

Online Continual Learning for Satellite Scene Detection Using YOLO-based Incremental Adaptation

Mrs. Kurapati Praveena
Assistant Professor, Dept. of CSE
(AIML)
SR Gudlavalleru Engineering
College Gudlavalleru, India

Ms. Sappa Meghana
Dept. of CSE (AIML)
SR Gudlavalleru Engineering
College
Gudlavalleru, India

Ms. Morla Rajitha
Dept. of CSE (AIML)
SR Gudlavalleru Engineering
College
Gudlavalleru, India

Ms. Peyyala Hyny Suma Sri
Dept. of CSE (AIML)
SR Gudlavalleru Engineering College
Gudlavalleru, India

Ms. Thadikonda Sravanthi
Dept. of CSE (AIML)
SR Gudlavalleru Engineering College
Gudlavalleru, India

Abstract - Satellite imagery plays an important role in applications such as environmental monitoring, urban planning, and disaster management. However, deep learning models trained on a fixed dataset often experience performance degradation when applied to new environmental conditions due to domain shift. This work proposes an Online Continual Learning (OCL) framework to improve adaptive satellite scene detection using a YOLO-based object detection model. Initially, the detector is trained on an old-domain dataset containing water, desert, and green landscape scenes. When evaluated on a new-domain dataset consisting of cloudy scenes, a significant drop in detection accuracy and mean Average Precision (mAP) is observed, indicating distribution mismatch. To address this challenge, incremental fine-tuning is performed using new-domain samples while retaining previously learned knowledge to reduce catastrophic forgetting. Experimental results show improvements in precision, recall, and mAP after continual learning adaptation. The system is implemented in Python and deployed through a Streamlit-based web interface to support real-time satellite scene detection. The results demonstrate that integrating online continual learning improves the robustness and adaptability of deep learning models in dynamic remote sensing environments.

I. INTRODUCTION

Satellite imagery functions as a major and highly valuable data source for environmental monitoring, disaster assessment, urban planning, and climate analysis. The continuous advancement of remote sensing technologies has led to a rapid increase in high-resolution satellite data, necessitating efficient and automated analysis methods. Deep learning models, particularly convolutional neural networks (CNNs), have demonstrated strong performance in visual recognition and object detection tasks. Among these, the YOLO (You Only Look Once) detection framework has become widely adopted due to its real-time detection capability and computational efficiency. However, models trained on a fixed data distribution often experience performance degradation when exposed to new environmental conditions, a challenge commonly referred to as domain shift. In satellite imagery, such distributional variations may arise from seasonal changes, atmospheric conditions, sensor differences, or geographic diversity. Traditional retraining strategies require rebuilding the entire model whenever new data becomes available whenever new data becomes available, which is computationally intensive and unsuitable for dynamic remote sensing systems.

To address this limitation, this work implements an Online Continual Learning (OCL) strategy to enable adaptive satellite scene detection. The YOLO-based detector was initially trained on an old domain dataset consisting of water, desert, and green landscape scenes. When evaluated on a new-domain dataset containing cloudy scenes, a noticeable decline in detection performance was observed, confirming the impact of domain shift. Incremental fine-tuning was then performed using both previously seen and newly introduced data to facilitate adaptation while preserving learned representations, thereby reducing catastrophic forgetting. Experimental analysis revealed enhanced precision, recall, and mean Average Precision (mAP) after continual learning. The trained

model was deployed via a Streamlit web application for real-time satellite scene detection. The results confirm that integrating online continual learning enhances robustness and adaptability of deep learning models in evolving remote sensing environments.

II. RELATED WORK

Satellite image analysis has experienced significant transformation with the rise of deep learning techniques, especially convolutional neural networks (CNNs), which have delivered strong performance in tasks such as object detection, semantic segmentation, and scene classification. Earlier remote sensing approaches primarily relied on handcrafted features combined with classical machine learning algorithms; however, these methods often struggled to handle variations in illumination, spatial resolution, atmospheric interference, and complex environmental patterns. The adoption of deep learning models addressed many of these limitations, enabling automatic feature extraction and improved generalization. In particular, The YOLO framework introduced a single-stage detection mechanism that performs object localization and classification simultaneously, enabling efficient real-time detection capable of real-time object localization and classification (Redmon et al., 2016). Due to its computational efficiency and accuracy, YOLO and its variants have been widely utilized in remote sensing applications for detecting buildings, vehicles, and land cover categories from high-resolution satellite imagery.

