 Convolutional Neural Network model. They applied this model to a dataset featuring 178 images of eight distinct freshwater species found in Malaysia. To address potential overfitting due to image variability, they implemented an augmentation process involving transformations like zoom, rotation, and flipping. The model achieved an accuracy ranging from 60% to 80% across the eight fish species. However, when compared to the models established in this current research, both the ResNet18 (91% accuracy) and YOLO v8 (84% accuracy) models outperformed the FishDeTec model quite significantly.

Another research study conducted by Jalal et al., (2020) was about developing an effective and automatic system for detecting and classifying fish species in underwater videos, addressing the various challenges posed by underwater environments, such as the ones mentioned above by Salman (2016), that result in poor performance metrics. Their proposed solution combines optical flow and Gaussian mixture models with the YOLO deep neural network to improve fish detection and classification, particularly in scenarios where fish are freely moving or camouflaged. This hybrid model was evaluated on two underwater video datasets: the LifeCLEF 2015 benchmark from the Fish4Knowledge repository and a dataset from the University of Western Australia. The model achieved fish detection F-scores of 95.47% and 91.2% and fish species classification accuracies of 91.64% and 79.8% on these datasets, respectively.

The study also conducted a comparative analysis against other deep learning models, such as AlexNet and VGGNet, using various functional parameters like the number of convolutional and fully connected layers, iterations, batch size, and the inclusion of dropout layers. The proposed model, although having fewer layers and therefore less computational complexity, achieved a testing accuracy of 90.48%, outperforming AlexNet, which achieved 86.65%. However, the model did not surpass the validation accuracy of VGGNet, which benefits from being a much larger and pre-trained model. It was concluded that the proposed automatic fish species classification system which was based on a simplified version of AlexNet, was effective for classifying freshwater fish species with less computational power and fewer training images. Although the model does not outperform VGGNet in validation accuracy, it demonstrates superior performance over AlexNet in testing accuracy. The inclusion of a dropout layer before the softmax classifier significantly improved the accuracy of the model, a plausible feature that should be added in future works to try to improve the accuracy of image classifiers for underwater imagery.

In a similar study, Raufa et al., (2019) wanted to address the challenges associated with morphological-based fish species identification, which is often flawed and timeconsuming due to the close resemblance between various fish species, a challenge that was mentioned before for this current research.

The paper proposed a deep learning framework based on CNN architecture to improve the accuracy of fish species identification. Specifically, the authors developed a 32-layer CNN model, designed to extract valuable and discriminating features from fish images, thus enhancing classification performance. The proposed model builds upon the VGGNet architecture, following the approach taken by Jalal et al. (2020), by adding four convolutional layers to each level in the network, introducing deep supervision to boost its classification capabilities.

The study made use of a dataset named Fish-Pak, consisting of 915 images across six fish species, with three different image views (head region, body shape, and scale). This dataset was used to test the effectiveness of the proposed CNN architecture. The performance of the 32-layer CNN model was compared with other well-known deep learning frameworks, including VGG- 16 (used for transfer learning), various configurations of VGG (one block, two blocks, three blocks), LeNet-5, AlexNet, GoogleNet, and ResNet-50.

The results demonstrate that the proposed 32-layer CNN architecture significantly outperforms these existing models. Specifically, the model achieved an accuracy of 96.94% on body view images with a learning rate of 0.001 and momentum of 0.9 over 350 iterations, surpassing the ResNet-50, which was the second-best performing model with an accuracy of 90.76% under similar conditions.

5.3.1 Addressing Key Challenges in Fish Species Identification

The results of this study reflect the persistent challenges in fish species classification, which were also identified in the reviewed literature. These challenges play a crucial role in influencing the overall accuracy and effectiveness of the classification algorithms. Specifically, the complexities of computer vision tasks in the context of underwater environments, such as the automatic recognition of fish within video sequences, are amplified by environmental variability, visual distortions, and image noise. These factors, as seen in the previous chapter, negatively impacted the effectiveness of the machine learning models by degrading the quality of the features extracted for image representation. To overcome these challenges, research has shown that non-linear automatic learning systems, such as multi-layer neural networks like CNNs, are effective. As explained by Salman et al. (2016), non-linear automatic learning systems, such as multi-layer neural networks like CNNs, are used. In these networks, each non-linear hidden layer serves as the input for the next, increasing the complexity of the network and enabling it to effectively encode the non-linearity of the input data. As a result, the network and its parameters are designed to match the non-linear nature of the data, and the supervised training process ensures the automatic learning of complex, non-linear, and discriminative features. With reference to the results of this study, some simpler modifications to enhance performance metrics should be considered. These include using larger training datasets with more images preferably even higher resolution ones in comparison to the lower quality ones used for this study, and incorporating more species not just fish but other marine genus, could heavily increase robustness.

The following points are several key studies that propose innovative techniques to overcome hurdles related to poor accuracy scores. From deep learning models and traditional machine learning approaches to texture analysis and shape recognition methods, these studies have proposed various methods to address the challenges in fish species identification, some of which include:

- Spampianto et al. (2010) attempted to classify fish using texture features extracted from Gabor filtering, grey-level histograms, and Fourier descriptors for shape features, achieving a 92% accuracy on a dataset of 360 images.
- Takakazu Ishimatsu et al. (1998) used speckle patterns and scale forms of fish, employing morphological algorithms and filters to differentiate between three species: Pilchard Sardine, Japanese Horse Mackerel, and Common Mackerel, with respective accuracy rates of 90%, 88%, and 90%.

- Rova et al. (2007) explored the method of warping images before classification, using Support Vector Machine (SVM) techniques on a dataset of 320 images, resulting in a 90% accuracy.
- Pornpanomchai et al. (2013) developed a fish recognition system based on shape and texture, comparing Artificial Neural Networks (ANN) and the Euclidean Distance Method (EDM). Their system, tested on 300 images with a training set of 600 images, achieved accuracies of 81.67% with EDM and 99.00% with ANN. Rodrigues et al. (2015) recommended an algorithm based on SIFT feature extraction combined with Principal Component Analysis, though their study was conducted on a relatively small dataset of 162 images covering six species, achieving a 92% accuracy.
- Hernandez-Serna et al. (2014) used ANNs for the automatic identification of fish species, working with a dataset of 697 images and achieving an accuracy of 91.65%.

The following studies indicate that while SVM techniques can be useful, they may not always provide the highest performance for fish species classification, particularly in complex or noisy environments. Future research should explore alternative methods or hybrid models to overcome these limitations.

- Joo et al. (2013) focused on extracting stripe and colour patterns of wild cichlids and used Random Forests and SVM for classification. Despite their efforts, they achieved a low accuracy of 72% on a dataset of 594 wild cichlids, indicating that the method may not be suitable for all fish species.
- Ogunlana et al. (2015) classified fish based on shape features using SVM, working with a training dataset of 76 fish and a testing dataset of 74 fish, resulting in a modest accuracy of 78.59%.
- Not SVM related, but Cadieux et al. (2000) developed a method that generates fish contours in an unconstrained environment using an infrared silhouette sensor. However, this system struggled with noisy inputs, leading to a relatively low classification accuracy of around 78%.

