Haberdashers' School

Occasional Papers Series

Occasional Paper Number Ninety

Harvesting Intelligence: How Computer Vision is Transforming Food Security in Agriculture

Enora Hauduc

Haberdashers' Girls' School
Extended Research Project Prize Winner

November 2025

Haberdashers' School Occasional Paper Number Ninety
November 2025

Harvesting Intelligence: How Computer Vision is Transforming Food Security in Agriculture

Enora Hauduc

Abstract

With a rising population and greater concerns surrounding the effects of climate change, ensuring food security has become one of the most pressing challenges of the 21st century. To meet this demand, new data-driven and intelligent approaches to agriculture have emerged. This paper explores how computer vision, a branch of artificial intelligence that trains computers to interpret visual information, is transforming agricultural practices today.

By combining hyperspectral imaging with 3D Convolutional Neural Networks, computer vision allows for plant diseases to be detected earlier, reducing pesticide use and preventing crop loss. For crop planning, Fully Convolutional Neural Networks applied to satellite imagery allow for farmland to be mapped more precisely, improving yield estimation and land management. Within farms, YOLO-based object detection methods automate cattle health monitoring, allowing farmers to manage larger and healthier herds.

Despite some accessibility and computational challenges remaining, these applications collectively demonstrate how computer vision systems can enhance agricultural resilience and contribute to a more sustainable and food-secure future.

1.1 Introduction

With 733 million people facing food insecurity in 2023, hunger remains one of humanity's most urgent crises. An increasing population, combined with unpredictable weather, shrinking farmland, and rising temperatures, is placing the global food system under more pressure than ever before (United Nations, 2024).

In an attempt to solve this, in September 2015, all United Nations Member States adopted a 15-year plan aimed at eradicating poverty, improving global equality, and tackling climate change. The plan, known as the 17 Sustainable Development Goals (SDGs), centres on sustainability and long-term global development (United Nations, 2024). To address these challenges and support the UN's SDGs, modern and advanced agricultural innovation is essential. One such emerging technology, computer vision – a rapidly growing subfield of artificial intelligence focused on interpreting visual data – is already showing great potential for application in agriculture (Khan and Al-Habsi,2020).

This paper will explore the existing and emerging applications of computer vision and assess how they are currently impacting, and could further support, progress towards the UN's SDGs.

1.2.1 Application One: Plant Disease Detection

Crops, which account for the largest percentage of global calorie intake, suffer annual losses of up to 40% due to pests and plant diseases (National Geographic, 2025). This significantly impacts the availability of food and affects progress towards Zero Hunger (SDG 2), a goal aimed at preventing hunger and malnutrition (FAO, n.d.).

By detecting plant diseases earlier, their impact can be substantially reduced before the pathogen can spread to neighbouring plants. That said, it is time-consuming and labour-intensive for a farmer to constantly monitor crops for signs of disease, and most crop diseases remain asymptomatic until they have already spread.

Figure 1 – The Septoria tritici: The most damaging wheat disease in the UK (Alemayehu et al., 2012).

One traditional solution is to use pesticides as a preventive measure before any pathogens can affect the crops (Alavanja, 2010). However, pesticides harm local pollinators and wildlife. They also pose health risks, including asthma and cancer in humans, meaning they are not a suitable solution for safe and sustainable farming (Pesticide Collaboration, 2025).

A promising alternative uses a computer vision model to monitor crops for signs of disease. Unlike their human counterparts, these models provide continuous, highly accurate data that can alert farmers about possible outbreaks earlier, without negatively impacting the local environment (Shoaib et al., 2023).

1.2.2 Applying Hyperspectral Imaging

To detect these asymptomatic diseases, imaging methods that capture wavelengths beyond the visible light spectrum are required to highlight subtle indicators of plant disease. Unlike standard cameras, which only capture spectral data for red, green, and blue bands from the Electromagnetic spectrum (EM), a technology known as Hyperspectral Imaging (HSI) captures data from hundreds of narrow wavelength bands, creating a continuous spectral map across the entire EM spectrum.

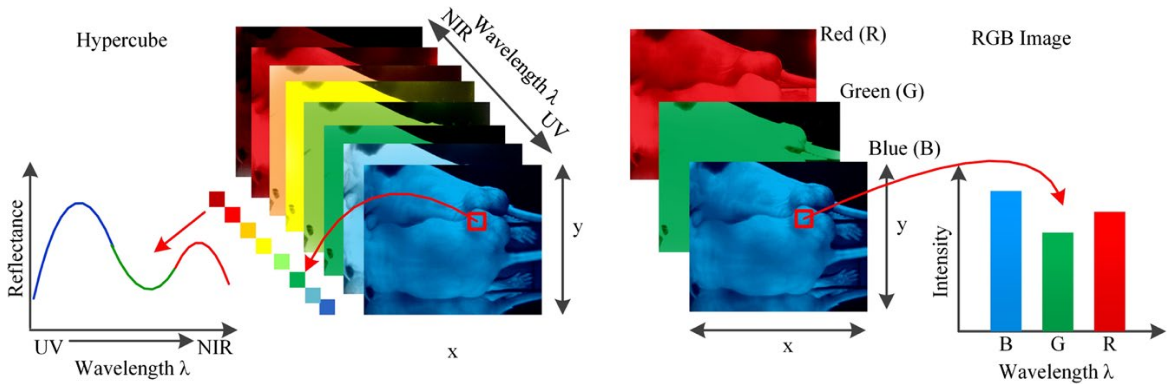

Figure 2 - Hyperspectral imagery spans the full wavelength range compared to just three wavelength bands (Bhargava et al., 2024)

HSI works by using hyperspectral sensors to record the unique pattern of light known as a ‘spectral signature’, which is reflected from plant leaves (Bhargava et al., 2024). Since every ‘spectral signature’ is unique to each material, HSI allows for the detection of subtle differences between healthy and infected plant tissue.

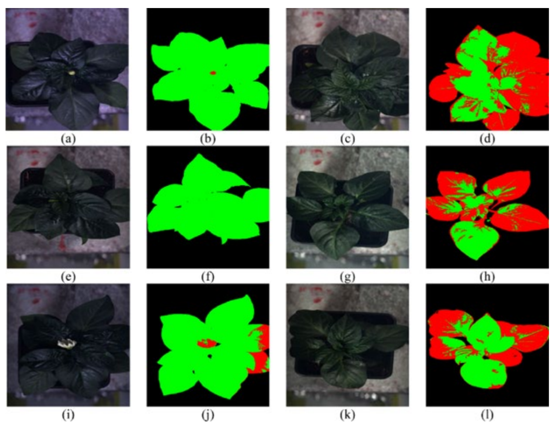

Figure 3 - Early detection of Tomato Spotted Wilt Virus using hyperspectral images. Healthy plants (a,b), (e,f), (i,j); TSWV-infected plants (c,d), (g,h), (k,l). Green: healthy pixels. Red: infected pixels. Note how it is difficult to distinguish between healthy and infected pixels without the hyperspectral analysis. (Wang et al., 2019)

1.2.3 An Introduction to 3D Convolutional Neural Networks

While HSI captures detailed spectral data, by itself, it cannot directly detect plant diseases. To interpret this data, it must be processed by a computer vision model, such as a 3D Convolutional Neural Network (3D CNN).