In spite of these developments, deep learning models are vulnerable to domain shift, resulting performance degradation when the test data distribution differs from the training distribution. In satellite imagery, such shifts commonly arise from seasonal variation, cloud cover, sensor discrepancies, and geographic diversity. Conventional retraining strategies require rebuilding or extensively fine-tuning the model whenever new data becomes available, resulting in substantial computational cost and limited scalability for real-time systems. Consequently, research efforts have focused on domain adaptation and transfer learning approaches to reduce distribution mismatch between source and target datasets. Techniques such as feature alignment and unsupervised adaptation have demonstrated effectiveness in mitigating performance degradation without requiring large volumes of labeled target-domain data. More recently, continual learning has emerged as a promising paradigm for handling evolving data distributions. Unlike traditional transfer learning, continual learning enables models to incrementally update their knowledge from streaming data while minimizing catastrophic forgetting — the loss of previously learned information during adaptation. Methods including replay buffers, regularization based constraints, and parameter isolation techniques have shown encouraging results in classification tasks. However, the application of continual learning to object detection remains comparatively limited, particularly within remote sensing contexts. Some studies have explored incremental detection using exemplar replay and adaptive weight-sharing mechanisms to preserve old-class performance while learning new categories.

Within remote sensing specifically, incremental learning has been investigated for land cover mapping and temporal change detection, emphasizing the need for adaptive models capable of evolving over time. Nevertheless, online continual learning for satellite scene detection — especially under environmental variations such as cloud cover — remains insufficiently explored. To address this research gap, the present work proposes a YOLO-based continual learning framework that gradually adapts a baseline satellite scene detector to new domain data while retaining knowledge acquired from the original domain. Compared to conventional offline domain adaptation methods, the proposed approach supports realtime incremental updates using fine-tuning combined with knowledge retention mechanisms, resulting in improved detection performance on new-domain imagery without significant degradation on previously learned classes. This capability is especially important for real-world satellite applications, where environmental conditions change continuously and models must adapt efficiently without full retraining.

III. SYSTEM OVERVIEW

The proposed system implements an Online Continual Learning (OCL) framework to enable adaptive satellite scene detection using a YOLO-based deep learning model. The main objective of the framework is to handle domain shift by gradually updating the model when new domain data becomes available, thereby removing the need for complete retraining. The overall architecture is structured into five major stages: dataset preparation, baseline model training, domain shift evaluation, continual adaptation, and deployment.

During dataset preparation, satellite images were organized into two domains. The old-domain dataset consisted of water, desert, and green landscape scenes, while the new-domain dataset included cloudy scenes representing previously unseen environmental conditions. All images were labeled and formatted according to the YOLO object detection standard, with matching bounding box labels for each scene category.

In the baseline training phase, the YOLO model was trained only on the old-domain dataset. Its performance was evaluated using validation data from the same domain to establish reference metrics. To examine the effect of domain shift, the trained baseline

model was later tested on new-domain cloudy images. A noticeable decline in detection performance was observed, indicating a clear discrepancy between the training and testing data distributions

To overcome this limitation, an online continual learning strategy was applied through incremental fine-tuning using a combination of old-domain and new-domain samples. This approach enabled the model to adapt to the newly introduced cloudy scenes while retaining previously learned representations, thereby mitigating catastrophic forgetting. Finally, the adapted model was deployed using a Streamlit based web interface. The system allows users to upload the satellite images, after which real-time scene detection is performed and the predicted results are displayed visually. This architecture ensures efficient adaptation to evolving satellite data distributions while maintaining practical deployment capability.

IV. SYSTEM ARCHITECTURE

The proposed system architecture is designed as a modular framework to support online continual learning for adaptive satellite scene detection. The overall architecture is illustrated in Figure. The process begins with dataset acquisition, which includes satellite images collected from two distinct domains: an old domain comprising water, desert, and green landscape scenes, and a new domain containing cloudy scenes that introduce distributional variation.

During preprocessing, the Images are organized into structured directories and converted into YOLO-compatible detection format, including bounding box annotations and corresponding label files.