5.4 Way Forward to Reduce the Spread of IAS

To effectively stop the spread of invasive alien species, it is crucial to enhance our understanding of their impacts on ecosystems and biodiversity, which is currently limited and often based on weak evidence. The complexity of species interactions, along with the varied effects of invasive species in different regions, makes environmental management decisions difficult and sometimes controversial. A major obstacle is the lack of precise knowledge about the life history traits and invasive strategies of these species, which is essential for understanding their roles in ecosystems and their impacts on ecosystem services and biodiversity (Katsanevakis et al., 2014).

Conducting thorough cost-benefit analyses of biological invasions and mitigation measures is a cornerstone of invasion economics, but these analyses require high-quality information on all effects related to species introduction (Iacona, 2014). The impact of a species can differ across regions, as demonstrated by the comb jelly *Mnemiopsis leidyi*, which had significant effects in some areas, such as reducing zooplankton biomass on the Swedish west coast, but not in others like the Black Sea (Bilio & Niermann, 2004). Additionally, the impacts of invasive species can vary over time due to the dynamic interactions between the alien species and the recipient ecosystem, and there may be significant time lags between a species' introduction and its observed impacts (David et al., 2017).

To improve prevention and mitigation efforts, it is necessary to quantify and map the impacts of invasive species and to understand how human activities facilitate these invasions. Current constraints include the lack of comprehensive natural and socioeconomic data, gaps in assessments of marine ecosystem services, and the inherent complexity of the problem (Liquete et al., 2013). Recommended next steps include improving methods for assessing the impacts of alien species, developing suitable indicators, enhancing mapping of species distribution and abundance, and shifting from approaches that offer only weak evidence to those that provide strong evidence and estimates of impact magnitude (Thiele et al., 2010). All of the above actions will greatly assist managers and policymakers in making informed decisions to prevent and mitigate the spread of alien invasive species, especially for the local Maltese context.

To effectively prevent and manage the spread of invasive alien species in marine environments, coordinated global and regional efforts are essential. While national measures can temporarily reduce the introduction of these species, the problem often escalates as they establish themselves in neighbouring regions and trading ports. This challenge is exacerbated by the limitations of current methods for treating ballast water and preventing hull fouling, which are either marginally effective or too costly to implement on a large scale (Karanasiou, 2018).

Once alien species establish themselves, they can adapt to new environments, increasing the risk of invasion in other regions. As elaborated in the study of Bax et al. (2003), the North Pacific seastar, originally suited to the northern hemisphere, has adapted to the southern seasons of the hemispheres, raising the threat of its spread to ports in Australia, New Zealand, South Africa, and South America.

Global and regional strategies are crucial because uncoordinated national efforts can complicate international trade. As global trade and marine traffic continue to grow, along with the development of new trade routes, it is vital to balance the benefits of trade liberalisation with the need to minimise the risks of biological invasions (Burgiel et al., 2006).

Effective management of marine invasive species should involve three key strategies: prioritising prevention over control due to its cost-effectiveness, addressing all significant pathways of introduction, and recognising that while complete prevention is unlikely, the focus should be on minimising risks and managing occurrences (Angell, 2023). Policy frameworks should include six key intervention points: prevention, detection, quarantine, eradication, control, and mitigation (García-Díaz et al., 2021; Angell, 2023).

Progress is being made in improving ballast water treatment methods, such as heat treatment, chemical treatment, ultrafiltration, and ultraviolet light, but each has its limitations. Current practices, like ballast water exchange at sea, aim for a 95% exchange rate but are not always reliable (Endresen et al., 2004). Therefore, new treatment techniques and a risk management framework to identify and monitor high-risk vessels are needed.

A comparative risk management system is also necessary to evaluate the costs and benefits of different management options. This system would consider the environmental, social, and economic impacts of invasive species, and highlight the need for early warning systems and regional responses, supported by an international monitoring network (Magaletti et al., 2018).

Despite efforts to manage vectors that carry alien marine species, it is impossible to completely prevent their entry into new areas. Continuous monitoring is essential for detecting new invasions early, allowing for rapid response and control efforts.

Port surveys and continuous monitoring also provide opportunities for quarantine and local eradication efforts. Effective quarantine can facilitate eradication, but achieving it can be challenging (Giakoumi et al., 2019). Reducing marine traffic between invaded and non-invaded areas can help minimise further impacts, as demonstrated by the spread of the North Pacific seastar from the Derwent estuary to Port Phillip Bay in Australia.

For the local Maltese context, not many studies have been done to prevent or mitigate the impact of marine IAS, but one collective document labelled as "Monitoring Marine Invasive Species in Mediterranean Marine Protected Areas (MPAs), A strategy and practical guide for managers" takes into consideration multiple non-indigenous and invasive alien species. In this document, Otero et al. (2013) discuss most (not all) of the alien invasive marine species that occur within the Mediterranean, their introduction pathways, ecological and economic impacts, reproduction patterns, and how to identify them. Some examples of the species present in this document include the *Fistularia commersonii, Caulerpa taxifolia, Siganus luridus, Pomadasys incisus*, and more.

If invasive species are detected soon after their arrival, eradication may be possible. However, if this window is missed, as with *Caulerpa taxifolia* in the Mediterranean, the species can become widely distributed, and proven methods to mitigate or eradicate them are limited (Ruesink & Collado-Vides, 2006). This is a clear opportunity to take before the lionfish (*Pterois miles*) is first spotted within Maltese waters, a highly invasive fish with no natural predators. According to the study of Bottacini et al. (2024), there has been no official sighting (most probably misclassification) of the lionfish yet in Malta, hence it is crucial to implement preventive measures before it establishes itself. While biological and genetic control techniques offer some promise, they are still in the early stages of development and require careful evaluation and international oversight due to their potential widespread impacts.

6. CONCLUSION

This study aimed to address the gap in research concerning image classification of invasive alien fish species in the Mediterranean, specifically the Maltese islands, focusing on five specific species: *Fistularia commersonii*, *Lobotes surinamensis*, *Pomadasys incisus*, *Stephanolepis diaspros*, and *Siganus luridus*. The analysis of the confusion matrices for the three models (ResNet18, YOLO, and TF.Keras) revealed varying levels of performance across species. Notably, ResNet18 consistently performed well, achieving high precision and recall values, particularly for *Lobotes surinamensis* and *Fistularia commersonii*, while YOLO v8 demonstrated slightly lower but still competitive results. This directly responds to the aim and the first research question of this study, showcasing how algorithms can be tailored for accurate identification despite the environmental challenges. Contrastingly, TF.Keras exhibited the lowest overall performance, struggling particularly with species like *Pomadasys incisus* and *Siganus luridus*. The superior performance of ResNet18, with an accuracy score of 91% and precision values up to 0.98, marks it as the most effective model. These results have led to further development of ResNet18 and enhancements to the database, laying the groundwork for the second and third objectives: integrating the algorithm into a user-friendly website for by-catch identification and establishing a foundational database documenting invasive species findings. Although the database is evolving, it provides a platform for continuous updates, contributing to regional monitoring efforts and early warning systems, fulfilling the final objective of proactive responses to extreme invasions and protecting local marine ecosystems. Notably, this research study has already been published in a special issue journal of MDPI focused on "Machine Learning and Artificial Intelligence with Applications" (see 'Appendix A' for more details).