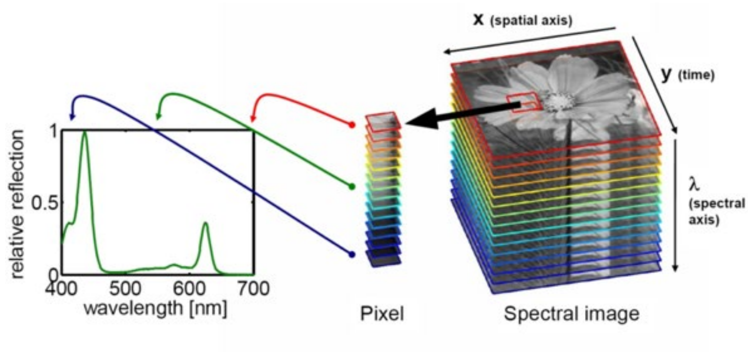

Figure 4 - Each pixel contains a stack of data for the entire reflective spectrum (Polder & Pekkeriet, 2013)

3D CNNs are designed to process three-dimensional data, so plant images captured using HSI technology must first be converted into a 'data cube' - stacks of plant images with each slice representing a specific wavelength. By using this format, it is possible to analyse each pixel individually through a spectrum of wavelengths, revealing both spatial and spectral connections between different areas of the plant. Once trained, the 3D CNN can then extract these patterns, ultimately performing an image classification role by sorting inputted images of plant tissues as healthy or diseased.

1.2.4 Inside the 3D CNN

To extract these patterns, the 3D CNN must process the inputted cube data through a series of layers. This process centres on the use of convolutional layers, which contain a set of 3D filters known as 'kernels' that scan through the data to identify significant features.

Each kernel is a three-dimensional matrix of numbers that slides across sections of the input data, multiplying the value in the kernel by the pixel values in the cube. The result of this multiplication is summed, producing a single value that captures the key features of that block of input data.

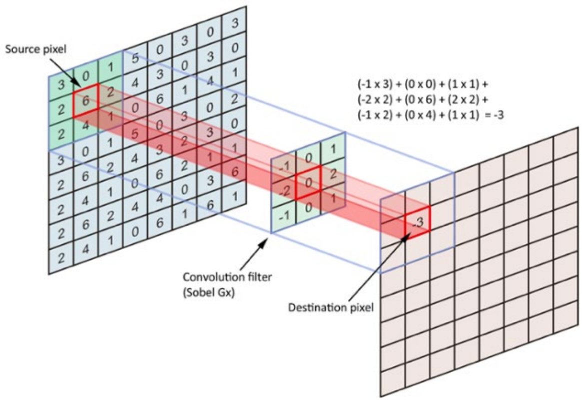

Figure 5 - The kernel multiplies each value in the kernel by the corresponding pixel value and then sums the result (Costa, 2019)

Repeating this across the entire data cube reduces the data size and creates a 3D 'feature map' that highlights specific patterns such as unusual textures or edges, which could indicate disease.

The model then passes this feature map through an ‘activation layer’ to introduce non-linearity into the data, allowing even more complex patterns to be detected. The most common method is ReLU (Rectified Linear Unit), which replaces all negative pixel values with a zero while only keeping the important positive values. Without non-linearity, the model would only be able to predict simple straight-line relationships between the input and output. In the case of plant diseases, introducing non-linearity allows the model to make connections across different dimensions (height, width, depth) and wavelengths so that it can learn intricate patterns such as the shape and texture of diseased plants.

In between each stage, the spatial dimensions of the cube are repeatedly reduced through ‘pooling’. Sections of pixels are replaced with one pixel representing the average value of that area, reducing the computational cost of the model, so that subsequent layers will require fewer computations to process.

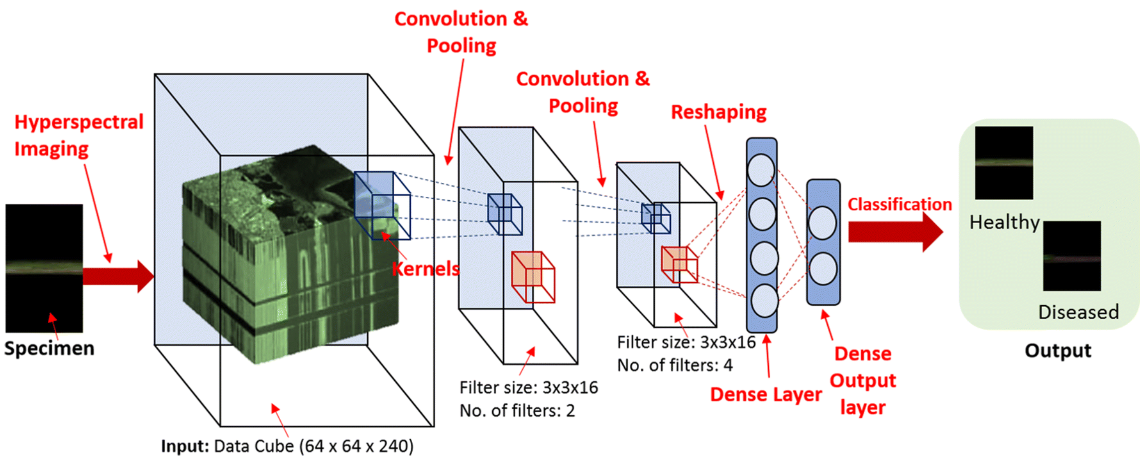

Figure 6 - The HSI cube is repeatedly convoluted, pooled, and activated before being flattened and classified into ‘healthy’ or ‘diseased’ (Nagasubramanian et al., 2019)

After multiple rounds of convolution, activation and pooling layers, the final cube is then ‘flattened’ into a 1-dimensional (1D) vector. This crucial stage allows for the data to act as an input into the ‘fully connected layer’, a layer that only accepts 1D inputs. Here, each of the elements from the 1D vector input directly impacts the output (compared to convolutional layers, where only some elements impact the next output). Finally, a calculation is performed on the 1-dimensional vector, leading to a single classification of ‘diseased’ or ‘healthy’ (Montesinos López, Montesinos López and Crossa, 2022; Yamashita et al., 2018).

1.2.5 Why Should We Use 3D CNNs?

Whilst 3D CNNs are essential for analysing 3D data such as HSI cubes, a major disadvantage of using high-dimensional models is the vast amount of computing power and time required to process the data. For example, a study found that training a 3D CNN to detect strawberry mould disease took 100 times as long as training a 2D CNN to perform the same task (Jung et al., 2022).

Some academics have proposed a solution to this problem by identifying and isolating only a few key wavelength bands that most effectively distinguish between diseased and healthy plant tissue (Nagasubramanian et al., 2019).

