Computer Vision for Environmental Safety

Learning Outcomes

After completing this unit, students will be able to:

- Explain how computer vision contributes to environmental safety and monitoring.
- Describe key concepts of object detection and image classification in environmental contexts.
- Identify how drones, satellite data, and sensors are integrated for situational awareness.
- Apply a simple computer-vision model to classify satellite or drone images.
- Discuss how automated image analysis supports environmental decision-making and risk mitigation (e.g., pollution detection, wildfire monitoring, UXO detection, biodiversity tracking).

Introduction

Today, we are going to talk about computer vision not as an abstract AI concept from a machine learning textbook, but as a practical tool that is already being used to support environmental safety.

In simple terms, computer vision is about teaching machines to “see” - to analyze satellite imagery, drone footage, aerial photographs, or video streams, and to extract meaningful information from them. This is not just about recognizing objects in an image, but about detecting changes, patterns, and risks at a scale that is impossible to handle manually.

In environmental and safety-related applications, this capability becomes critical. Many of the areas we need to monitor are difficult or dangerous to access: industrial zones, remote forests, conflict-affected regions, or areas contaminated with unexploded ordnance. Traditional field inspections in such contexts are slow, expensive, and often unsafe.

Computer vision, combined with machine learning, allows us to continuously monitor these regions using visual data. Automated systems can process large volumes of satellite and drone imagery and identify early signals of environmental threats - such as pollution leaks, illegal logging, wildfire outbreaks, or structural damage - long before the situation escalates.

Throughout this lecture, we will look at how these systems work in practice: what kinds of data they rely on, which computer vision tasks are most relevant for environmental monitoring, and how the results of automated image analysis are translated into real-world decisions and risk mitigation actions. Figure 7.1 shows how computer vision differs from regular human vision.

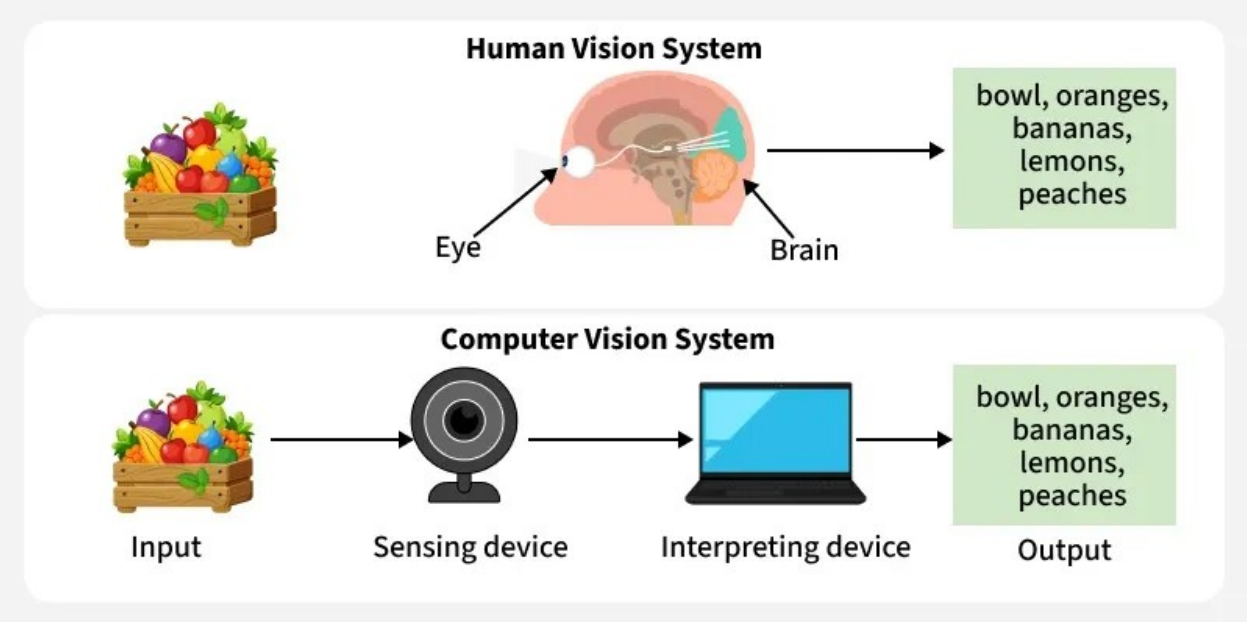

Figure 7.1 Computer vision schematic

The Role and Significance of Computer Vision

When we talk about the role of computer vision in environmental monitoring, the key idea is scalability. Traditional monitoring methods rely heavily on people - inspectors, field surveys, on-site measurements. These approaches work, but they do not scale well and they often fail exactly where monitoring is most needed: in large, remote, or dangerous areas.

Computer vision provides a more scalable and objective alternative. Instead of relying on human observation, we transform raw pixels from images into structured information that can be analyzed automatically and consistently over time.

At the core of most environmental computer vision systems are two fundamental tasks: image classification and object detection. Although they are often mentioned together, they answer different practical questions.

Image classification focuses on *what is present* in an image. For example, a model can determine whether a satellite image represents forest, water, urban areas, or degraded land. This type of analysis is especially useful when we need to monitor large regions and track changes over time.

Object detection, in contrast, focuses on *where exactly something is located*. It identifies individual objects or features and their boundaries - such as fallen trees, waste accumulations in rivers, or damaged buildings. This level of detail is essential when results must be linked to specific locations for field intervention or decision-making.

By combining these tasks, computer vision systems provide situational awareness in near real time. Models can be trained to recognize spectral patterns associated with oil spills, smoke signatures from early-stage wildfires, or structural changes visible in satellite and radar imagery.

The practical significance of this approach is that it shortens the time between environmental change and response. Instead of reacting after damage becomes visible on the ground, authorities and organizations can receive early warnings and act before the situation escalates.

Data Acquisition and Integration Strategies

The effectiveness of computer vision in environmental monitoring depends directly on the data it receives. In practice, no single data source is sufficient. Environmental safety systems rely on combining several types of visual data, each addressing different constraints of scale, resolution, and operating conditions.

Satellite imagery provides the foundation for large-scale monitoring. Platforms such as Sentinel and Landsat enable continuous observation of vast territories over long periods of time. Multispectral data, especially in the near-infrared and short-wave infrared bands, allows us to assess vegetation health, detect deforestation, and identify land-use changes that are not visible in standard RGB imagery.

Drones complement satellite data by offering much higher spatial resolution and flexible deployment. They are typically used for localized inspections, such as identifying small pollution sources or detecting unexploded ordnance in areas where ground access is unsafe as shown on Figure 7.2.