The research also highlighted key challenges and limitations, as identified in the second research question, that could have impacted the outcome of the study. These include the dataset quality and size, which was not only limited but also featured low image resolutions. This likely contributed to the lower performance metrics of the

TF.Keras model, and could have introduced biases in model training and evaluation, furthermore, the simpler architecture of the TF.Keras model was insufficient for accurately classifying species with complex morphological features. Moreover, the underwater environment introduces a range of challenges, such as varying light conditions, background noise, and the presence of various species in a single image. These factors were not fully accounted for in the models, potentially leading to misclassification. Undoubtedly, the risk of overfitting was a concern due to the limited and potentially unbalanced dataset. Lastly, the models struggled with species that share similar morphological characteristics, such as *Siganus luridus* and *Stephanolepis diaspros*, highlighting the difficulty in distinguishing species with overlapping visual features.

While this study primarily focused on algorithm development, it also laid foundational work for integrating these algorithms into a citizen science platform, supporting the third and final research question. Future studies should focus on creating larger and more diverse datasets, including high-resolution images and a broader range of species. This would help improve the robustness of models and reduce biases caused by underrepresented species. There should also be integration of advanced techniques such as optical flow and Gaussian mixture models (Jalal et al., 2020). This could enhance model performance in detecting and classifying fish species in challenging underwater environments. Additionally, the inclusion of dropout layers before the softmax classifier (Jalal et al., 2020) could reduce overfitting and improve accuracy. Beyond classification, there is a need for research focused on the prevention and mitigation of invasive species. Studies could explore the development of real-time monitoring systems and the implementation of early warning systems (Magaletti et al., 2018) that use image classification technologies to detect and respond to new invasions before it is too late.

REFERENCES

Akyol, O., Ceyhan, T., Özgül, A., & Ertosluk, O. (2018). Maximum size of reticulated leatherjacket, Stephanolepis diaspros Fraser-Brunner, 1940 (Tetraodontiformes: Monacanthidae), for the Turkish Seas. *Mar. Sci*, *14*, 463-480.

Al Muksit, A., Hasan, F., Emon, M. F. H. B., Haque, M. R., Anwary, A. R., & Shatabda, S. (2022). YOLO- Fish: A robust fish detection model to detect fish in realistic underwater environment. *Ecological Informatics*, *72*, 101847.

AL-Nahdi, A. (2021). Biological Characteristics, Population Dynamics and Fisheries Management of Pomadasys commersonnii (Lacepède, 1802) in the Arabian Sea Coast of Oman. *The Arabian Seas: Biodiversity, Environmental Challenges and Conservation Measures*, 779-828.

Alotaibi, B., & Alotaibi, M. (2020). A Hybrid Deep ResNet and Inception Model for Hyperspectral Image Classification. *PFG – Journal of Photogrammetry, Remote Sensing and Geoinformation Science*, *88*(6), 463–476. https://doi.org/10.1007/s41064-020-00124-x

An, D., Hao, J., Wei, Y., Wang, Y., & Yu, X. (2021). Application of computer vision in fish intelligent feeding system—A review. *Aquaculture Research*, *52*(2), 423-437.

Ando, R. K., & Zhang, T. (2005). A framework for learning predictive structures from multiple tasks and unlabelled data. In Journal of Machine Learning Research, 6(Nov), 1817-1853.

Angell, N. R. (2023). *Cost-effectiveness of aquatic invasive species prevention techniques* (Master's thesis, University of Minnesota).

Arndt, E., Robinson, A., Hester, S., Woodham, B., Wilkinson, P., & Gorgula, S. (2021). Factors that influence vessel biofouling and its prevention and management. *Final report for CEBRA Project*, *190803*.

Azzurro, E., La Mesa, G., & Fanelli, E. (2013). The rocky-reef fish assemblages of Malta and Lampedusa islands (Strait of Sicily, Mediterranean Sea): a visual census study in a changing biogeographical sector. *Journal of the Marine Biological Association of the United Kingdom*, *93*(8), 2015-2026.

Azzurro, E., Smeraldo, S., & D'Amen, M. (2022). Spatio-temporal dynamics of exotic fish species in the Mediterranean Sea: Over a century of invasion reconstructed. *Global Change Biology*, *28*(21), 6268-6279.

Azzurro, E., Soto, S., Garofalo, G., & Maynou, F. (2012). Fistularia commersonii in the Mediterranean Sea: invasion history and distribution modelling based on presence-only records. *Biological Invasions*, *15*(5), 977–990. https://doi.org/10.1007/s10530-012-0344-4

Bai, X., Yang, X., & Latecki, L. J. (2008). Detection and recognition of contour parts based on shape similarity. *Pattern Recognition*, *41*(7), 2189-2199.

Barbedo, J. G. A. (2022). A Review on the Use of Computer Vision and Artificial Intelligence for Fish Recognition, Monitoring, and Management. *Fishes*, *7*(6), 335. https://doi.org/10.3390/fishes7060335

Bariche, M., Letourneur, Y., & Harmelin-Vivien, M. (2004). Temporal fluctuations and settlement patterns of native and Lessepsian herbivorous fishes on the Lebanese coast (eastern Mediterranean). *Environmental Biology of fishes*, *70*, 81-90.

Bax, N., Williamson, A., Aguero, M., Gonzalez, E., & Geeves, W. (2003). Marine invasive alien species: a threat to global biodiversity. *Marine policy*, *27*(4), 313-323.

Bilio, M., & Niermann, U. (2004). Is the comb jelly really to blame for it all? Mnemiopsis leidyi and the ecological concerns about the Caspian Sea. *Marine Ecology Progress Series*, *269*, 173-183.

Bodilis, P., Crocetta, F., Langeneck, J., & Francour, P. (2013). The spread of an Atlantic fish species,Pomadasys incisus(Bowdich, 1825) (Osteichthyes: Haemulidae), within the Mediterranean Sea with new additional records from the French Mediterranean coast. *Italian Journal of Zoology*, *80*(2), 273–278. https://doi.org/10.1080/11250003.2012.730555

Bonney, R., Phillips, T. B., Ballard, H. L., & Enck, J. W. (2015). Can citizen science enhance public understanding of science? *Public Understanding of Science*, *25*(1), 2–16. https://doi.org/10.1177/0963662515607406

Bottacini, D., Pollux, B. J., Nijland, R., Jansen, P. A., Naguib, M., & Kotrschal, A. (2024). Lionfish (Pterois miles) in the Mediterranean Sea: a review of the available knowledge with an update on the invasion front. *NeoBiota*, *92*, 233-257.