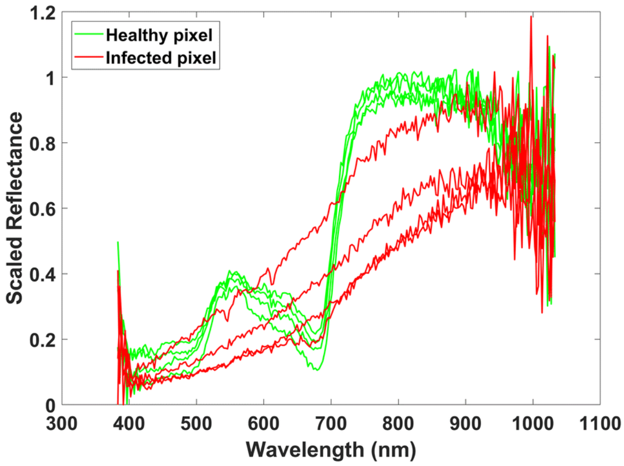

Figure 7 - Reflectance differences between healthy and infected plant tissue with a noticeable contrast in the 700-800nm range (Nagasubramanian et al., 2019)

These wavelengths could be captured using cheaper multispectral cameras tuned to this specific wavelength band and then analysed using a 2D CNN.

Currently, 3D CNNs must run on expensive processors known as Graphical Processing Units (GPUs) due to their high computational requirements. So, by reducing the processing power required, the model could become more accessible in rural areas that lack access to high-end hardware (Iqbal, Davies and Perez, 2024). However, a study in South Korea found that even at the most effective wavelength band (760-800nm), the multispectral camera lacked sufficient accuracy for diagnostics (Jung et al., 2022).

The study proposed an alternative solution by allowing the user to manually select the desired areas of the plant to be processed instead of processing an entire leaf or stem (Jung et al., 2022). This would reduce the dimensions of the inputted data, decreasing the computational cost. Nevertheless, the requirement for manual intervention by the user would offset the benefits of reduced processing power, meaning neither solution would be practical.

Overall, this means that using a HSI cube with a 3D CNN remains the most viable solution for plant disease detection, given that the alternative options have not achieved the same level of accuracy and efficiency.

1.2.6 The Future of Plant Disease Detection

Despite these challenges, the future usage of 3D CNNs for plant disease detection remains promising due to advancing research in the computational efficiency of neural networks. One new technique, known as Neural Architecture Search (NAS), can automate the design of neural network architectures to create more efficient models (Lin et al., 2024). Applying it to the disease-detection model would reduce the amount of computational power needed for training and execution, addressing the overarching issue of 3D CNNs' high computational requirements. Improved network efficiency would reduce the battery and memory requirements of the GPU running the model, allowing the 3D CNNs to run on edge devices such as smart cameras or drones. This would reduce reliance on the cloud for processing, making the future of plant disease detection more portable and accessible, especially in remote locations without internet access (Dong et al., 2023).

1.3.1 Application Two: Field Boundary Segmentation

In developing nations, where smallholder farms are prevalent and food security is a concern, tools for accurate monitoring of agricultural land are essential. These farms tend to have more irregularly shaped and smaller field boundaries than those on large commercial farms, making precise and consistent measurements difficult (Yang et al., 2020).

However, obtaining accurate estimates of field sizes is crucial for local governments to effectively determine yield per acre, a key metric for evaluating crop performance. For farmers, this data informs best practices and helps identify factors that limit production, enabling them to improve agricultural efficiency. Even a small error of ±0.1ha in the field size estimation of a 0.25ha field can deviate the actual yield by ±30%, potentially leading to poor decisions on productivity strategies.

Field boundary data also plays a direct role in calculating the quantity of seeds or fertiliser required. This is especially important in smallholder farms, where input costs have a significant impact on the amount of food that can be grown. Improving efficiency through better data analysis supports Life on Land (SDG 15), promoting better land management through precision agriculture (Flor et al., 2024).

Traditionally, field areas have been measured using boundary maps; though these become outdated, are not available in most smallholder farms, and are expensive to update regularly (Wagner and Natascha, 2020).

A better solution is to use satellite imagery: a standardised, large-scale imaging technique that captures images at regular intervals, allowing for continuous monitoring of fields. While satellite images alone do not compute field sizes, combining them with a computer vision model allows boundaries to be detected consistently and precisely, leading to more accurate measurements of area (Wagner and Natascha, 2020).

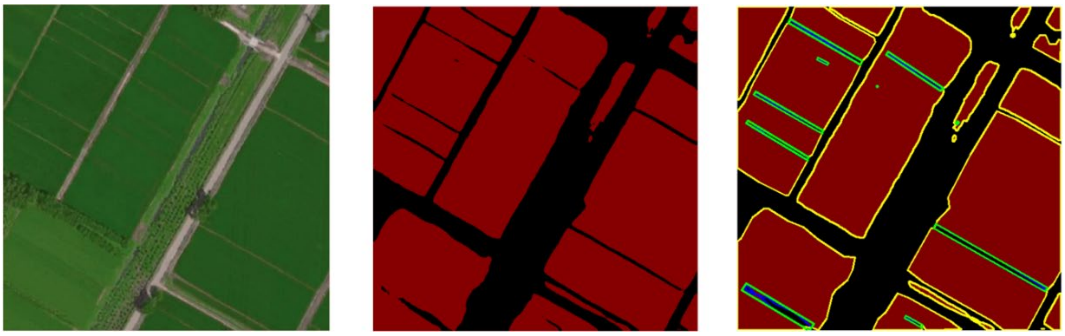

Figure 8 – Input field images (left), processing (middle), segmented field unit output (right) (Xu et al., 2023)

1.3.2 Instance and Semantic Segmentation

At the core of field boundary detection is image segmentation, which involves dividing sections of an image into meaningful regions. It is comprised of two key segmentation techniques: semantic and instance segmentation. Unlike object classification, which simply produces a label, segmentation creates a new image with the relevant boundaries highlighted.

Instance segmentation assigns each pixel in the image to a pre-defined category, distinguishing between separate instances of objects from the same category. The alternative approach, semantic segmentation, similarly classifies pixels into distinct categories but does not distinguish between multiple instances of the same object. For example, all fields would be simply labelled as 'field', without distinguishing between individual plots (Wołk and Tatara, 2024).

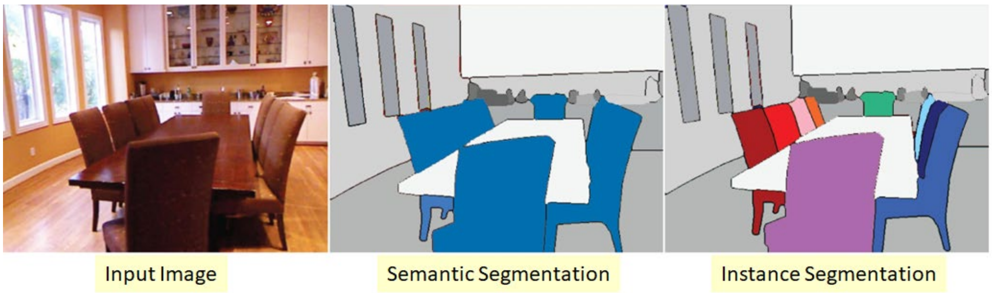

Figure 9 – Instance segmentation differentiates between different objects from the same class (as seen by different colours for each chair) (Solawetz, 2024)

For field boundary segmentation, it is unnecessary to differentiate between multiple instances from the same category. Instead of classifying the fields as 'field 1', 'field 2', etc., it is sufficient to distinguish between 'field' and 'soil'. Thus, semantic segmentation is the most suitable segmentation technique for this application.