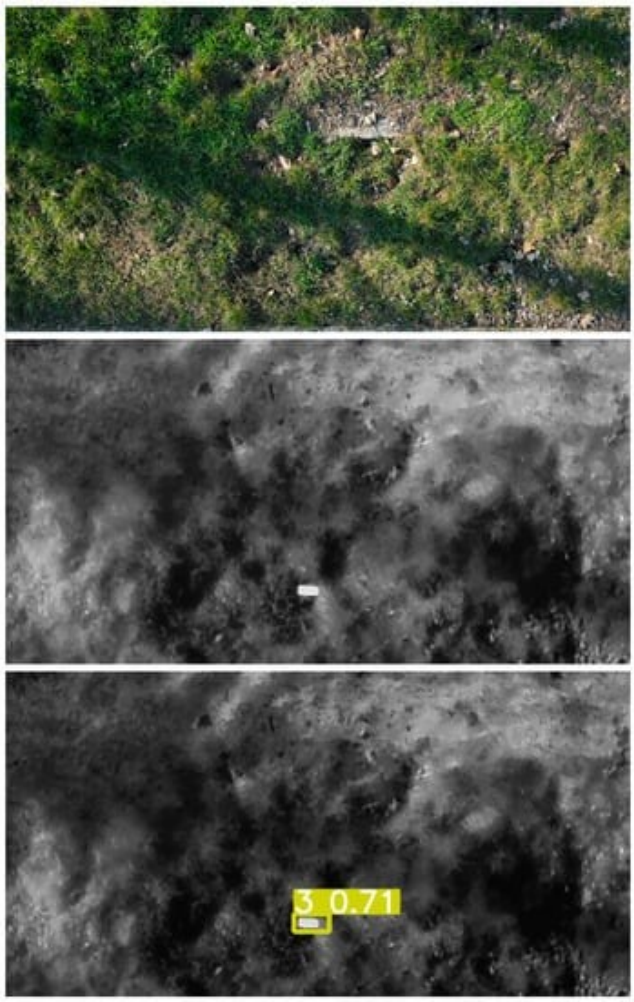

Figure 7.2 Drone with computer vision

Radar-based systems, particularly Synthetic Aperture Radar, address the limitations of optical imagery. Because SAR does not depend on sunlight and can penetrate clouds and smoke, it enables continuous monitoring regardless of weather or time of day. SAR data is especially valuable for detecting structural changes and building damage. Figure 7.3 displays an example of SAR image.



Figure 7.3 SAR image

In practice, these data sources are used together. The outputs of computer vision models are integrated with Geographic Information Systems, where detected risks and changes are mapped spatially. This integration transforms raw imagery into actionable information that supports environmental decision-making and risk mitigation.

Ethical, Safety, and Data Considerations

Before we go deeper into models and metrics, we need one practical checkpoint: environmental computer vision is not only a technical topic. The way we collect, store, and share visual data can create real ethical and safety risks.

The first issue is privacy and data responsibility. Satellite and drone imagery keeps improving in resolution. That means the same datasets that are useful for monitoring forests, rivers, or industrial sites can also unintentionally capture sensitive human activity. So the question is not only *“Can we detect it?”* but also *“Should we collect or publish it at this level of detail?”*

The second issue is dual-use risk. Environmental data is often assumed to be harmless, but that’s not always true. For example, precise locations of endangered species, or detailed information about infrastructure vulnerabilities, can be exploited. So practitioners need to balance scientific transparency with safeguards that prevent misuse.

The takeaway is simple: good governance is part of good engineering. Responsible access control, careful sharing policies, and awareness of dual-use scenarios are essential if we want computer vision to strengthen environmental resilience rather than create new risks.

Core Concepts & Methods

Conceptual Framework of Computer Vision in Environmental Safety

Now that we have looked at the data - where it comes from and what limitations it has - we can talk about how computer vision actually extracts information from these images.

In environmental monitoring, computer vision is not about building the most complex model possible. It is about choosing methods that are robust, interpretable, and well aligned with the properties of the data.

Most environmental computer vision pipelines are built around a small set of core ideas. At a high level, these methods take raw visual input - satellite pixels, drone images, or radar signals - and transform them into structured representations of the environment: land-cover maps, detected objects, or change indicators.

The important point here is that the choice of method is driven by the task and the data, not by fashion or model complexity. Large-scale monitoring, change detection, and safety-critical applications require approaches that are stable over time and tolerant to noise, incomplete labels, and imperfect ground truth.

With this in mind, we can now look at the two most common computer vision tasks used in environmental safety: image classification and object detection.

Image Classification versus Object Detection

In environmental computer vision, most practical systems rely on two closely related tasks: image classification and object detection. Although they are often mentioned together, they serve different purposes and answer different questions.

Image classification focuses on *what is present* in an image. The model assigns labels to an entire image or to predefined regions, allowing us to understand overall patterns - such as land cover types, vegetation condition, or damaged versus undamaged areas. This approach works particularly well for large-scale monitoring, where the goal is to track changes over time rather than individual objects. Figure 7.4 displays illustrates the example of image segmentation approach.

Object detection, in contrast, focuses on *where exactly something is located*. Instead of labeling regions, the model identifies individual objects and their precise positions. This level of detail becomes essential when analysis must be linked to concrete actions on the ground - for

example, locating unexploded ordnance, identifying specific waste accumulations, or mapping individual damaged buildings. Figure 7.5 displays the result of object detection algorithm applied to cars in the parking lot.

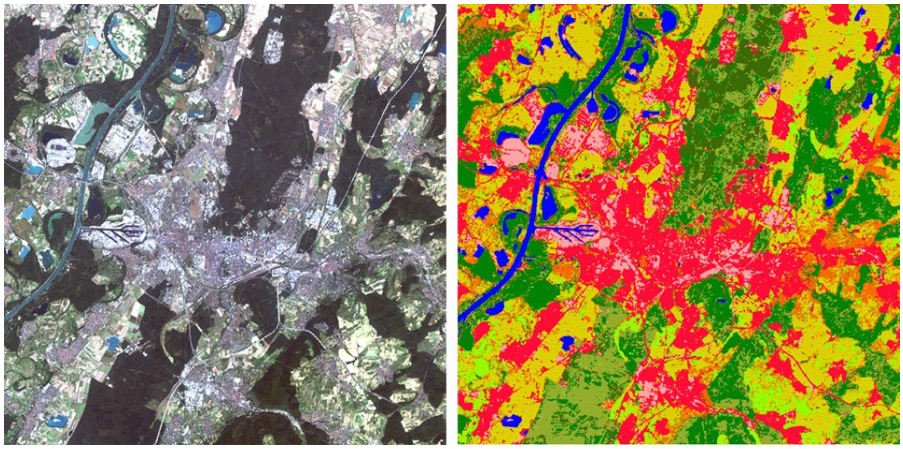

Figure 7.4 Example of image segmentation

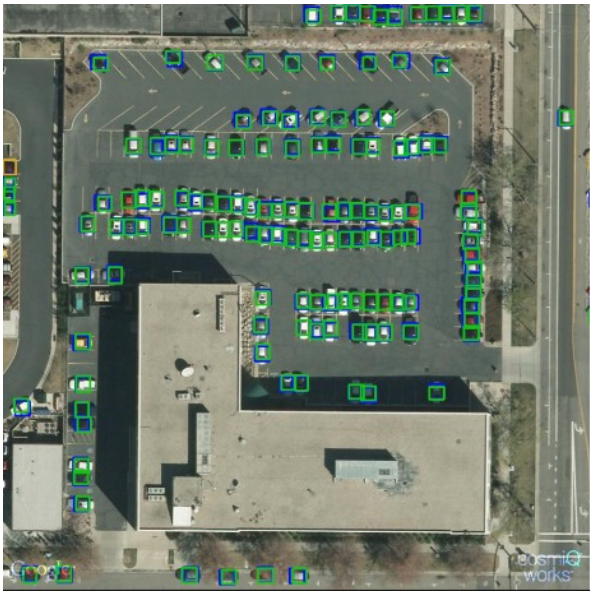

Figure 7.5 Object detection

The key difference between these approaches is not technical complexity, but intent. Classification helps us understand patterns at scale, while object detection enables targeted intervention. In many real-world environmental monitoring systems, both methods are combined: classification is used to identify areas of interest, and object detection is applied to analyze those areas in detail.