Bowser, A., Hansen, D., He, Y., Boston, C., Reid, M., Gunnell, L., & Preece, J. (2013). Using gamification to inspire new citizen science volunteers. *Proceedings of the First International Conference on Gameful Design, Research, and Applications*. https://doi.org/10.1145/2583008.2583011

Burgiel, S., Foote, G., Orellana, M., & Perrault, A. (2006). Invasive alien species and trade: integrating prevention measures and international trade rules. *The Center for International Environmental Law and Defenders of Wildlife, Washington, DC*, 66-74.

C. Pornpanomchai, B. Lurstwut, P. Leerasakultham, W. Kitiyanan, Shape and Texture based fish image recognition system, Kasetsart J. (Nat. Sci.) 47 : 624 - 634 (2013)

Catalán, I. A., Amaya Álvarez-Ellacuría, José-Luis Lisani, Josep Sánchez, Vizoso, G., Antoni Enric Heinrichs-Maquilón, Hinz, H., Josep Alós, Signarioli, M., Jacopo Aguzzi, Francescangeli, M., & Palmer, M. (2023). Automatic detection and classification of coastal Mediterranean fish from underwater images: Good practices for robust training. *Frontiers in Marine Science*, *10*. https://doi.org/10.3389/fmars.2023.1151758

Chicho, Bahzad Taha, and Amira Bibo Sallow. “A Comprehensive Survey of Deep Learning Models Based on Keras Framework.” Journal of Soft Computing and Data Mining, vol. 2, no. 2, 24 Oct. 2021, pp. 49–62, publisher.uthm.edu.my/ojs/index.php/jscdm/article/view/8732.

Cooperativa, V. M., Mytilene, L. I., Salento, V. L., Anton, C. S. Z. S., & MARIA, C. F. (n.d.) Mediterranean Marine Science.

Cui, Y., Jia, M., Lin, T. Y., Song, Y., & Belongie, S. (2019). Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 9268-9277).

David, P., Thebault, E., Anneville, O., Duyck, P. F., Chapuis, E., & Loeuille, N. (2017). Impacts of invasive species on food webs: a review of empirical data. *Advances in ecological research*, *56*, 1-60.

Deidun, A., & Sciberras, A. (2017). *Unearthing marine biodiversity through citizen science - the Spot the Jellyfish and the Spot the Alien Fish campaign case studies from the Maltese Islands (Central Mediterranean)*. https://core.ac.uk/download/pdf/132620063.pdf

Deidun, A., Insacco, G., Galdies, J., Balistreri, P., & Zava, B. (2021). Tapping into hard-to-get information: the contribution of citizen science campaigns for updating knowledge on range- expanding, introduced and rare native marine species in the Malta-Sicily Channel. *BioInvasions Records*, *10*(2), 257–269. https://doi.org/10.3391/bir.2021.10.2.03

Deidun, A., Vella, P., Sciberras, A., & Sammut, R. (2010). New records of Lobotes surinamensis (Bloch, 1790) in Maltese coastal waters. *Aquatic Invasions*, *5*(Supplement 1), S113–S116. https://doi.org/10.3391/ai.2010.5.s1.023

Demertzis, K., Iliadis, L. S., & Anezakis, V.-D. (2018). Extreme deep learning in biosecurity: the case of machine hearing for marine species identification. *Journal of Information and Telecommunication*, *2*(4), 492–510. https://doi.org/10.1080/24751839.2018.1501542

Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., & Fei-Fei, L. (2009, June). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition* (pp. 248-255). Ieee.

Dickinson, J. L., Zuckerberg, B., & Bonter, D. N. (2010). Citizen science as an ecological research tool: challenges and benefits. *Annual review of ecology, evolution, and systematics*, *41*(1), 149-172.

Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., & Darrell, T. (2014, June). DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition. In ICML (pp. 647-655).

Dosemagen, S., & Parker, A. (2019). Citizen science across a spectrum: Building partnerships to broaden the impact of citizen science. *Science & Technology Studies*, *32*(2), 24-33.

El Naqa, I., & Murphy, M. J. (2015). *What is machine learning?* (pp. 3-11). Springer International Publishing.

Endresen, Ø., Behrens, H. L., Brynestad, S., Andersen, A. B., & Skjong, R. (2004). Challenges in global ballast water management. *Marine pollution bulletin*, *48*(7-8), 615-623.

Evans, J., Barbara, J., & Schembri, P. J. (2015). Updated review of marine alien species and other "newcomers" recorded from the Maltese Islands (Central Mediterranean). *Mediterranean Marine Science*, *16*(1), 225. https://doi.org/10.12681/mms.1064

Fawzi A., Samulowitz H., Turaga D. and Frossard P., "Adaptive data augmentation for image classification," *2016 IEEE International Conference on Image Processing (ICIP)*, Phoenix, AZ, USA, 2016, pp. 3688-3692, doi: 10.1109/ICIP.2016.7533048.

Fisher, J. P. (2005). An overview of international initiatives, treaties, agreements and management actions addressing alien invasive species. *Phillips, M. Bueno, P*.

Fleuré, V., Magneville, C., Mouillot, D., & Sébastien Villéger. (2024). Automated identification of invasive rabbitfishes in underwater images from the Mediterranean Sea. *Aquatic Conservation*, *34*(1). https://doi.org/10.1002/aqc.4073

Galanidi, M., Zenetos, A., & Bacher, S. (2018). Assessing the socio-economic impacts of priority marine invasive fishes in the Mediterranean with the newly proposed SEICAT methodology. *Mediterranean Marine Science*, *19*(1), 107. https://doi.org/10.12681/mms.15940

Galil, B. S. (2007). Loss or gain? Invasive aliens and biodiversity in the Mediterranean Sea. *Marine Pollution Bulletin*, *55*(7-9), 314–322. https://doi.org/10.1016/j.marpolbul.2006.11.008

García-Díaz, P., Cassey, P., Norbury, G., Lambin, X., Montti, L., Pizarro, J. C., ... & Tomasevic, J. A. (2021). Management policies for invasive alien species: addressing the impacts rather than the species. *BioScience*, *71*(2), 174-185.

García-Llorente, M., Martín-López, B., González, J. A., Alcorlo, P., & Montes, C. (2008). Social perceptions of the impacts and benefits of invasive alien species: Implications for management. *Biological Conservation*, *141*(12), 2969–2983. https://doi.org/10.1016/j.biocon.2008.09.003

Gauci, A., Deidun, A., & Abela, J. (2020). Automating Jellyfish Species Recognition through Faster Region-Based Convolution Neural Networks. *Applied Sciences*, *10*(22), 8257. https://doi.org/10.3390/app10228257

Giakoumi, S., Katsanevakis, S., Albano, P. G., Azzurro, E., Cardoso, A. C., Cebrian, E., ... & Sghaier, Y. R. (2019). Management priorities for marine invasive species. *Science of the total environment*, *688*, 976-982.