1.3.3 Models for Field Boundary Segmentation

To effectively perform semantic segmentation, a specialised neural network known as a Fully Convolutional Neural Network (F-CNN) is required. Unlike traditional CNNs, which produce an overall label showing what the object is (but not where it is), F-CNNs label each pixel individually, producing a segmentation map that preserves the spatial layout of the input. This is crucial for delineating the fields and not just stating whether they are present in the inputted image (Huang et al., 2022).

1.3.4 An Overview of F-CNNs

These F-CNNs consist of two main components: an encoder and a decoder. The encoder repeatedly reduces the spatial dimensions of the image, filtering out unnecessary noise and extracting key features from each image. As the total number of pixels decreases, the image becomes more abstract, eventually consisting of very few pixels, which compactly represent critical information such as the edges or shapes of the field boundary.

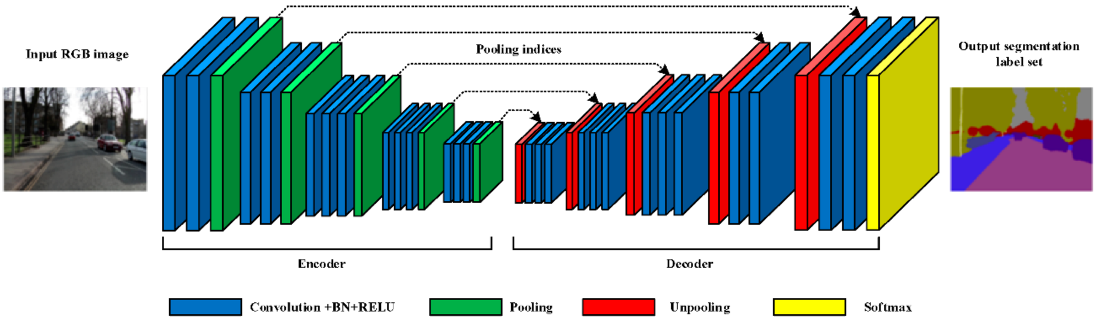

Figure 10 - Layers are encoded and decrease spatially, and then decoded by adding information from previously encoded layers (Huang et al., 2020)

The decoder section then reconstructs the image using 'skip connections' to reintroduce important features from earlier layers of the encoder until the image reaches its original size. Rather than simply 'undoing' the work of the encoder, it selectively reintroduces key pixels which recover essential details while avoiding unnecessary complexity.

As a result, encoding reduces detailed resolution, while during upsizing, the decoder reconstructs a broader view of the fields, preserving the general shape and layout of objects without reintroducing minor details (Pound, 2018). This encoder-decoder structure is crucial for field boundary segmentation, where it is essential for the model to generalise field shapes and not to focus on specific textures and patterns. This is important as it avoids overfitting - where the model memorises patterns without generalising to unseen data - as this leads to poor performance on unseen satellite images (IBM, 2021).

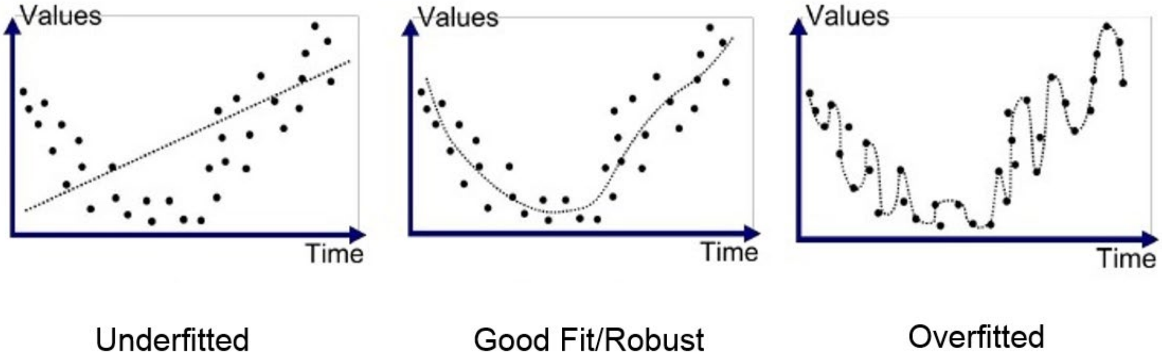

Figure 11 – The left image fails to capture patterns in the data, the middle one shows a well-fitting model, and the right one is overfitted and has memorised noise and unnecessary detail (Bhande, 2018)

1.3.5 Applying F-CNNs to Field Boundary Segmentation

The real-world applications of F-CNNs have already demonstrated promising results. In rural Bangladesh, where fields average just 0.105 hectares, an F-CNN model achieved a 90% accuracy in predicting field boundaries. Similarly, in Nigeria, where, despite trees obscuring some field boundaries, the F-CNN-based method managed to achieve high precision regardless, showing its robustness across different landscapes (Yang et al., 2020).

The main benefit of using F-CNNs compared to traditional CNNs is their ability to preserve spatial context during segmentation by producing an image output rather than a single prediction. Because traditional CNNs flatten data in their final layer, which loses these spatial relationships, maintaining spatial relationships within the image would be impossible without the use of an F-CNN (Kim, 2024). Additionally, F-CNNs lack fully connected layers: these are the final layers with a large number of parameters, where each pixel influences the overall prediction. As a result, their computational complexity is reduced, allowing them to process high-resolution images more efficiently, using less memory and training time (Pound, 2018).

Whilst F-CNNs have their advantages, they also have some limitations. The nature of the encoding and decoding process means that, despite the use of skip connections, some finer details will still be lost. However, because field boundary segmentation distinguishes between major field divisions, rather than capturing small, intricate details, some loss of detail is acceptable. Compared to medical imaging, where high precision is crucial, the advantages of the F-CNN, through their computational efficiency and preservation of spatial context, are more favourable in this context.

1.3.6 The Future of Field Boundary Segmentation

While F-CNNs form the basic model, several advanced variants, such as U-Net and PSPNet, have been created to resolve issues like loss of detail. U-Net, for example, reduces the loss of detail by using more extensive skip connections, creating a continuous flow of spatial information throughout the network (Kim, 2024).