Computer Vision Model Concepts for Environmental Safety

When we talk about models in environmental computer vision, the main question is not “what is the most advanced architecture,” but “what can we deploy reliably with the data and constraints we have.” Environmental imagery is noisy, ground truth is expensive, and conditions change over time - seasons, lighting, cloud cover, and even sensor types.

In practice, model choice is driven by the output we need: segmentation or detection

Segmentation: U-Net

If the goal is to map *areas* - for example the extent of floodwater, forest boundaries, or damaged zones - segmentation is the natural choice.
A widely used architecture here is U-Net, because it is designed to produce pixel-level masks and capture both local details and broader spatial context. Figure 7.6 shows exactly why it is called U-Net. In environmental monitoring, that makes it useful whenever we care about *where the boundary is* and *how large the affected area is*.

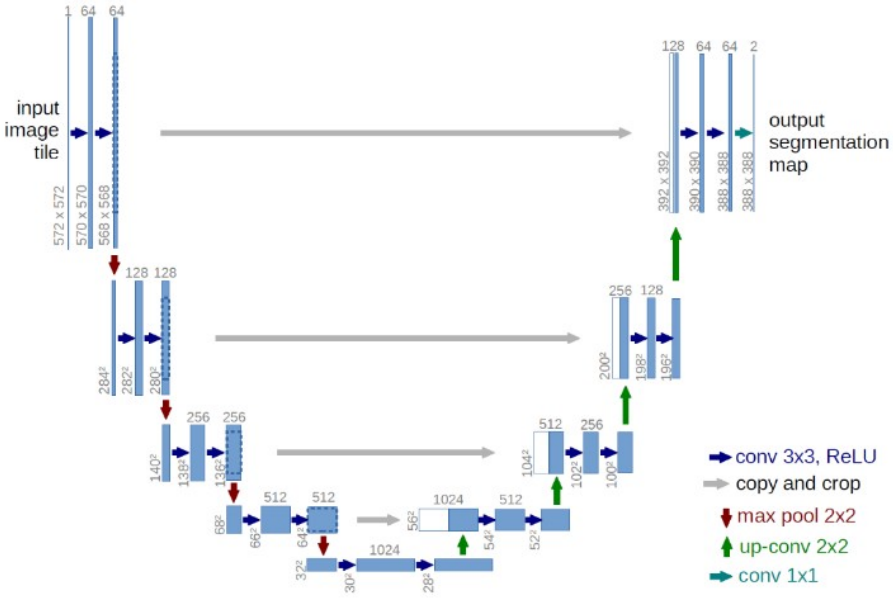

Figure 7.6 U-Net model architecture

Detection: MobileNet / YOLO-lite

If the goal is to find *individual objects* - such as pieces of waste, specific damaged structures, or UXO candidates - we typically use object detection models.
In many operational settings, especially on drones or edge devices, we prefer lightweight architectures such as MobileNet-based detectors or YOLO-lite, because they balance speed and accuracy and can run under limited compute.

Pre-trained models and transfer learning

A recurring challenge in environmental safety is limited labeled data. In hazardous or inaccessible regions, collecting and labeling imagery is costly or impossible. This is why transfer learning is so common: we start with a model pre-trained on a large general dataset and then fine-tune it on a smaller environmental dataset - for example adapting an urban imagery model to recognize specific building damage patterns, or adapting a detector to identify UXO-like objects from drone imagery.

Combining complementary algorithms

Environmental pipelines often combine multiple algorithmic approaches to improve robustness.

For example, Random Forest models can reduce “salt-and-pepper” noise in large homogeneous regions, which is useful for land-cover classification. Meanwhile, simpler neural models such as an MLP can sometimes separate minor or rare classes that are underrepresented - for instance subtle industrial anomalies or small-scale changes that get lost in broader classes.

The practical idea is straightforward: rather than betting everything on one model, we combine methods that fail in different ways.

Key takeaway

So, the model “concept” here is not a single architecture - it is an operational design choice. U-Net is a strong default for segmentation maps, MobileNet / YOLO-lite are practical choices for fast detection, and transfer learning is often the difference between “we can build it” and “we have no dataset to start with.”

Statistical Performance Metrics

Once we have defined what our models are supposed to do - classify regions or detect individual objects - the next question is how we evaluate their performance. In environmental monitoring, this is not as straightforward as it may seem.

Many environmental datasets are highly unbalanced. For example, damaged areas may cover only a small fraction of a large region, while the majority remains unchanged. In such cases, simple accuracy can be misleading, because a model can achieve high accuracy while failing to detect the phenomena we actually care about.

For this reason, environmental computer vision relies on metrics that better reflect spatial agreement and detection quality.

Intersection over Union (IoU)

One of the most widely used metrics for segmentation and area-based analysis is Intersection over Union, or IoU.

IoU measures how well the predicted area overlaps with the reference, or ground-truth, area as shown in Figure 7.7.

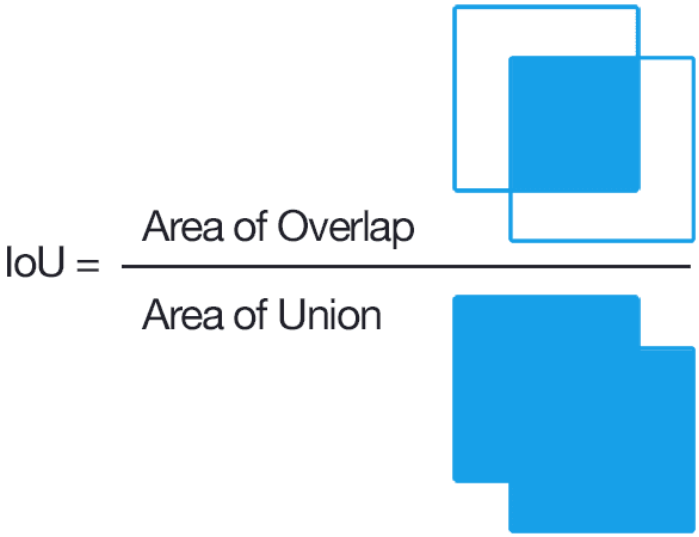

Figure 7.7 Intersection over Union calculation illustration

Intuitively, IoU answers a simple question: *how much do these two areas overlap?*
A higher IoU indicates stronger spatial agreement between the model output and expert-labeled data. In tasks such as deforestation mapping or building damage assessment, IoU values close to one indicate high-quality predictions.

Precision, Recall, and F1-Score

When we move from area-based segmentation to object detection, additional metrics become important.

Precision tells us how many of the detected objects are actually correct. This is critical in safety-related tasks, where false positives can lead to unnecessary or even dangerous interventions.

Recall measures how many relevant objects the model successfully detects. Missing an object, such as unexploded ordnance or early wildfire smoke, can have serious consequences.

The **F1-score** combines precision and recall into a single value, balancing false positives and false negatives. This makes it particularly useful in environmental applications, where class imbalance is common and both types of errors matter.

Why these metrics matter in practice?

In environmental monitoring, metrics are not just numbers for model comparison. They influence real-world decisions.

For example, increasing detection thresholds may improve precision but reduce recall. The appropriate balance depends on the task: environmental safety applications often prioritize minimizing missed detections, even at the cost of accepting some false alarms.

Understanding these trade-offs is essential for deploying computer vision systems responsibly and effectively.

Factors Influencing Measured Efficiency

When we talk about model performance metrics, it is important to understand one key point: the measured efficiency of a computer vision system is not determined by the model alone.