Goreau, T. J., & Hilbertz, W. (2005). Marine ecosystem restoration: costs and benefits for coral reefs. *World resource review*, *17*(3), 375-409.

Goren, M., & Galil, B. S. (2005). A review of changes in the fish assemblages of Levantine inland and marine ecosystems following the introduction of non-native fishes. *Journal of Applied Ichthyology*, *21*(4), 364-370.

Gulli, A., & Pal, S. (2017). *Deep learning with Keras*. Packt Publishing Ltd.

Hamzaoui, M., Mohamed Ould-Elhassen Aoueileyine, Lamia Romdhani, & Ridha Bouallègue. (2023). An Improved Deep Learning Model for Underwater Species Recognition in Aquaculture. *Fishes*, *8*(10), 514–514. https://doi.org/10.3390/fishes8100514

Han, D., Liu, Q., & Fan, W. (2018). A new image classification method using CNN transfer learning and web data augmentation. *Expert Systems with Applications*, *95*, 43–56. https://doi.org/10.1016/j.eswa.2017.11.028

Haubrock, P. J., Turbelin, A. J., Cuthbert, R. N., Novoa, A., Taylor, N. G., Angulo, E., ... & Courchamp, F. (2021). Economic costs of invasive alien species across Europe. *NeoBiota*, *67*, 153-190.

He, K., Zhang, X., Ren, S., & Sun, J. (2016a). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770- 778).

He, K., Zhang, X., Ren, S., & Sun, J. (2016b). Identity mappings in deep residual networks. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14* (pp. 630-645). Springer International Publishing.

Hernández-Serna, Andrés, and Luz Fernanda Jiménez-Segura. “Automatic Identification of Species with Neural Networks.” Ed. Mark Costello. PeerJ 2 (2014): e563. PMC. Web. 30 Sept. 2016.

Hridayami, P., Putra, I. K. G. D., & Wibawa, K. S. (2019). Fish species recognition using VGG16 deep convolutional neural network. *Journal of Computing Science and Engineering*, *13*(3), 124-130.

Huang, C., Li, Y., Loy, C. C., & Tang, X. (2016). Learning deep representation for imbalanced classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 5375-5384).

Iacona, G. D. (2014). The economic costs and ecological benefits of protected areas for biodiversity conservation.

Ikechukwu, A. V., Murali, S., Deepu, R., & Shivamurthy, R. C. (2021). ResNet-50 vs VGG-19 vs training from scratch: A comparative analysis of the segmentation and classification of Pneumonia from chest X-ray images. *Global Transitions Proceedings*, *2*(2), 375-381.

Iman, M., Arabnia, H. R., & Rasheed, K. (2023). A review of deep transfer learning and recent advancements. *Technologies*, *11*(2), 40.

Iqbal, M. A., Wang, Z., Ali, Z. A., & Riaz, S. (2019). Automatic Fish Species Classification Using Deep Convolutional Neural Networks. *Wireless Personal Communications*. https://doi.org/10.1007/s11277-019-06634-1

Jalal, A., Salman, A., Mian, A., Shortis, M., & Shafait, F. (2020). Fish detection and species classification in underwater environments using deep learning with temporal information. *Ecological Informatics*, *57*, 101088.

Jogin, M., Madhulika, M. S., Divya, G. D., Meghana, R. K., & Apoorva, S. (2018, May). Feature extraction using convolution neural networks (CNN) and deep learning. In *2018 3rd IEEE international conference on recent trends in electronics, information & communication technology (RTEICT)* (pp. 2319-2323). IEEE.

Johnston, E. L., Dafforn, K. A., Clark, G. F., Rius, M., & Floerl, O. (2017). Anthropogenic activities promoting the establishment and spread of marine non-indigenous species post-arrival.

Joo D1, Kwan YS, Song J, Pinho C, Hey J, Won YJ, Identification of Cichlid Fishes from Lake Malawi using Computer Vision, 2013 Oct 25;8(10):e77686. doi: 10.1371/journal.pone.0077686

Joseph, F. J. J., Nonsiri, S., & Monsakul, A. (2021). Keras and TensorFlow: A hands-on experience. *Advanced deep learning for engineers and scientists: A practical approach*, 85- 111.

Kalogirou, S., Azzurro, E., Bariche, M., & Lameed, A. G. (2012). The ongoing shift of Mediterranean coastal fish assemblages and the spread of non-indigenous species. *Biodiversity Enrichment in a Diverse World. InTech, Rijeka, Croatia*, 263-280.

Kalogirou, S., Corsini, M., Kondilatos, G., & Wennhage, H. (2007). Diet of the invasive piscivorous fish Fistularia commersonii in a recently colonized area of the eastern Mediterranean. *Biological invasions*, *9*, 887-896.

Kapoor, A., Gulli, A., Pal, S., & Chollet, F. (2022). *Deep Learning with TensorFlow and Keras: Build and deploy supervised, unsupervised, deep, and reinforcement learning models*. Packt Publishing Ltd.

Karanasiou, A. P. (2018). Environmental liability in business management: Ballast Water Management Systems.

Katsanevakis, S., Coll, M., Piroddi, C., Steenbeek, J., Ben Rais Lasram, F., Zenetos, A., & Cardoso, A. C. (2014). Invading the Mediterranean Sea: biodiversity patterns shaped by human activities. *Frontiers in Marine Science*, *1*, 32.

Katsanevakis, S., Olenin, S., Puntila-Dodd, R., Rilov, G., Stæhr, P. A., Teixeira, H., ... & Tidbury, H. J. (2023). Marine invasive alien species in Europe: 9 years after the IAS Regulation. *Frontiers in Marine Science*, *10*, 1271755.

Katsanevakis, S., Poursanidis, D., Hoffman, R., Rizgalla, J., Rothman, S. B. S., Levitt-Barmats, Y. A., ... & Zenetos, A. (2020). Unpublished Mediterranean records of marine alien and cryptogenic species. *BioInvasions Records*, *9*(2), 165-182.

Katsanevakis, S., Tempera, F., & Teixeira, H. (2016). Mapping the impact of alien species on marine ecosystems: the Mediterranean Sea case study. *Diversity and Distributions*, *22*(6), 694–707. https://doi.org/10.1111/ddi.12429

Katsanevakis, S., Wallentinus, I., Zenetos, A., Leppäkoski, E., Çinar, M. E., Oztürk, B., Grabowski, M., Golani, D., & Cardoso, A. C. (2014). Impacts of invasive alien marine species on ecosystem services and biodiversity: a pan-European review. *Aquatic Invasions*, *9*(4), 391–423. https://doi.org/10.3391/ai.2014.9.4.01

Katsanevakis, S., Zenetos, A., Belchior, C., & Cardoso, A. C. (2013). Invading European Seas: assessing pathways of introduction of marine aliens. *Ocean & Coastal Management*, *76*, 64-74.