Finer segmentation, however, is only as effective as the resolution of the satellite image used. To provide higher-quality images, a company known as Albedo Space intends to launch satellites with a resolution of 10cm for public use in 2025, a significant improvement to current resolutions of 15m for NASA's Landsat satellites (The Economist, 2025), (Maxar Technologies, n.d). With each pixel soon representing just a 10 x 10cm square and the resolution set to increase drastically over the coming years, field boundary segmentation models will become increasingly precise.

Figure 12 - A visual representation of different satellite resolutions on photo clarity (IPTSAT, 2024)

This means that using F-CNNs to segment field boundaries will continue to remain a vital use of computer vision, contributing to more resilient food systems and acting as an innovative tool for improving data-backed agricultural practices.

1.4.1 Application Three: Precision Cattle Farming

Over the past 50 years, the demand for meat and dairy has surged, with global production of cattle more than doubling since 1961 (Ritchie, Rosado, and Roser, 2017). Our reliance on cattle for meat and dairy means that to improve food security, meat and dairy must continue to remain available.

The main indicator of cattle productivity is their Body Condition Score (BCS), a 5-point scoring system that assesses body fat levels. Cattle with a low BCS have an increased risk of diseases such as mastitis, which shortens their life span and reduces their milk yield and meat quality (Krogstad and Bradford, 2025).

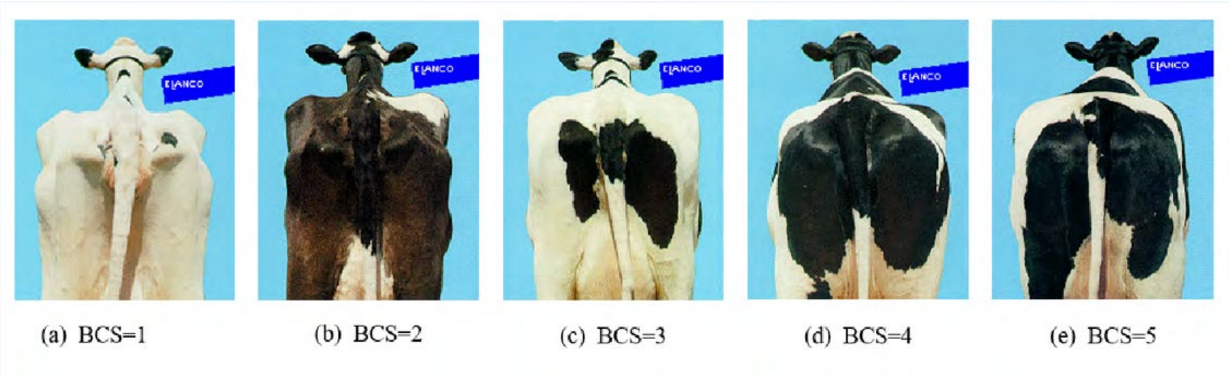

Figure 13 - Cow Body Condition Score (Li et al., 2019)

Equally, higher BCS indicates obese cows, a sign of feed inefficiency. For farmers, there is an economic incentive to ensure cows maintain a suitable BCS score as the efficient conversion of feed optimises resources, in turn affecting responsible consumption and production (SDG 12).

Traditionally, BCS must be conducted through thorough and time-intensive hands-on assessments, which are performed at least five times annually per cow (Megalac, 2024). However, with the average farmer in the UK managing a herd of 152 cows (an increasing number due to economic pressures), regular and accurate assessments are difficult to maintain (Prior, M., 2024).

For this reason, integrating a computer vision system that continuously monitors cattle BCS could provide farmers with valuable feedback regarding their herd's health. Unlike periodic and manual checks, automated systems would provide real-time analysis, increasing the efficiency and accuracy of cattle rearing (Li, G. et al., 2021).

1.4.2 An Overview of Object Detection

To identify a cow's BCS, we must use object detection: a computer vision technique that identifies and locates objects in an image using a square known as a 'bounding box'. This is particularly useful when tracking and identifying individual cows within crowded environments such as barns.

There are two main approaches to object detection: two-shot detection and single-shot detection. Two-shot detection processes the inputted image twice: first, it finds potential object locations, before refining this prediction in a second pass to establish these locations. The result of this is an object detection which is more accurate but also more computationally expensive. Using it in a barn, where surveillance occurs 24/7, would be impractical as the thousands of frames generated per hour would need to be processed twice, meaning it would quickly become unsustainable in the long term.

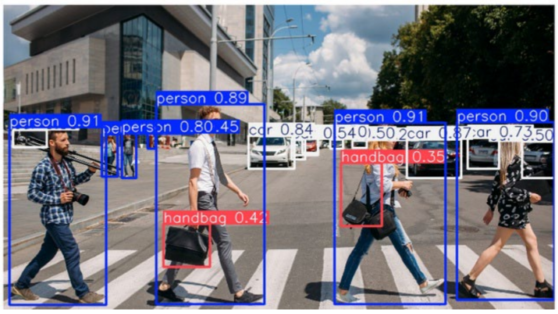

Figure 14 - Each bounding box prediction contains a confidence level rating (Morgunov, 2024)

Instead, a more efficient approach known as single-shot object detection processes the entire image in one pass, simultaneously making predictions about the position and occurrence of the cows. While this model is more computationally efficient, it is also less accurate at detecting smaller objects. The most commonly used single-shot model is YOLO (You Only Look Once), which balances both speed and accuracy, making it ideal for real-time cow monitoring (Parab et al., 2022).

1.4.3 Usage of YOLO for Object Detection

YOLO is a CNN that takes image inputs and outputs predicted bounding boxes, along with corresponding class labels. It uses a grid-based detection system, dividing the image into smaller cells known as 'grid cells', which individually predict the presence and location of an object in their area. The multiple grid cells mean the model can process predictions in parallel instead of scanning the image linearly, making YOLO incredibly fast.

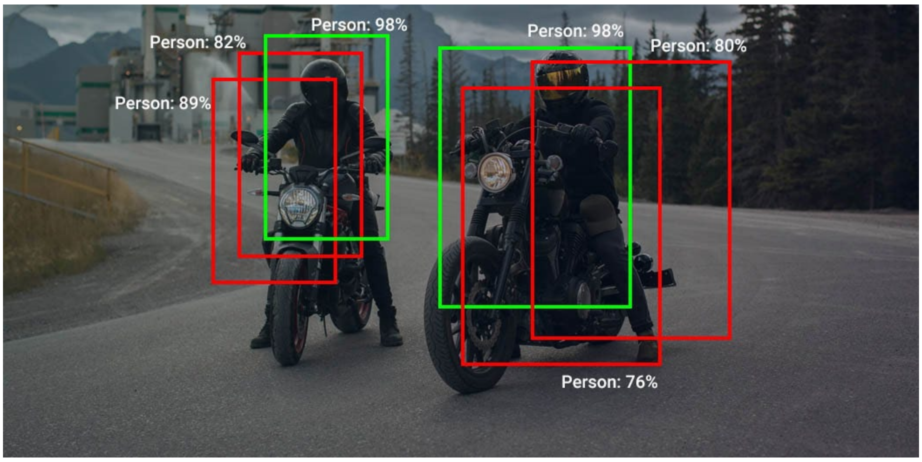

Figure 15 - If there are multiple bounding boxes, the box with the highest confidence score is selected (shown by the green box) (Bhalerao, C., 2023)

After predictions have been made, a post-processing technique known as 'Non-maximum Suppression' (NMS) is used to remove any duplicate bounding boxes. NMS compares the probability of each overlapping bounding box, keeping only the bounding box with the highest confidence score, thus reducing false positives from the image (Diwan, Anirudh and Tembhurne, 2022).