In environmental monitoring, performance is strongly influenced by the properties of the data and by the conditions under which that data is collected.

Spatial Resolution Differences

One common limitation comes from differences in spatial resolution. Ground-truth labels are often created using very high-resolution imagery or manual interpretation, while models may operate on moderate-resolution satellite data.

As a result, even a well-performing model rarely achieves perfect overlap with reference data. This is not necessarily a modeling error, but a consequence of comparing information captured at different spatial scales.

Environmental Noise and Atmospheric Conditions

Optical imagery is highly sensitive to environmental noise. Clouds, shadows, smoke, and varying illumination conditions can all affect image quality and reduce model performance.

Radar-based data, such as SAR, helps address some of these limitations by providing weather-independent observations. However, radar imagery introduces its own challenges, including higher signal noise and less intuitive visual patterns, which can complicate object boundaries and classification.

Temporal Stability and Seasonal Effects

Another important factor is temporal variability. Environmental conditions change over time due to seasons, weather patterns, and land-use dynamics.

To improve stability in long-term monitoring, techniques such as normalized difference indices are often used. These indices reduce the influence of illumination differences and seasonal variation, making it easier to detect meaningful environmental changes rather than temporary fluctuations.

Key takeaway

The main takeaway is that performance metrics must always be interpreted in context. In environmental computer vision, lower metric values do not automatically imply poor model quality - they often reflect real-world data limitations.

Understanding these factors helps practitioners set realistic expectations and design monitoring systems that are robust rather than artificially optimized.

Applications in Environmental Risk Mitigation

After discussing data, methods, metrics, and their limitations, we can now focus on the most practical question: *how is computer vision actually used to reduce environmental risk?*

In real-world scenarios, automated image analysis supports decision-making by detecting risks early, across large areas, and often in places where direct human inspection is difficult or dangerous.

Pollution and Waste Monitoring

One of the most common applications is pollution monitoring. Computer vision models can detect oil spills on water surfaces or identify illegal waste dumping in rivers and remote areas.

The key advantage here is scale. Large territories can be monitored continuously, allowing authorities to focus inspections and cleanup efforts on the most critical locations rather than relying on sporadic field reports [7].

Unexploded Ordnance (UXO) Detection

In conflict-affected and post-conflict regions, computer vision contributes directly to human safety. By analyzing drone imagery and high-resolution satellite data, models can help identify areas that potentially contain unexploded ordnance.

These systems do not replace demining experts. Instead, they reduce risk by narrowing down search areas and prioritizing zones that require immediate attention.

Wildfire and Smoke Detection

Computer vision is also used for early wildfire detection. Models can recognize smoke patterns or thermal anomalies before fires become large and uncontrollable.

Even small improvements in detection time can have a significant impact, reducing environmental damage, response costs, and risks to nearby communities.

Building Damage Assessment

Another critical application is damage assessment after natural disasters or during armed conflict.

Radar-based systems, particularly SAR, are especially valuable in these scenarios because they operate independently of weather and lighting conditions. Structural damage causes measurable changes in radar signal reflections, allowing damage to be detected even when optical imagery is unavailable. Figures 7.8 and 7.9 illustrate how damaged assessment can be visualized.



Figure 7.8 Build damage assessment after earthquake in Turkey



Figure 7.9 Damages of the city Rubizhne (Ukraine) due to combat operations

These outputs are typically integrated into GIS platforms, where they support emergency response, humanitarian aid planning, and recovery prioritization [9].

Deforestation and Biodiversity Monitoring

Long-term environmental risks such as deforestation and biodiversity loss are monitored by comparing satellite imagery over time.

Computer vision enables consistent, repeatable detection of gradual changes in land cover, supporting conservation efforts and policy decisions - especially in regions where continuous field monitoring is not feasible [6], [7], [8]. The result of such a monitoring is displayed in Figure 7.10.

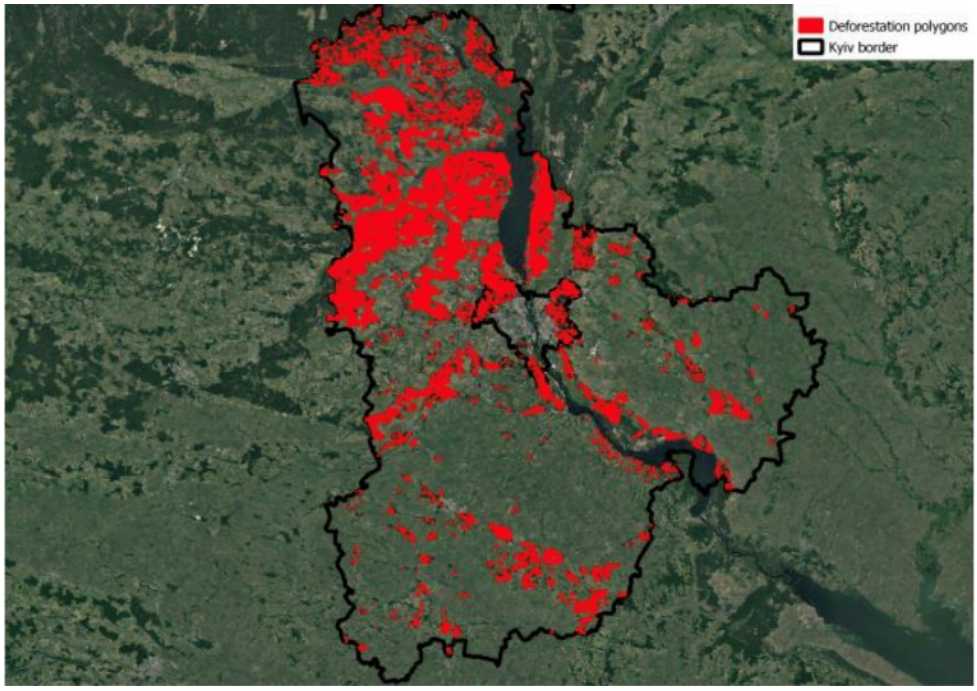

Figure 7.10 Deforestation detection visualization

Key takeaway

Across all these applications, the role of computer vision is not to replace environmental experts, but to extend their capacity. Automated analysis makes monitoring faster, safer, and more systematic, allowing human expertise to be applied where it matters most.

Data Sources and Integration for Environmental Monitoring

In practice, environmental monitoring systems are built as multi-layered pipelines rather than relying on a single data source.

Satellite imagery provides the top layer - continuous, large-scale coverage that allows entire regions to be monitored consistently over time. This layer is responsible for early detection and trend analysis.

Once areas of interest are identified, higher-resolution data sources are used. Drones and aerial imagery provide detailed, localized views that support inspection, verification, and precise mapping.

Radar-based data plays a complementary role by ensuring continuity. When optical data is unavailable due to clouds, smoke, or darkness, SAR imagery allows monitoring to continue without interruption.

The integration point for all these data streams is Geographic Information Systems. GIS provides the spatial framework that aligns satellite observations, drone inspections, and model outputs into a single operational picture.

This layered approach allows environmental monitoring systems to balance scale, detail, and reliability, transforming heterogeneous data into actionable situational awareness.

From Detection to Environmental Action

The implementation of computer vision (CV) in environmental safety provides a technical framework for transforming raw visual data into actionable spatial intelligence. By automating the interpretation of digital photographs, satellite maps, and drone imagery, these systems facilitate the early detection of ecological threats, the monitoring of hazardous regions, and the enhancement of data-driven decision-making processes. The transition from visual identification to mitigation involves the synthesis of Artificial Intelligence (AI), Remote Sensing, and Geographic Information Systems (GIS) to create automated monitoring loops that save time and reduce human risk.