Kelta, Z. (2022, September). *YOLO Object Detection Explained: A Beginner's Guide*. Www.datacamp.com. https://www.datacamp.com/blog/yolo-object-detection-explained

Kim, J. S., & Hong, K. S. (2009). Color–texture segmentation using unsupervised graph cuts. *Pattern Recognition*, *42*(5), 735-750.

Kishore, R., Jute, A., & Phillip, K. (2018). *Ecological Assessment Of The Marine Invasive Alien Species (Ias) Pernaviridis In Trinidad And Tobago*. Institute of Marine Affairs.

Kolar, C. S., Courtenay Jr, W. R., Nico, L. G., & Hubert, W. (2010). Managing undesired and invading fishes. *Inland fisheries management in North America, 3rd edition. American Fisheries Society, Bethesda, Maryland*, 213-259.

Kourantidou, M., Cuthbert, R. N., Haubrock, P. J., Novoa, A., Taylor, N. G., Leroy, B., ... & Courchamp, F. (2021). Economic costs of invasive alien species in the Mediterranean basin. *NeoBiota*, *67*, 427-458.

Kumar, G. S., Painumgal, U. V., Kumar, M. C., & Rajesh, K. H. V. (2018). Autonomous underwater vehicle for vision based tracking. *Procedia computer science*, *133*, 169-180.

Kumar, G., & Bhatia, P. K. (2014, February). A detailed review of feature extraction in image processing systems. In *2014 Fourth international conference on advanced computing & communication technologies* (pp. 5-12). IEEE.

Kuris, A. M. (2002). Eradication of introduced marine pests. In *Managing for healthy ecosystems* (pp.549-556). CRC Press.

Le, N. V. L., Nguyen, D. T., Al-Tawaha, A. R. M., & Vo, D. T. (2021). A Study on Legal Policies and Solutions for Ship Ballast Water Treatment. *Water Conserv. Manag.*, *5*, 114-120.

Lei, Y., Jia, F., Lin, J., Xing, S., & Ding, S. X. (2016). An intelligent fault diagnosis method using unsupervised feature learning towards mechanical big data. *IEEE Transactions on Industrial Electronics*, *63*(5), 3137-3147.

Li, J., Gray, R. M., & Olshen, R. A. (2000). Multiresolution image classification by hierarchical modeling with two-dimensional hidden Markov models. *IEEE transactions on information theory*, *46*(5), 1826-1841.

Li, Z., Liu, F., Yang, W., Peng, S., & Zhou, J. (2021). A survey of convolutional neural networks: analysis, applications, and prospects. *IEEE transactions on neural networks and learning systems*, *33*(12), 6999-7019.

Liquete, C., Piroddi, C., Drakou, E. G., Gurney, L., Katsanevakis, S., Charef, A., & Egoh, B. (2013). Current status and future prospects for the assessment of marine and coastal ecosystem services: a systematic review. *PloS one*, *8*(7), e67737.

Liu, H., Liu, T., Gu, Y., Li, P., Zhai, F., Huang, H., & He, S. (2021). A high-density fish school segmentation framework for biomass statistics in a deep-sea cage. *Ecological Informatics*, *64*, 101367.

Lorenzo Nespereira, J. M., Gonzalez-Pajuelo, J. M., & Gregoire, M. (2003). Age and growth of the bastard grunt (Pomadasys incisus: Haemulidae) inhabiting the Canarian archipelago, Northwest Africa. *Fishery Bulletin*.

M. Hassoon, I. (2022). Fish Species Identification Techniques: A Review. *Al-Nahrain Journal of Science*, *25*(2), 39–44. https://doi.org/10.22401/anjs.25.2.08

M., Cebrian, E., Francour, P., Galil, B., & Savini, D. (2013). Monitoring marine invasive species in Mediterranean marine protected areas (MPAs): a strategy and practical guide for managers. *IUCN, Malaga*, *136*.

Ma, Y.-X., Zhang, P., & Tang, Y. (2018). Research on Fish Image Classification Based on Transfer Learning and Convolutional Neural Network Model. *Research on Fish Image Classification Based on Transfer Learning and Convolutional Neural Network Model*.https://doi.org/10.1109/fskd.2018.8686892

Magaletti, E., Garaventa, F., David, M., Castriota, L., Kraus, R., Luna, G. M., ... & Gollasch, S. (2018). Developing and testing an early warning system for non indigenous species and ballast water management. *Journal of sea research*, *133*, 100-111.

Mahmood, A., Bennamoun, M., An, S., Sohel, F., & Boussaid, F. (2020). ResFeats: Residual network based features for underwater image classification. *Image and Vision Computing*, *93*, 103811.

Manaswi, N. K., & Manaswi, N. K. (2018). Understanding and working with Keras. *Deep learning with applications using Python: Chatbots and face, object, and speech recognition with TensorFlow and Keras*, 31-43.

Mannino, A. M., Balistreri, P., & Deidun, A. (2017). The marine biodiversity of the Mediterranean Sea in a changing climate: the impact of biological invasions. *Mediterranean identities-environment, society, culture*, 101-127.

Marom, N.D.; Rokach, L.; Shmilovici, A. Using the confusion matrix for improving ensemble classifiers. In Proceedings of the 2010 IEEE 26th Convention of Electrical and Electronics Engineers in Israel, Eilat, Israel, 17–20 November 2010; pp. 555–559.

Mazza, G., Tricarico, E., Genovesi, P., & Gherardi, F. (2014). Biological invaders are threats to human health: an overview. *Ethology Ecology & Evolution*, *26*(2-3), 112-129.

Mifsud Scicluna, B., Gauci, A., & Deidun, A. (2024). AquaVision: AI-Powered Marine Species Identification. *Information*, *15*(8), 437.

Minchin, D. (2006). The transport and the spread of living aquatic species. In *The ecology of transportation: managing mobility for the environment* (pp. 77-97). Dordrecht: Springer Netherlands.

Minchin, D., Gollasch, S., & Wallentinus, I. (2005). *Vector pathways and the spread of exotic species in the sea*. ICES Cooperative Research Reports (CRR).

Mittal, S., Srivastava, S., & Jayanth, J. P. (2022). A survey of deep learning techniques for underwater image classification. *IEEE Transactions on Neural Networks and Learning Systems*, *34*(10), 6968-6982.

Modasshir, M., & Rekleitis, I. (2020, May). Enhancing coral reef monitoring utilizing a deep semi- supervised learning approach. In *2020 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 1874-1880). IEEE.

Molina-Molina, J. C., Salhaoui, M., Guerrero-González, A., & Arioua, M. (2021). Autonomous Marine Robot Based on AI Recognition for Permanent Surveillance in Marine Protected Areas. *Sensors*, *21*(8), 2664. https://doi.org/10.3390/s21082664

Moradi, R., Berangi, R., & Minaei, B. (2020). A survey of regularization strategies for deep models. *Artificial Intelligence Review*, *53*(6), 3947-3986.