1.4.4 Applications of YOLO to BCS

A 2024 study in Türkiye successfully applied YOLOv8 to 1,270 cows, using experts to manually create the training data (Dandıl, Çevik, and Boğa, 2024). The researchers set up cameras in a barn to capture images of cow tails, achieving an overall accuracy of 0.81. The study considered two different breeds of cow, Holstein and Simmental, which have different appearances due to their genetics. Holsteins, typically leaner and used for milk, had more visible skeletal features, resulting in a slightly higher accuracy of 0.82. By contrast, Simmental cows, larger and bred for both milk and meat, were harder to classify accurately, scoring 0.80 (Koknaroglu et al., 2021).

Figure 16 – Correct classifications are denoted by a green box, whilst incorrect classifications by a red box (Singh, A,2020)

The study found that misclassification occurred between the 'Good' and 'Fat' categories, possibly caused by Holsteins remaining leaner even when they had gained weight. In Simmental cows, the misclassification occurred between the 'Fat' and 'Obese' categories, due to their body structure naturally including more fat and muscle.

A solution to this problem could be to implement breed-specific models, adjusting BCS thresholds for each breed of cow. Overall, whilst the accuracy was not perfect, BCS is ultimately a partially subjective scoring system, and the model's ability to differentiate between 'Obese' and 'Emaciated' cows already provides significant improvement to herd monitoring (DairyNZ, 2024).

1.4.5 Why Do We Use YOLO?

YOLO is a highly successful one-shot object detection algorithm known for its speed and real-time detection accuracy. With a frame per second (FPS) rate of 45, object tracking accuracy is increased in environments where changes may happen frequently, such as when cows move around their shed.

However, the main disadvantage of YOLO is its inability to handle small objects. Since YOLO works by dividing the image into a grid, each cell can predict only a limited number of bounding boxes. This means very small objects may not be detected if many other objects are already present in the same cell. In this context, though, cows are large objects and, if spread out in the pen, are impacted minimally by this disadvantage, making YOLO a suitable model, nonetheless.

Another challenge is that, like most models, YOLO operates as a 'black box model' where the code, and therefore decision-making processes, are not available to farmers and veterinarians. Without understanding how these labels are applied, farmers and veterinarians may find it more difficult to trust or interpret these results (Diwan, Anirudh and Tembhurne, 2022).

1.4.6 The Future of BCS

Since its launch in 2015, the YOLO family has progressed significantly. And with the latest YOLOv11 released in 2024, each iteration brings continual improvements in speed, accuracy, and efficiency (Ali and Zhang, 2024).

Additionally, as higher-level camera resolutions of 6k and 8k become more standard, future YOLO models will be able to leverage this higher resolution to detect objects that cover a smaller proportion of the frame, improving the model's accuracy (Minami and Nishikawa, 2024), (Pepe et al., 2022). For the application of cow BCS monitoring, this means cows in the background will be detected without needing close-range cameras, making the system more flexible and effective.

1.5 Conclusion

Each of the applications of computer vision outlined above offer scalable, non-invasive solutions to improving food security by reducing food waste, improving agricultural planning, and increasing livestock efficiency.

That said, each application has varying levels of impact on the SDG goals. Whilst cattle production is increasing globally, we are seeing a shift to more conscious meat consumption closer to home (AHDB, 2025). In the UK, the domestic cattle supply is already falling at a rate of 5% each year, meaning the impact of implementing a BCS detection model in the UK would be minimal (AHDB, 2025). Similarly, field boundary segmentation, while an invaluable tool in smallholder farms, has limited effects on large-scale industrial farms, which tend to be built in regular geometric shapes with known field measurements.

By contrast, plant diseases affect farms in both developed and developing countries globally. With climate change increasing temperatures, diseases will spread quicker, making their early detection more vital than ever. Early detection is critical to reducing pesticide use, preventing crop loss, and improving food security. Therefore, investing in plant disease detection would have the greatest impact on improving global food security in agriculture.

Applications of machine learning are still a relatively new part of technology, so they are constantly evolving and improving. Ultimately, by making these intelligent tools available on a global scale, the usage of computer vision not only addresses food insecurity but also contributes to a more resilient, efficient, and productive global food system that benefits all.

Bibliography:

1. AHDB: *What's the long-term outlook for beef and sheep meat consumption?* [online]. Agriculture and Horticulture Development Board (AHDB). Available at: https://ahdb.org.uk/knowledge-library/what-s-the-long-term-outlook-for-beef-and-sheep-meat-consumption [Accessed 8 June 2025]
2. AHDB. 2025 *Beef Market Outlook - February 2025*. [online]. Agriculture and Horticulture Development Board (AHDB). Available at: https://ahdb.org.uk/beef-market-outlook [Accessed 8 June 2025]