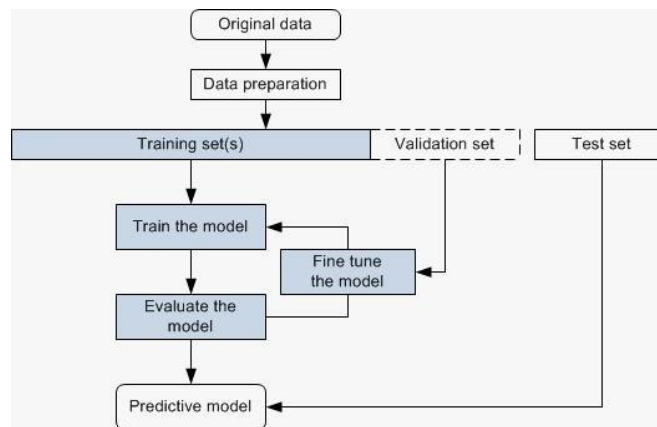


Fig. 1. Online Continual Learning Framework

cloudy images to assess the impact of domain shift. Due to environmental and distributional differences between the domains, a measurable decline in detection performance is observed, confirming the model's limited generalization capability.

To overcome this challenge, an Online Continual Learning module is incorporated into the architecture. Incremental finetuning is performed using a combination of old-domain and new-domain samples, allowing the model to adapt to newly introduced environmental conditions while retaining previously learned representations. This strategy mitigates catastrophic forgetting and ensures balanced performance across both domains.

The resulting updated YOLO model serves as an adaptive detector capable of handling diverse satellite scene variations. For practical implementation, the final model is deployed through a Streamlit-based web interface that allows users to upload satellite images and receive real-time detection results. Overall, the modular architecture enhances scalability, flexibility, and efficient real-time deployment, making the system suitable for dynamic remote sensing environments where data distributions continuously change.

V. WORKFLOW

The proposed system follows a structured workflow to enable adaptive satellite scene detection through Online Continual Learning. The workflow begins with dataset division into two domains: an olddomain dataset containing water, desert, and green landscape scenes, and a new-domain dataset consisting of cloudy scenes that introduce environmental variation. All images are organized and annotated in YOLO detection format, ensuring compatibility with the object detection framework.

In the initial phase, a baseline YOLO model is trained exclusively on the old-domain dataset to establish reference performance. When tested on the new-domain cloudy images, the model decreased detection accuracy, emphasizing the effect of domain shift caused by environmental and distributional variations.

To address this limitation, incremental fine-tuning is performed using a combination of old-domain and new-domain samples. This Online Continual Learning strategy allows the model to adapt to previously unseen conditions while preserving learned representations and mitigating catastrophic forgetting.

Finally, the adapted model is deployed using a Streamlit-based interface, enabling real-time satellite scene detection through user-uploaded images. This workflow ensures efficient adaptation to evolving data distributions while maintaining practical deployment capability.

VI. ALGORITHM

Input:

Old-domain dataset D_{old} New-domain dataset D_{new}

Output:

Adapted YOLO model $M_{adapted}$

Step 1: Dataset Preparation

1.1 Collect satellite images and divide them into:

- Old-domain dataset D_{old} (water, desert, green scenes)
- New-domain dataset D_{new} (cloudy scenes)

1.2 Convert all images into YOLO detection format.

1.3 Generate corresponding bounding box label files.

Step 2: Model Initialization

2.1 Initialize pretrained YOLO model M_0 .

2.2 Specify the training hyperparameters like learning rate, epochs, batch size.

Step 3: Baseline Training

3.1 Train M_0 using D_{old} .

3.2 Obtain trained baseline model M_{base} .

3.3 Save trained weights.

Step 4: Domain Shift Evaluation

4.1 Evaluate M_{base} on D_{new} .

4.2 measure performance metrics through:

- Precision
- Recall
- mean Average Precision (mAP)

4.3 Observe performance degradation to confirm domain shift.

Step 5: Dataset Combination

5.1 Combine datasets:

$D_{combined} = D_{old} \cup D_{new}$

Step 6: Online Continual Learning

6.1 Fine-tune M_{base} using $D_{combined}$.

6.2 Update model parameters incrementally. 6.3 Prevent catastrophic forgetting using balanced training samples.

6.4 Obtain adapted model $M_{adapted}$.

Step 7: Deployment

7.1 Deploy $M_{adapted}$ using Streamlit.

7.2 Allow users to upload satellite images.

7.3 Perform real-time scene detection and display results.

VII. IMPLEMENTATION

The proposed system is implemented in Python using modern deep learning frameworks to support satellite scene detection with Online Continual Learning. The overall implementation pipeline includes dataset preparation, baseline training, domain shift evaluation, incremental fine-tuning, performance analysis, and deployment.