Identification of Environmental and Humanitarian Hazards

Computer vision models provide situational awareness by answering critical questions regarding the presence and location of environmental risks. Automated systems are deployed across various domains to identify specific threats that require intervention:

- **Pollution and Waste:** AI-instrumented camera systems and drones track plastic pollution in river systems with reported high detection accuracies under controlled conditions. This technology identifies characteristic spectral absorption peaks of plastics, allowing for the detection of contamination plumes.
- **Ecological Degradation:** Automated processing of multispectral imagery, such as Sentinel-2 data, enables the detection of deforestation in protected forests [6], [7], [8]. Systems utilize indices such as NDVI to track felling dynamics and assess vegetation health.
- **Conflict-Related Risks:** In areas where direct access is restricted due to safety concerns, CV is used to identify unexploded ordnance (UXO) and quantify structural damage to buildings and infrastructure [9].
- **Early Warning Systems:** Modern AI frameworks automatically analyze visual data for the early detection of wildfire smoke or industrial leaks, generating alerts for communities before ecological damage reaches an unmanageable point of escalation.

Supporting Strategic Recovery and Resource Allocation

The integration of CV outputs with GIS layers allows stakeholders to visualize the spatial distribution of risks, which is essential for strategic planning and the prioritization of recovery efforts.

- **Cleanup and Remediation:** In scenarios involving aquatic pollution, automated mapping enables organizations to identify garbage accumulation zones in remote areas. This spatial data is used for planning targeted cleanup expeditions and extraction activities.
- **Infrastructural Restoration:** Building damage detection systems provide estimates of the number of affected structures and pre-war population densities in specific zones. These "Damage Proxy Maps" assist rescue operations and humanitarian aid organizations in determining where resources are most critically needed.
- **Source Identification:** By georeferencing pollution plumes and tracing them back to their primary sources on a map, CV supports industrial inspections and the enforcement of environmental regulations.

Evidence-Based Policy and Economic Support

Computer vision serves as an objective tool for national-level decision-making, particularly when traditional statistical data collection is hindered by conflict or environmental disasters.

- **Agricultural Policy Support:** High-resolution land-use maps are utilized to identify uncultivated or abandoned agricultural lands. In regions like Ukraine, these maps have been integrated into national registries to support government decisions, such as granting tax waivers to farmers whose land is contaminated or near front lines [4].
- **Food Security Monitoring:** Long-term monitoring of crop types and productivity through satellite imagery helps policymakers address food security challenges by identifying shifts in agricultural dynamics.
- **Post-War Reconstruction:** Continuous time-series data provides a spatial blueprint for reforestation and reconstruction efforts in areas deemed unsafe for human activity [9].

From Detection to Environmental Action: Real-World Examples of Computer Vision in Environmental Safety

When a computer detects an environmental threat, action must follow. This section examines how Computer Vision transforms detection into real environmental decisions—across three

critical phases: preventing threats before they escalate (*pre-event*), monitoring while crises unfold (*during-event*) and assessing damage and planning recovery (*post-event*).

Phase 1: Pre-Event Detection and Early Warning

Water Quality Monitoring with HydroVision

Imagine authorities want to know whether river water is safe before people get sick. Traditionally, they collect water samples manually - expensive, slow, and often too late.

HydroVision is a system that analyzes photographs of water (taken by drones or satellites) and reveals contamination or water quality degradation. It searches for color, turbidity, and algal presence - telltale signs of pollution:

- Green water = algal bloom (potentially toxic)
- Brown turbidity = suspended organics (factory discharge)
- Color anomaly = warning signal

Instead of one person at the riverbank, the system analyzes hundreds of kilometers of water bodies daily. By applying deep convolutional neural networks to RGB images, HydroVision estimates optically active parameters such as chlorophyll concentration and colored dissolved organic matter (CDOM), which serve as proxies for water quality and contamination trends. This non-invasive approach reduces dependency on costly spectral sensors and accelerates initial threat identification in diverse water bodies [12].

Biodiversity and Habitat Monitoring

Computer Vision also supports pre-event readiness by enabling automated environmental monitoring at scale—identifying biodiversity patterns, tracking species distributions, and mapping habitat conditions. Vision-based systems can detect land cover changes, species presence, and habitat features, contributing to broader climate action and biodiversity protection efforts. These automated detections provide important baseline datasets that reveal long-term trends which might otherwise go unnoticed until they manifest as acute threats [13].

Phase 2: During-Event Monitoring in Real Time

Real-Time Monitoring During Wildfires

A similar principle applies during forest fires:

- 🔥 Drones scan the region, searching for smoke plumes and thermal hotspots,
- 🗺️ Fire is located at the forest edge—it can be stopped before spreading,
- 📊 Firefighters receive not “somewhere there’s a fire” but exact ignition coordinates

Modern Computer Vision architectures, especially when combined with Internet of Things (IoT) sensors, allow for continuous real-time surveillance of dynamic environmental events. Hybrid sensing systems integrate cameras with analytical models to maintain robust tracking and detection even in challenging conditions—variable lighting or obscured views - enabling effective detection of changes or anomalies in ecosystems under stress [14].

Detecting Aquatic Plastic Pollution

When rivers become choked with plastic waste, Computer Vision can pinpoint every piece floating in the water:

- 🎯 Automatic detection of plastic water bottles in river environments,
- 📍 Precise coordinates - exactly where debris is drifting,
- ⏱️ Real-time updates to contamination maps.

Cleanup workers and volunteers know exactly where to search—no wasted effort in the dark. Researchers have developed models that automatically detect plastic water bottles in river environments by training on extensive datasets of aquatic imagery. These models can identify and localize plastic debris within complex natural scenes with high precision. Such capabilities allow authorities and environmental organizations to monitor pollution hotspots continuously, improving situational awareness during pollution episodes and informing cleanup tactics and resource allocation in near real time [15].

Phase 3: Post-Event Assessment and Recovery

Burned-Area Mapping After Wildfires

After a forest fire, flood, or storm, authorities need to know:

- Where exactly burned? (Which areas require restricted access?)
- How severe? (Are trees simply leafless gray, or charred?)
- How many resources for recovery? (Personnel, equipment, time?)

Traditionally, people hiked through forests for months, measured manually, and consulted locals - slow and expensive.

Deep learning models on satellite imagery complete this work in days. Deep learning models applied to satellite imagery, such as U-Net and other convolutional architectures, allow for automated burned-area mapping following wildfires. These systems use pre-fire and post-fire Sentinel-2 imagery to classify burn severity and delineate fire scars with high accuracy, providing objective, reproducible assessments essential for planning rehabilitation strategies and allocating restoration resources [16].

The output includes:

- 📍 Burn maps: exactly where fire burned/is burning
- 🎨 Color-coded severity: yellow = light, orange = moderate, black = severe
- 📊 Reports: quantitative estimates of forest area requiring recovery

Post-Disaster Damage Assessment

The same approach works for floods, storms, and other catastrophes:

- 🏘️ Which buildings are flooded?
- 🛣️ Which roads are blocked?
- 👥 Which populations are at risk?

Computer Vision contributes to broader post-disaster environmental evaluation, where deep CNN models detect and categorize damage patterns in natural and built environments immediately after storms, floods, or other ecological disturbances. By analyzing large volumes of post-event aerial imagery, these models support rapid classification of affected zones, enabling faster decision-making and improving coordination between environmental agencies, emergency responders, and restoration programs [17].