Morales, E. F., & Escalante, H. J. (2022). A brief introduction to supervised, unsupervised, and reinforcement learning. In *Biosignal processing and classification using computational learning and intelligence* (pp. 111-129). Academic Press.

Muslim, M., & Zulkifli, M. F. R. (2021). Convolutional neural network architectures performance evaluation for fish species classification. *Journal of Sustainability Science and Management*, *16*(5), 124-139.

Occhipinti-Ambrogi, A., & Galil, B. (2010). Marine alien species as an aspect of global change. *Advances in Oceanography and Limnology*, *1*(1), 199–218. https://doi.org/10.1080/19475721003743876

Otero, M., Cebrian, E., Francour, P., Galil, B., & Savini, D. (2013). Monitoring marine invasive species in Mediterranean marine protected areas (MPAs): a strategy and practical guide for managers. *IUCN, Malaga*, *136*.

Ovalle, J. C., Vilas, C., & Antelo, L. T. (2022). On the use of deep learning for fish species recognition and quantification on board fishing vessels. *Marine Policy*, *139*, 105015.

Pajuelo, J. G., Lorenzo, J. M., Gregoire, M., & Domínguez-Seoane, R. (2003). Life history of the Pomadasys incisus (Osteichthyes: Haemulidae) of the Canarian Archipelago. *Scientia Marina*, *67*(2), 241-248.

Pan, S. J. (2020). Transfer learning. *Learning*, *21*, 1-2.

Park, S. S., Tran, V. T., & Lee, D. E. (2021). Application of various yolo models for computer vision- based real-time pothole detection. *Applied Sciences*, *11*(23), 11229.

Pešić, A., Marković, O., Joksimović, A., Ćetković, I., & Jevremović, A. (2020). Invasive Marine Species in Montenegro Sea Waters. *The Handbook of Environmental Chemistry*, 547–572. https://doi.org/10.1007/698_2020_700

Pateman, R. M., Dyke, A., & West, S. E. (2021). The diversity of participants in environmental citizen science. *Citizen Science: Theory and Practice*.

Ramzan, F., Khan, M. U. G., Rehmat, A., Iqbal, S., Saba, T., Rehman, A., & Mehmood, Z. (2019). A Deep Learning Approach for Automated Diagnosis and Multi-Class Classification of Alzheimer's Disease Stages Using Resting-State fMRI and Residual Neural Networks. *Journal of Medical Systems*, *44*(2). https://doi.org/10.1007/s10916-019-1475-2

Rathi, D., Jain, S., & Indu, S. (2017, December). Underwater fish species classification using convolutional neural network and deep learning. In *2017 Ninth international conference on advances in pattern recognition (ICAPR)* (pp. 1-6). Ieee.

Rauf, H. T., Lali, M. I. U., Zahoor, S., Shah, S. Z. H., Rehman, A. U., & Bukhari, S. A. C. (2019). Visual features based automated identification of fish species using deep convolutional neural networks. *Computers and electronics in agriculture*, *167*, 105075.

Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). *You Only Look Once: Unified, Real-Time Object Detection*. https://arxiv.org/pdf/1506.02640

Richardson, M., & Domingos, P. (2006). Markov logic networks. *Machine learning*, *62*, 107-136.

Riley, S. (2009). Preventing Transboundary Harm From Invasive Alien Species. *Review of European Community & International Environmental Law*, *18*(2), 198–210. https://doi.org/10.1111/j.1467-9388.2009.00641.x

Rim, Z. K., & Mohamed-Nejmeddine, B. (2011). Reproductive biology of the lessepsian reticulated leatherjacket Stephanolepis diaspros (Fraser-Brünner, 1940) in the Gulf of Gabes (Eastern Mediterranean Sea). *Reviews in Fish Biology and Fisheries*, *21*(3), 641-648.

Rizzini, D. L., Kallasi, F., Oleari, F., & Caselli, S. (2015). Investigation of vision-based underwater object detection with multiple datasets. *International Journal of Advanced Robotic Systems*, *12*(6), 77.

Rodrigues, M.T.A., Freitas, M.H.G., Pádua, F.L.C. et al. Pattern Anal Applic (2015) 18: 783. doi:10.1007/s10044-013-0362-6

Rova, A., Mori, G., & Dill, L. M. (2007, May). One fish, two fish, butterfish, trumpeter: Recognizing fish in underwater video. In *MVA* (pp. 404-407).

Roy, H., Pauchard, A., Stoett, P., Renard Truong, T., Bacher, S., Galil, B., ... & Vandvik, V. (2023). Summary for Policymakers of the Thematic Assessment Report on Invasive Alien Species and their Control. *IPBES Invasive Alien Species Assessment*, 1-56.

Ruesink, J. L., & Collado-Vides, L. (2006). Modeling the increase and control of *Caulerpa taxifolia*, an invasive marine macroalga. *Biological Invasions*, *8*, 309-325.

Rum, S. N. M., & Az, F. (2021). FishDeTec: A Fish Identification Application using Image Recognition Approach. *International Journal of Advanced Computer Science and Applications*, *12*(3). https://doi.org/10.14569/ijacsa.2021.0120312

S. Cadieux, F. Lalonde, and F. Michaud, “Intelligent System for Automated Fish Sorting and Counting,” IEEE IROS, pp. 1279–1284, 2000

S.O. gunlana, O. Olabode, S.A.A. Oluwadare, G.B. Iwasokun, Fish Classification Using Support Vector Machine, African Journal of Computing & ICT, Vol 8. No. 2 June, 2015

Salman, A., Jalal, A., Shafait, F., Mian, A., Shortis, M., Seager, J., & Harvey, E. (2016). Fish species classification in unconstrained underwater environments based on deep learning. *Limnology and Oceanography: Methods*, *14*(9), 570–585. https://doi.org/10.1002/lom3.10113

Salman, A., Siddiqui, S. A., Shafait, F., Mian, A., Shortis, M. R., Khurshid, K., ... & Schwanecke, U. (2020). Automatic fish detection in underwater videos by a deep neural network-based hybrid motion learning system. *ICES Journal of Marine Science*, *77*(4), 1295-1307.

Sanchez, S. A., Romero, H. J., & Morales, A. D. (2020). A review: Comparison of performance metrics of pretrained models for object detection using the TensorFlow framework. *IOP Conference Series: Materials Science and Engineering*, *844*, 012024. https://doi.org/10.1088/1757- 899x/844/1/012024

Sarang, P. (2020). Deep Dive in tf.keras. *Apress EBooks*, 71–132. https://doi.org/10.1007/978-1-4842- 6150-7_3

Sarkar, P., De, S., & Gurung, S. (2023). Fish Detection from Underwater Images Using YOLO and Its Challenges. *Advances in Intelligent Systems and Computing*, 149–159. https://doi.org/10.1007/978-981-99-1472-2_13

Sarker, I. H. (2021). Deep learning: a comprehensive overview on techniques, taxonomy, applications and research directions. *SN computer science*, *2*(6), 420.