3. Alavanja, M.C.R., 2010. Pesticides use and exposure: Extensive worldwide. *Rev Environ Health*, [online] Available at: https://pmc.ncbi.nlm.nih.gov/articles/PMC2946087/ [Accessed 6 March 2025].
4. Alemayehu, W., Orozco, D., Tache, M., et al., 2012. Genetic basis of heterosis for yield in hybrid rice. *PLOS Genetics*, [online] Available at: https://journals.plos.org/plosgenetics/article?id=10.1371/journal.pgen.1002070 [Accessed 27 January 2025].
5. Ali, M.L. and Zhang, Z., 2024. The YOLO Framework: A Comprehensive Review of Evolution, Applications, and Benchmarks in Object Detection. *Computers*, 13(12), 336. Available at: https://doi.org/10.3390/computers13120336 [Accessed 18 March 2025].
   Available at: https://doi.org/10.3390/heritage5030099 [Accessed 23 March 2025]
6. Bhalerao, C., 2023. A Deep Dive Into Non-Maximum Suppression (NMS). [online] Available at: https://builtin.com/machine-learning/non-maximum-suppression [Accessed 18 March 2025].
7. Bhande, A., 2018. What is underfitting and overfitting in machine learning and how to deal with it. *GreyAtom*. [online] Available at: https://medium.com/greyatom/what-is-underfitting-and-overfitting-in-machine-learning-and-how-to-deal-with-it-6803a989c76 [Accessed 17 March 2025]
8. Bhargava, A., Sachdeva, A., Sharma, K., Alsharif, M.H., Uthansakul, P. and Uthansakul, M., 2024. Hyperspectral imaging and its applications: A review. *Heliyon*, [online] 10(6), p.e23999. Available at: https://doi.org/10.1016/j.heliyon.2024.e23999 [Accessed 23 March 2025].
9. C. Costa, V., 2019. Understanding the structure of a CNN (image) *Medium*. [online] Available at: https://viniciuscantocosta.medium.com/understanding-the-structure-of-a-cnn-b220148e2ac4 [Accessed 17 March 2025].
10. DairyNZ (2024) Body Condition Score Reference Guide 2024. Available at:https://www.dairynz.co.nz/media/hjpl5oj2/body-condition-score-reference-guide-2024.pdf [Accessed: 17 March 2025].
11. Dandıl, E., Çevik, K.K., and Boğa, M., 2024. Automated Classification System Based on YOLO Architecture for Body Condition Score in Dairy Cows. Available at: https://pmc.ncbi.nlm.nih.gov/articles/PMC11435893/ [Accessed 17 March 2025].
12. Diwan, T., Anirudh, G. and Tembhurne, J.V., 2022. Object detection using YOLO: challenges, architectural successors, datasets and applications. *Multimedia Tools and Applications*, 81(6), pp.9243–9275. [online] Available at: https://doi.org/10.1007/s11042-022-13644-y [Accessed 23 March 2025].
13. Dong, D., Jiang, H., Wei, X., Song, Y., Zhuang, X., and Wang, J., 2023. ETNAS: An energy consumption task-driven neural architecture search. *Sustainable Computing: Informatics and Systems*, [online] Available at: https://www.sciencedirect.com/science/article/abs/pii/S2210537923000811 [Accessed 27 January 2025].
14. FAO, n.d. *Plant Production and Protection Division*. [online] Available at: https://www.fao.org/plant-production-protection/about/en [Accessed 27 January 2025].
15. Flor, R.J., Yuan, S., Rizzo, G., Agus, F., Pame, A.R.P., Ouk, A., Then, R., Keo, S., Kumar, V. and Grassini, P., 2024. Accurate measurement of field size is essential for analysis of smallholder survey data. *Field Crops Research*, [online] Available at: https://www.sciencedirect.com/science/article/pii/S0378429024001369 [Accessed 8 June 2025].

16. Huang, H., Wu, Y., Yu, M., Shi, X., Qiao, F., Luo, L., Wei, Q., and Liu, X., 2020. EDSSA: An Encoder-Decoder Semantic Segmentation Networks Accelerator on OpenCL-Based FPGA Platform (image) *Sensors*, 20(14), p.3969. Available at: https://doi.org/10.3390/s20143969 [Accessed 17 March 2025].
17. Huang, S.-Y., Hsu, W.-L., Hsu, R.-J. and Liu, D.-W., 2022. Fully convolutional network for the semantic segmentation of medical images: A survey. *Diagnostics*, 12(11), p.2765. [online] Available at: https://doi.org/10.3390/diagnostics12112765 [Accessed 23 March 2025].
18. IBM, 2021. What is Overfitting? [online] Available at: https://www.ibm.com/think/topics/overfitting [Accessed 17 March 2025]
19. IPTSAT, 2024. High and Very High-Resolution Data. IPTSAT, [online] Available at: https://www.iptsat.com/en/high-and-very-high-resolution-data/ [Accessed 18 March 2025].
20. Jung, D.H., Kim, J.D., Kim, H.Y., Lee, T.S., Kim, H.S. and Park, S.H., 2022. A hyperspectral data 3D convolutional neural network classification model for diagnosis of gray mold disease in strawberry leaves. *Frontiers in Plant Science*, 13. [online] Available at: https://doi.org/10.3389/fpls.2022.837020 [Accessed 27 January 2025].
21. Iqbal, U., Davies, T. and Perez, P., 2024. A Review of Recent Hardware and Software Advances in GPU-Accelerated Edge-Computing Single-Board Computers (SBCs) for Computer Vision Sensors. *Sensors*, 24(15), p.4830. [online] Available at: https://doi.org/10.3390/s24154830 [Accessed 9 June 2025].
22. Khan, A.I. and Al-Habsi, S., 2020. Machine Learning in Computer vision. *Procedia Computer Science*, 167, pp.1444–1451. [online] Available at: https://doi.org/10.1016/j.procs.2020.03.355 [Accessed 23 March 2025]
23. Kim, D., 2024. Fully Convolutional Network Overview. *Computer vision*. Available at: https://www.doptsw.com/posts/post_2024-11-11_f23973 [Accessed 17 March 2025].
24. Koknaroglu, H., Demircan, V., Dalgic, A. and Cinar, I., 2021. Comparison of Holstein and Simmental cows in terms of performance, cost and profitability. [online] Available at: https://www.researchgate.net/publication/351711321_Comparison_of_Holstein_and_Simmental_cows_in_terms_of_performance_cost_and_profitability [Accessed 18 March 2025].
25. Krogstad, K.C. and Bradford, B.J., 2024. Associations of body condition score, body condition score change, and hyperketonemia with mastitis, reproduction, and milk production. *Journal of Dairy Science.* [online] Available at: https://doi.org/10.3168/jds.2024-25387. [Accessed 18 March 2025].
26. Li, X., Hu, Z., Huang, X., Feng, T., Yang, X., and Li, M., 2019. Cow Body Condition Score Estimation with Convolutional Neural Networks. *2019 IEEE 4th International Conference on Image, Vision and Computing (ICIVC)*, pp. 433-437. Available at: https://www.semanticscholar.org/paper/Cow-Body-Condition-Score-Estimation-with-Neural-Li-Hu/15f7d28ad100ca50e1188b615331f13abf785186 [Accessed 17 March 2025].
27. Lin, C.-H., Chen, T.-Y., Chen, H.-Y. and Chan, Y.-K., 2024. Efficient and lightweight convolutional neural network architecture search methods for object classification. *Pattern Recognition*, 156, p.110752. . [online] Available at: https://doi.org/10.1016/j.patcog.2024.110752 [Accessed 23 March 2025].
28. Lu, G. and Fei, B., 2014. Medical hyperspectral imaging: a review (image). *Journal of Biomedical Optics*. [online] Available