Case Study: Vegetation Health Assessment Using NDVI from Sentinel-2 Data

Context

In this case study, multispectral satellite imagery had been utilized to assess vegetation health and to analyze spatial patterns of environmental degradation. The analysis had been grounded in the Normalized Difference Vegetation Index (NDVI), a well-established spectral indicator derived from the red and near-infrared (NIR) bands of Sentinel-2 imagery.

The case study had been designed to demonstrate how relatively simple computer vision and remote sensing techniques could yield actionable environmental insights when applied to real satellite data.

Environmental relevance

NDVI-based analysis had been widely adopted in environmental safety applications, including deforestation monitoring, agricultural land-use assessment, drought impact analysis, and post-

disturbance environmental evaluation. The use of satellite-derived spectral indices had enabled large-scale, repeatable observations over extended time periods.

This approach had been particularly effective in regions where direct field measurements had been limited due to safety constraints or accessibility challenges. By relying on orbital data, continuous environmental monitoring had been maintained without physical presence in potentially hazardous areas.

Data and Spectral Inputs

The analysis had been based on Sentinel-2 multispectral imagery. From the available spectral channels, the red (B4) and near-infrared (B8) bands had been selected due to their sensitivity to vegetation reflectance properties.

To support transparency and reproducibility, intermediate processing steps had been documented directly within the analysis workflow.

Code excerpt: Sentinel-2 band loading and preprocessing

```
images = (ee.ImageCollection('COPERNICUS/S2_SR_HARMONIZED')
            .filterBounds(aoi)
            .filterDate(start_date, end_date)
            .median()
            .clip(ai))
```

Method Overview

NDVI values had been computed using the standard normalized difference formulation:

$$NDVI = \frac{NIR - Red}{NIR + Red}$$

This formulation had emphasized contrast between vegetation-covered surfaces and non-vegetated areas by exploiting the differential reflectance behavior of chlorophyll in the red and near-infrared portions of the spectrum.

Code excerpt: NDVI calculation

```
def addNDVI(image):
    ndvi = image.normalizedDifference(['B8', 'B4']).rename('NDVI')
    return ndvi.copyProperties(image, ['system:time_start'])
```

Visualization and Spatial Representation

Following computation, NDVI maps had been generated to visualize spatial variations in vegetation condition.

In these maps, **low NDVI values** (typically below ~0.1) were represented by red or brown tones and corresponded to **non-vegetated surfaces**, such as water bodies, bare soil, or built-up areas. **Intermediate NDVI values** (approximately 0.1–0.3), visualized using yellow or orange tones, indicated **sparse or stressed vegetation**, often associated with transitional land-cover types or seasonal effects. **High NDVI values** (above ~0.3), shown in green tones, represented **dense and healthy vegetation**, reflecting areas with high photosynthetic activity.

These visual representations had served as the primary interface for environmental interpretation, enabling rapid identification of anomalous or degraded zones. Figure 7.11 shows an example of NDVI map visualization.

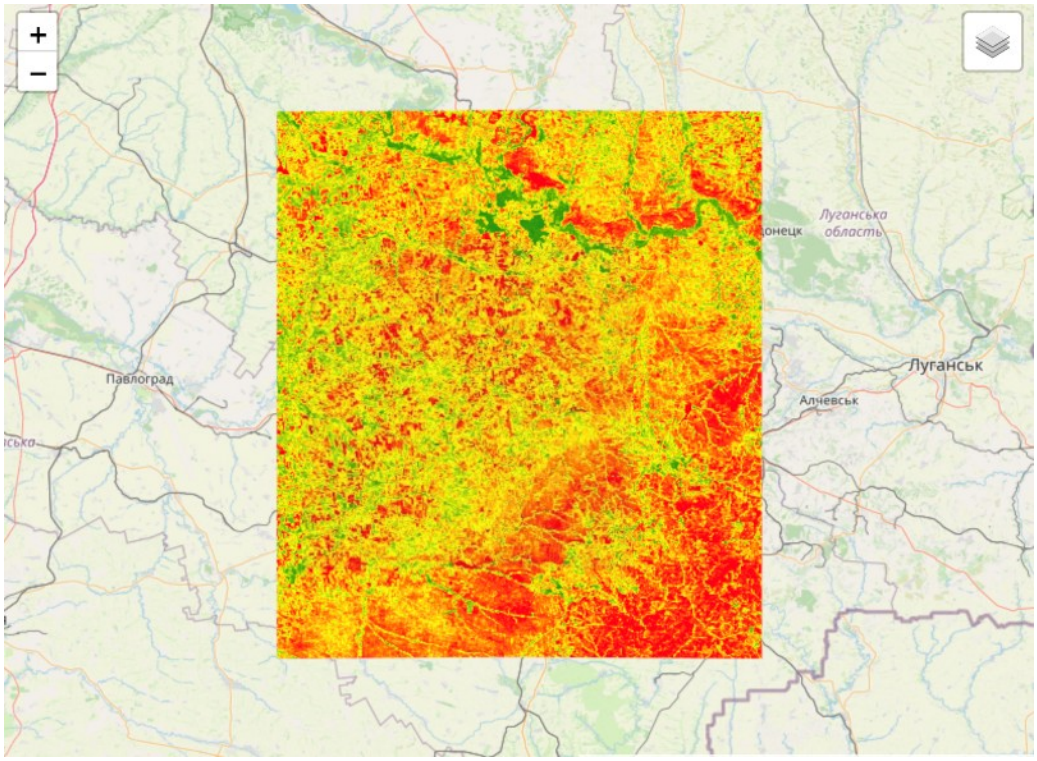

Figure 7.11 NDVI map visualization

Interpretation of NDVI Patterns

Threshold-based interpretation had been applied to categorize vegetation condition:

- NDVI values below approximately 0.1 had been associated with water bodies, bare soil, or severely degraded land,
- Intermediate NDVI values (approximately 0.1–0.3) had indicated sparse or stressed vegetation,
- Higher NDVI values had corresponded to dense and healthy vegetation cover.

Spatial clusters of persistently low NDVI values had been interpreted as potential indicators of environmental stress, anthropogenic disturbance, or land-use change.

Results and Environmental Implications

The resulting NDVI outputs had revealed distinct spatial patterns in vegetation health across the observed region. Extended areas of reduced vegetation activity had suggested possible deforestation, land abandonment, or post-disturbance ecological degradation.

By documenting intermediate results and visual outputs alongside computational steps, the case study had demonstrated how spectral analysis could support transparent environmental assessment and reproducible analysis pipelines.

Practical Exercise

Objective

To gain hands-on experience with satellite imagery analysis and understand how NDVI can be used for environmental safety and monitoring: compute an NDVI map, visualize vegetation health patterns, and interpret the results in an environmental context.

Step-by-Step Guide: How to Run the Notebook

The provided notebook “ML_CV_NDVI.ipynb” is a complete, fully functional workflow that is ready-to-run. It is pre-configured to **access Sentinel-2 data via Google Earth Engine**, compute NDVI composites, visualize spatial patterns of vegetation condition, and perform temporal comparisons. The notebook produces **interactive maps and visual outputs automatically** once executed. No additional coding is required to reproduce the results, although parameters such as the Area of Interest and time range can be modified for exploratory analysis.

We will use Google Colab to run the provided notebook in a browser environment (no local setup required), following the same “upload → run all → interpret outputs” flow as in the course template.