Schaaf, N., de Mitri, O., Kim, H. B., Windberger, A., & Huber, M. F. (2021). Towards measuring bias in image classification. In *Artificial Neural Networks and Machine Learning–ICANN 2021: 30th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 14– 17, 2021, Proceedings, Part III 30* (pp. 433-445). Springer International Publishing.

Schembri, P. J. (1997). The Maltese Islands: climate, vegetation and landscape. *GeoJournal*, *41*, 1-11. Sciberras, M., & Schembri, P. J. (2007). A critical review of records of alien marine species from the Maltese Islands and surrounding waters (Central Mediterranean). *Mediterranean Marine Science*, *8*(1), 41. https://doi.org/10.12681/mms.162

Shakman, E. A. (2008). *Lessepsian migrant fish species of the coastal waters of Libya: Status, biology, ecology* (Doctoral dissertation, Rostock, Univ., Diss., 2008).

Shaltout, M., & Omstedt, A. (2014). Recent sea surface temperature trends and future scenarios for the Mediterranean Sea. *Oceanologia*, *56*(3), 411-443.

Siddiqui, S. A., Salman, A., Malik, M. I., Shafait, F., Mian, A., Shortis, M. R., & Harvey, E. S. (2018). Automatic fish species classification in underwater videos: exploiting pre-trained deep neural network models to compensate for limited labelled data. *ICES Journal of Marine Science*, *75*(1), 374-389.

Spampinato, C., Giordano, D., Di Salvo, R., Chen-Burger, Y. H. J., Fisher, R. B., & Nadarajan, G. (2010, October). Automatic fish classification for underwater species behavior understanding. In *Proceedings of the first ACM international workshop on Analysis and retrieval of tracked events and motion in imagery streams* (pp. 45-50).

Targ, S., Almeida, D., & Lyman, K. (2016). Resnet in Resnet: Generalizing Residual Architectures. *ArXiv:1603.08029 [Cs, Stat]*. https://arxiv.org/abs/1603.08029

Templado, J. (2014). Future trends of Mediterranean biodiversity. *The Mediterranean Sea: Its history and present challenges*, 479-498.

Thiele, J., Kollmann, J., Markussen, B., & Otte, A. (2010). Impact assessment revisited: improving the theoretical basis for management of invasive alien species. *Biological Invasions*, *12*, 2025- 2035.

Thrun, S. (1996). Is learning the n-th thing any easier than learning the first?. In Advances in neural information processing systems, 640-646.

Torrey, L., & Shavlik, J. (2010). Transfer learning. In *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques* (pp. 242-264). IGI global.

Vabalas, A., Gowen, E., Poliakoff, E., & Casson, A. J. (2019). Machine learning algorithm validation with a limited sample size. *PloS one*, *14*(11), e0224365.

Vilà, M., Espinar, J. L., Hejda, M., Hulme, P. E., Jarošík, V., Maron, J. L., Pergl, J., Schaffner, U., Sun, Y., & Pyšek, P. (2011). Ecological impacts of invasive alien plants: a meta-analysis of their effects on species, communities and ecosystems. *Ecology Letters*, *14*(7), 702–708. https://doi.org/10.1111/j.1461-0248.2011.01628.x

Villon, S.; Iovan, C.; Mangeas, M.; Claverie, T.; Mouillot, D.; Villéger, S.; Vigliola, L. Automatic underwater fish species classification with limited data using few-shot learning. *Ecol. Inform.* **2021**, *63*, 101320. [**Google Scholar**] [**CrossRef**]

Wäldchen, J., & Mäder, P. (2018). Machine learning for image based species identification. *Methods in Ecology and Evolution*, *9*(11), 2216–2225. https://doi.org/10.1111/2041-210x.13075

Weiss, K., Khoshgoftaar, T. M., & Wang, D. (2016). A survey of transfer learning. *Journal of Big data*, *3*, 1-40.

Wu, X., & Li, T. (2024). A deep learning-based car accident detection approach in video-based traffic surveillance. *Journal of Optics*, 1-9.

Y. Nagashima, Takakazu Ishimatsu, A Morphological Approach to Fish Discrimination, IAPR Workshop on Machine Vision APPLICATIONS, Nov. 17-19,1998

Ying, X. (2019, February). An overview of overfitting and its solutions. In *Journal of physics: Conference series* (Vol. 1168, p. 022022). IOP Publishing.

Zarzyczny, K. M., Rius, M., Williams, S. T., & Fenberg, P. B. (2023). The ecological and evolutionary consequences of tropicalisation. *Trends in Ecology & Evolution*.

Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyals, O. (2021). Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, *64*(3), 107-115.

Zhang, W. (2023). Automated Fruit Grading in Precise Agriculture using You Only Look Once Algorithm. *International Journal of Advanced Computer Science and Applications*, *14*(10).

Zhang, W., Deng, L., Zhang, L., & Wu, D. (2022). A survey on negative transfer. *IEEE/CAA Journal of Automatica Sinica*, *10*(2), 305-329.

Zhong, J., Li, M., Qin, J., Cui, Y., Yang, K., & Zhang, H. (2022). Real-time marine animal detection using YOLO-based deep learning networks in the coral reef ecosystem. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *46*, 301-306.

Zhong, Z., Li, J., Luo, Z., & Chapman, M. (2017). Spectral–spatial residual network for hyperspectral image classification: A 3-D deep learning framework. *IEEE Transactions on Geoscience and Remote Sensing*, *56*(2), 847-858.

Zouari-Ktari, R., Bradai, M. N., & Bouain, A. (2008). The feeding habits of the Lessepsian fish Stephanolepis diaspros (Fraser-Brunner, 1940) in the Gulf of Gabes (eastern Mediterranean Sea). *Cahiers de Biologie marine*, *49*(4), 329.

APPENDICES

Appendix A

This research study has been peer-reviewed and published in a special issue journal of MDPI focused on “Machine Learning and Artificial Intelligence with Applications”. The publication details, including the link and citation, are provided below for reference.

The citation below provides the formal reference to the published research paper:

Mifsud Scicluna, B., Gauci, A., & Deidun, A. (2024). AquaVision: AI-Powered Marine Species Identification. Information, 15(8), 437. https://doi.org/10.3390/info15080437

The full text of the study can be accessed through the following link:

- Information | Free Full-Text | AquaVision: AI-Powered Marine Species Identification (mdpi.com)

This publication presents the full results of this study, including the methodologies, findings, and detailed analyses of the models used to classify invasive alien fish species. Readers are encouraged to consult the published paper for an in-depth exploration of the research and its implications.

Appendix B

To facilitate replication, validation, and further development of this research, access to the models used in this study, including TF.Keras, YOLO v8, and ResNet18, has been made available via the following GitFront repository:

- https://gitfront.io/r/bms1142002/gJZ5wNDmqqNg/AquaVision-Models