at: http://biomedicaloptics.spiedigitallibrary.org/article.aspx?articleid=1816617 [Accessed 27 January 2025].
29. Maxar Technologies, 2025. Clarity and Confidence: How High-Resolution Satellite Imagery Can Help You Extract Critical Information at Scale. Available at: https://go.maxar.com/Imagery-Leadership-Spatial-Resolution [Accessed 17 March 2025].
30. Megalac, 2024. Everything you need to know about body condition scoring dairy cows. *Megalac*, [online]. Available at: https://www.megalac.com/about/news/352-everything-you-need-to-know-about-body-condition-scoring-dairy-cows [Accessed 18 March 2025].
31. Minami, D. and Nishikawa, K., 2024. YOLO for High Resolution Images without Retraining. *Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, [online] Available at: https://www.apsipa.org/ [Accessed 18 March 2025].
32. Montesinos López, O.A., Montesinos López, A. and Crossa, J., 2022. Convolutional Neural Networks. In: O.A. Montesinos López, A. Montesinos López and J. Crossa, eds. *Multivariate Statistical Machine Learning Methods for Genomic Prediction*. Cham: Springer, [online] Available at: https://link.springer.com/chapter/10.1007/978-3-030-89010-0_13 [Accessed 23 March 2025]
33. Morgunov, A., 2024. Object Detection with YOLO: Hands-on Tutorial. Neptune Blog (image), [online] 28 October. Available at: https://neptune.ai/blog [Accessed 18 March 2025].
34. Musyoka, G.M., 2018. Automatic delineation of smallholder agricultural field boundaries using fully convolutional networks. [pdf] Available at: https://library.itc.utwente.nl/papers_2018/msc/gfm/musyoka.pdf [Accessed 17 March 2025].
35. Nagasubramanian, K., Jones, S., Singh, A.K., Singh, A., Ganapathysubramanian, B., and Sarkar, S., 2019. Plant disease identification using explainable 3D deep learning on hyperspectral images. *Plant Methods*. [online] Available at: https://plantmethods.biomedcentral.com/articles/10.1186/s13007-019-0479-8 [Accessed 27 January 2025].
36. National Geographic, 2025. What the world eats. [online] Available at: https://www.nationalgeographic.com/what-the-world-eats/ [Accessed 30 January 2025].
37. Parab, C.U., Mwitta, C., Hayes, M., Schmidt, J.M., Riley, D., Fue, K., Bhandarkar, S. and Rains, G.C., 2022. Comparison of single-shot and two-shot deep neural network models for whitefly detection in IoT web application. *AgriEngineering*, 4(2), pp.507–522. [online] Available at: https://doi.org/10.3390/agriengineering4020034 [Accessed 23 March 2025].
38. Pepe, M., Alfio, V.S., Costantino, D. and Herban, S., 2022. *Rapid and accurate production of 3D point cloud via latest-generation sensors in the field of cultural heritage: A comparison between SLAM and spherical videogrammetry*. Heritage, 5(3), pp.1910–1928. [online]
39. Pesticide Collaboration, 2025. The Problem with Pesticides. [online] Available at: https://pesticidecollaboration.org/the-problem-with-pesticides/#:~:text=Recent%20findings%20suggest%20that%2041,wildlife%2C%20and%20as%20crop%20pollinators[Accessed 27 January 2025].
40. Polder, G. and Pekkeriet, E., 2013. *A Spectral Imaging System for Detection of Botrytis in Greenhouses*. EFITA-WCCA-CIGR Conference "Sustainable Agriculture through ICT Innovation" [online] Available at:

https://www.researchgate.net/publication/259215766_A_Spectral_Imaging_System_for_Detection_of_Botrytis_in_Greenhouses [Accessed 27 January 2025]
41. Pound, M., 2018. Encoder Decoder Network - *Computerphile*. [online] Available at: https://www.youtube.com/watch?v=1icvxbAoPWc [Accessed 17 March 2025]
42. Prior, M., 2024. More cattle kept in UK 'megafarms', BBC finds. *BBC News*, [online] 14 August. Available at: https://www.bbc.co.uk/news/articles/cy4ldkpz1klo [Accessed 18 March 2025].
43. Ritchie, H., Rosado, P., and Roser, M., 2017. *Meat Production*. *Our World in Data*. Available at: https://ourworldindata.org/meat-production [Accessed 17 March 2025]
44. Shoaib, M., Shah, B., El-Sappagh, S., Ali, A., Ullah, A., Alenezi, F., Gechev, T., Hussain, T. and Ali, F., 2023. An advanced deep learning models-based plant disease detection: A review of recent research. *Front Plant Sci.*, [online] Available at: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10070872/ [Accessed 6 March 2025]
45. Singh, A., 2020. Selecting the Right Bounding Box Using Non-Max Suppression (with implementation). Analytics Vidhya, [online] Available at: https://www.analyticsvidhya.com/blog/2020/08/selecting-the-right-bounding-box-using-non-max-suppression-with-implementation/ [Accessed 18 March 2025].
46. Solawetz, J., 2024. What is Instance Segmentation? A Guide. *Roboflow Blog*. [online] Available at: https://blog.roboflow.com/instance-segmentation/ [Accessed 17 March 2025].
47. The Economist (2025) 'Spy-satellite-grade images could soon become available to everyone', *The Economist*, 25 February. Available at: https://www.economist.com/science-and-technology/2025/02/25/spy-satellite-grade-images-could-soon-become-available-to-everyone [Accessed: 17 March 2025].
48. United Nations, 2024. *The 17 Goals – Sustainable Development*. [online] United Nations Sustainable Development Goals. Available at: https://sdgs.un.org/goals [Accessed 23 March 2025].
49. Wagner, M.P. and Natascha, 2020. Extracting agricultural fields from remote sensing imagery using graph-based growing contours. *Remote Sensing for Crop Mapping*, 12(7), p.1205. [online] Available at: https://doi.org/10.3390/rs12071205 [Accessed 8 June 2025].
50. Wang, C.-Y. and Liao, H.-Y.M., 2024. YOLOv1 to YOLOv10: The fastest and most accurate real-time object detection systems. *arXiv*, [online] 18 August. Available at: https://arxiv.org/html/2408.09332v1 [Accessed 18 March 2025].
51. Wang, D., Vinson, R., Holmes, M., Seibel, G., Bechar, A., Nof, S. and Tao, Y., 2019. Early detection of Tomato Spotted Wilt Virus by hyperspectral imaging and outlier removal auxiliary classifier generative adversarial nets (OR-AC-GAN). *Scientific Reports Nature*, 9, Article number: 4377. [online] Available at: https://www.nature.com/articles/s41598-019-40975-7 [Accessed 8 June 2025].
52. Wołk, K. and Tatara, M.S., 2024. A review of semantic segmentation and instance segmentation techniques in forestry using LiDAR and imagery data. *Electronics*, 13(20), p.4139. [online] Available at: https://doi.org/10.3390/electronics13204139[Accessed 23 March 2025].
53. Xu, Y., Xue, X., Sun, Z., Gu, W., Cui, L., Jin, Y. and Lan, Y., 2023. Deriving agricultural field boundaries for crop management from satellite images using semantic feature pyramid network. (image) *Remote Sensing*, [online] 15(11), 2937. Available at: https://doi.org/10.3390/rs15112937 [Accessed 17 March 2025].

54. Yamashita, R., Nishio, M., Do, R.K.G. and Togashi, K., 2018. Convolutional neural networks: an overview and application in radiology. *Insights into Imaging*, 9. [online] Available at: https://doi.org/10.1007/s13244-018-0639-9 [Accessed 23 March 2025].
55. Yang, R., Ahmed, Z.U., Schulthess, U.C., Kamal, M. and Rai, R., 2020. Detecting functional field units from satellite images in smallholder farming systems using a deep learning-based computer vision approach: A case study from Bangladesh. *Remote sensing Applications: Society and environment*, [online] Available at: https://www.sciencedirect.com/science/article/pii/S2352938519301831 [Accessed 17 March 2025].