Notebook structure explained:

1. *Configuration & Libraries*. This section initializes the computational environment. Required Python libraries are imported, experiment parameters are defined in subsections:

   1a: The Area of Interest (AOI) is defined using GeoJSON format. By default, this is the territory of southeastern Ukraine, a region unfortunately engulfed in active combat operations.

   1b: Access to Google Earth Engine is configured.

   1c: Set experiment configuration parameters:

   - NDVI time window (NDVI_START_DATE and NDVI_END_DATE) representing the current analysis period.
   - NDVI historical baseline period (NDVI_HISTORICAL_START_DATE and NDVI_HISTORICAL_END_DATE) used for comparison.
   - Visualization parameters NDVI_VISUALIZATION_PARAMS.
   - Visualization parameters for NDVI difference DIFF_VISUALIZATION_PARAMS.

2. *Utility Functions*. This section contains reusable helper functions that encapsulate core operations such as:
   - building interactive Folium maps,
   - computing NDVI from Sentinel-2 imagery,
   - generating NDVI composites,
   - and supporting visualization logic.
3. *NDVI Composite Visualization*. In this section, a median NDVI composite is generated for the selected AOI and time period. The resulting NDVI map provides a spatial overview of vegetation condition and serves as the baseline representation of vegetation health during the selected season.
4. *NDVI Temporal Comparison*. This section performs temporal change detection by comparing NDVI composites from two different time periods. NDVI differencing is used to identify areas of vegetation improvement or degradation, enabling the detection of long-term environmental changes rather than isolated observations.
5. *Reflection Questions.* The final section focuses on qualitative interpretation. Students are guided to reflect on observed spatial and temporal patterns, consider

environmental drivers behind vegetation change, and discuss limitations and uncertainty in satellite-based analysis

Prerequisite: Google Earth Engine Setup

This practical exercise relies on Google Earth Engine for accessing and processing Sentinel-2 satellite imagery. Before running the notebook, Earth Engine access must be enabled for your Google account. You will find a detailed instruction in the notebook file in the **1b**. **Google Earth project setup** section.

After project is ready, update the EE_PROJECT variable as shown in Figure 7.12.

```
# Google Earth Engine project ID.
# This project must be enabled in Google Cloud Console and linked
# to the authenticated Google account to allow Earth Engine operations.
EE_PROJECT = "ml-cv-class-2026" # set your project ID here
```

Figure 7.12 EE_PROJECT variable definition

Notebook execution

When you are ready with the Google Earth Engine project setup, proceed to the notebook execution.

1. Open Google Colab.
2. Upload the notebook file.
3. Run the notebook: **Runtime** → **Run all**.
4. Scroll through outputs as cells execute. You should see intermediate results and final NDVI visualizations.
    - You will see NDVI visualization for the defined date range in section **3. NDVI Composite Visualization** as shown in Figure 7.13**.** Green regions contain healthy and dense vegetation, while red regions have less vegetation and more bare soil.

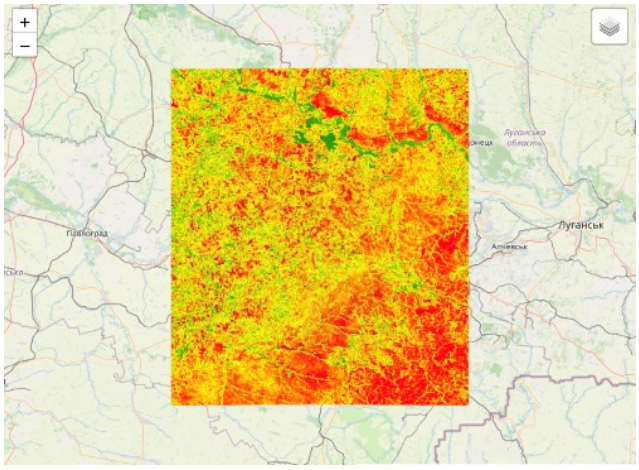

Figure 7.13 NDVI composite visualization

- And NDVI difference between defined date range and historical date range in section **4. NDVI Temporal Comparison.** Here green regions represent positive change in vegetation compared to the historical date range and red regions represent decrease in vegetation density. Figure 7.14 illustrates an example NDVI temporal comparison.

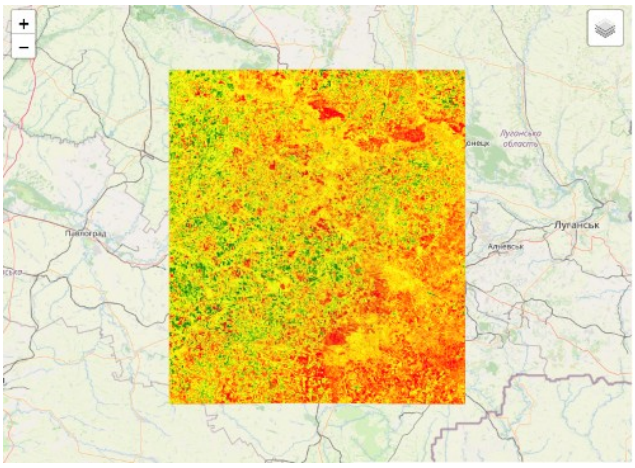

Figure 7.14 NDVI temporal comparison

Note: If Colab shows a warning (“Author not trusted”), click Run anyway.

Data Source

This exercise uses Sentinel-2 multispectral imagery. NDVI is computed from Red band (B4) to Near-Infrared band (B8).

The data is loaded from the Copernicus Sentinel-2 SR collection (e.g., *COPERNICUS/S2_SR_HARMONIZED*), filtered by area and date, then aggregated for analysis.

Tasks

Task 1: Run & Validate the Baseline Output

- Open and run the provided notebook.
- Confirm that Sentinel-2 data loads correctly for the selected AOI and date range.
- Compute the NDVI layer and render an NDVI map.

Questions:

- Where do you observe clusters of low and high NDVI values within the AOI?
- Which land-cover types are most likely associated with these clusters (e.g., water bodies, bare soil, urban areas, dense vegetation)?
- Do the NDVI patterns appear spatially coherent, or are they fragmented and noisy at the pixel level?
- What initial environmental interpretation can be made based on the spatial distribution of NDVI values?

Task 2: Threshold-Based Interpretation (Environmental Reasoning)

Use the NDVI interpretation thresholds and classify the AOI into “low / medium / high vegetation health”:

- NDVI < ~0.1
- NDVI ~0.1–0.3
- NDVI > ~0.3

Questions:

- Which NDVI-based vegetation health category (low, medium, or high) dominates the AOI, and why?
- Are low-NDVI areas spatially clustered or scattered? What does this suggest about potential environmental stress or land-use patterns?
- Which natural or anthropogenic factors could plausibly explain the presence of low-NDVI zones in this region?

Task 3: Sensitivity Check (Time & Location)

Change **one variable at a time**, re-run, and compare outputs:

- Change the **date range** (e.g., spring vs. late summer).
- Change the **AOI** to your home region or to the region you want to analyze (e.g., an urban edge versus an agricultural region)

Questions:

- How do NDVI patterns change when the observation period is modified?
- Which changes can be explained by seasonal vegetation dynamics, and which may indicate more persistent structural change?
- How does changing the AOI affect the interpretation of NDVI results, particularly when comparing urban-edge and agricultural regions?

Assessment

What is the main role of computer vision in environmental safety applications?

A) Replacing environmental experts with fully automated systems
B) Improving the visual quality of satellite images
C) Automatically extracting actionable information from large volumes of visual data
D) Manually labeling environmental images

Which computer vision task is best suited for identifying the exact location of illegal waste piles along a river?

A) Image classification
B) Object detection
C) Image enhancement
D) Data normalization

Why are drones often used together with satellite imagery in environmental monitoring systems?

A) Drones replace the need for satellite data
B) Drones provide higher-resolution, localized observations to complement large-scale satellite coverage
C) Drones eliminate the need for computer vision models
D) Satellite imagery cannot be used for environmental monitoring

What type of environmental information is most commonly derived from high NDVI values?

A) Urban infrastructure density

B) Presence of clouds or smoke
C) Dense and healthy vegetation
D) Areas affected by flooding

Why is Synthetic Aperture Radar (SAR) especially important for monitoring disasters such as floods or building damage?

A) It provides higher color resolution
B) It detects only vegetation changes
C) It works independently of weather conditions and daylight
D) It requires no integration with GIS

Which combination of technologies is most effective for building situational awareness in environmental safety systems?

A) Computer vision models only
B) Satellite imagery without machine learning
C) Computer vision + satellite/drone data + GIS integration
D) Manual image inspection with GPS devices

References / Further Reading

[1] R. Szeliski, *Computer Vision: Algorithms and Applications*, 2nd ed. Cham, Switzerland: Springer International Publishing, 2022. doi: 10.1007/978-3-030-34372-9. [Online]. Available: https://szeliski.org/Book/, Accessed: Dec. 15, 2025.

[2] S. J. D. Prince, *Computer Vision: Models, Learning, and Inference*. Cambridge, UK: Cambridge University Press, 2012. [Online]. Available: https://raw.githubusercontent.com/udlbook/cvbook/main/book.pdf, Accessed: Dec. 15, 2025.

[3] H. J. Escalante *et al.* (eds.), *Explainable and Interpretable Models in Computer Vision and Machine Learning*. Cham, Switzerland: Springer International Publishing, 2018. doi: 10.1007/978-3-319-98131-4. [Online]. Available: https://link.springer.com/book/10.1007/978-3-319-98131-4, Accessed: Dec. 15, 2025.

[4] N. Kussul, A. Shelestov, B. Yailymov, H. Yailymova, G. Lemoine, and K. Deininger, "Assessment of war-induced agricultural land use changes in Ukraine using machine learning applied to Sentinel satellite data," *International Journal of Applied Earth Observation and*

Geoinformation, vol. 140, p. 104551, 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1569843225001980, Accessed: Dec. 15, 2025.

[5] J. Song, A. Pashazadeh, S. Wu, V. R. Greene, K. O. Hallee, A. M. Tannehill, T. M. Saxon, M. Potter, and Z. Beneke, "Advances in microplastic detection and interception: A state-of-the-art review," *Cleaner Water*, vol. 4, p. 100141, 2025. [Online]. Available: https://doi.org/10.1016/j.clwat.2025.100141, Accessed: Dec. 15, 2025.

[6] O. Ballinger, "A New Tool Allows Researchers to Track Damage in Gaza," *Bellingcat*, Nov. 15, 2023. [Online]. Available: https://www.bellingcat.com/resources/2023/11/15/a-new-tool-allows-researchers-to-track-damage-in-gaza/, Accessed: Dec. 15, 2025.

[7] L. Shumilo, M. Lavreniuk, N. Kussul, and B. Shevchuk, "Automatic Deforestation Detection based on the Deep Learning in Ukraine," in *Proceedings of the 11th IEEE International Conference on Intelligent Data Acquisition and Advanced Computing Systems (IDAACS)*, 2021. doi: 10.1109/IDAACS53288.2021.9661008. [Online]. Available: https://doi.org/10.1109/IDAACS53288.2021.9661008, Accessed: Dec. 15, 2025.

[8] D. Mozghovyi, V. Korchynskyi, and D. Svynarenko, "Methodology of satellite monitoring of deforestation based on Sentinel-2A/B data," *Challenges and Issues of Modern Science*, vol. 3, 2024. [Online]. Available: https://cims.fti.dp.ua/j/article/view/252/207, Accessed: Dec. 15, 2025.

[9] O. Rumiantsev, Y. Oliinyk, and L. Oliinyk, "Damage Detection Based on Satellite Image Analysis," in *Lecture Notes in Data Engineering, Computational Intelligence, and Decision-Making*, vol. 2, *Lecture Notes on Data Engineering and Communications Technologies*, vol. 244, S. Babichev and V. Lytvynenko, Eds. Cham, Switzerland: Springer International Publishing, 2025, pp. 177–189. doi: 10.1007/978-3-031-88483-2. [Online]. Available: https://link.springer.com/book/10.1007/978-3-031-88483-2, Accessed: Dec. 15, 2025.

[10] European Space Agency, *Copernicus Sentinel-2* [Level-2A data set], 2026. [Online]. Available: https://dataspace.copernicus.eu/, Accessed: Dec. 15, 2025.

[11] U.S. Geological Survey, *Landsat 8–9 OLI/TIRS Level-2, Collection 2* [data set]. Earth Resources Observation and Science (EROS) Center, 2026. doi: 10.5066/P9OGBGM6. [Online]. Available: https://doi.org/10.5066/P9OGBGM6, Accessed: Dec. 15, 2025.

[12] S. L. Deshmukh, M. Wilchek, and F. A. Batarseh, "HydroVision: Predicting optically active parameters in surface water using computer vision," *arXiv preprint*, 2025. [Online]. Available: https://arxiv.org/abs/2509.01882, Accessed: Dec. 15, 2025.

[13] L. K. Jamili, K. Wu, K. Natarajan, A. K. Singh, M. Tee, and M. Jindal, “Computer vision in environmental monitoring: Automated detection for biodiversity and climate action,” *International Journal of Environmental Sciences*, vol. 11, no. 19s, pp. 870–878, 2025. doi: 10.64252/a24f9536. [Online]. Available: https://doi.org/10.64252/a24f9536, Accessed: Dec. 15, 2025.

[14] A. V. Voloshko and A. O. Makarenko, “Using sensor devices combined with computer vision algorithms to improve motion-tracking systems,” *Scientific Notes of the State University of Telecommunications*, 2024. doi: 10.31673/2786-8362.2024.010202. [Online]. Available: https://doi.org/10.31673/2786-8362.2024.010202, Accessed: Dec. 15, 2025.

[15] A. Heller, M. Jacobs, G. Acosta-González, *et al.*, “Plastic water bottle detection model using computer vision in aquatic environments,” *Scientific Reports*, vol. 15, p. 24851, 2025. doi: 10.1038/s41598-025-09300-8. [Online]. Available: https://doi.org/10.1038/s41598-025-09300-8, Accessed: Dec. 15, 2025.

[16] A. Singh, A. Sengupta, and P. Venkoparao, “Deep learning-based burned area mapping using pre- and post-fire Sentinel-2 imagery,” *Remote Sensing Applications: Society and Environment*, vol. 24, p. 100614, 2021. doi: 10.1016/j.rsase.2021.100614. [Online]. Available: https://doi.org/10.1016/j.rsase.2021.100614, Accessed: Dec. 15, 2025.

[17] S. Aksoy, V. Yilmaz, and S. Öztürk, “Deep convolutional networks for post-disaster environmental damage assessment using aerial imagery,” *Scientific Reports*, vol. 12, p. 1883, 2022. doi: 10.1038/s41598-022-06067-1. [Online]. Available: https://doi.org/10.1038/s41598-022-06067-1, Accessed: Dec. 15, 2025.