The YOLO (You Only Look Once) object detection framework provided by the Ultralytics library is employed as the core detection model. A pretrained YOLO architecture is initialized and trained using the old-domain dataset, which consists of water, desert, and green satellite scenes. All images are formatted according to the YOLO detection standard with each image having an associated annotation file containing normalized bounding box coordinates and class identifiers.

Model training is conducted using the PyTorch backend in a local CPU-based environment. During optimization, the YOLO framework optimizes a composite loss function of bounding box regression loss, classification loss, and distribution focal loss to achieve precise localization and class prediction. After completion of training, the best-performing weights are saved and designated as the baseline model for further analysis.

To evaluate domain shift, the baseline model is tested on the newdomain dataset containing cloudy satellite scenes. Detection effectiveness is assessed using standard object detection metrics, like Precision, Recall, and mean Average Precision at IoU threshold 0.5 (mAP50). A noticeable reduction in these metrics confirms the impact of distributional variation between training and testing domains.

For Online Continual Learning, incremental fine-tuning is performed by combining both old-domain and new-domain datasets. Rather than retraining the network from scratch, the previously trained baseline weights serve as initialization for adaptation. This strategy enables the detector to incorporate new environmental patterns while retaining previously learned representations, thereby mitigating catastrophic forgetting.

Following incremental adaptation, performance metrics are recomputed and compared with baseline results to validate improvement. Visualization of training curves and evaluation metrics is performed using Matplotlib and Pandas to provide graphical and tabular performance analysis.

In the final stage, the adapted detector is implemented via a web interface developed using Streamlit, allowing users to upload satellite images and obtain real-time detection outputs with predicted class labels and confidence scores. The entire system is developed and executed within a local machine environment using Python, Ultralytics YOLO, PyTorch, Matplotlib, Pandas, and Streamlit, ensuring a fully software-based and scalable implementation.

VIII. EXPERIMENTAL SETUP

A. Experimental Environment

Experiments were conducted to evaluate the performance of the proposed Online Continual Learning (OCL) framework in satellite scene detection. The implementation was carried out in Python using the Ultralytics YOLO framework with the PyTorch backend. All experiments were executed in a local machine environment using a CPU configuration. The objective of the evaluation was to analyze the impact of domain shift and measure the improvement achieved through incremental learning.

B. Dataset Configuration

The satellite dataset was divided into two distinct domains:

Old Domain: Water, Desert, and Green scenes
New Domain: Cloudy scenes

All images were preprocessed and converted into YOLO detection format. Each image was provided with a corresponding annotation file that include normalized bounding box details and class identifiers. Since each image represents a single scene category, full-image bounding box annotations were used for detection.

C. Model Configuration

A pretrained YOLO model was selected as the base detector. The baseline model was trained exclusively on the old-domain dataset using the following hyperparameters:

Image Size: 416×416

Batch Size: 16

Optimizer: AdamW (default Ultralytics configuration)

Baseline Training Epochs: 10

Incremental Fine-tuning Epochs: 2–3 Device: CPU

During training, the model optimized a composite loss function consisting of bounding box regression loss, classification loss, and distribution focal loss to ensure accurate localization and classification performance.

D. Domain Shift Evaluation

After baseline training, the model was evaluated on the newdomain (cloudy) dataset to assess generalization performance. A noticeable degradation in detection metrics was observed, confirming the presence of domain shift between the training and testing distributions. This evaluation established the need for adaptive learning.

E. Online Continual Learning Setup

To address domain shift, incremental fine-tuning was performed by combining both old-domain and new-domain datasets:

$$D_{combined} = D_{old} \cup D_{new}$$

Instead of retraining the network from scratch, the baseline model weights were used as initialization. Fine-tuning was conducted for a few additional epochs, enabling the model to incorporate new-domain knowledge while preserving previously learned representations. This strategy effectively mitigated catastrophic forgetting and improved cross-domain performance.

F. Evaluation Metrics

Model performance was evaluated using standard object detection metrics:

Precision

Recall

Mean Average Precision (mAP50)

A comparative analysis of these metrics before and after incremental learning was performed to quantify performance improvement and validate the effectiveness of the proposed OCL framework.

G. Deployment Environment

The final adapted model was deployed using a web-based interface developed with Streamlit. The deployment component allows users to upload satellite images and obtain real-time detection results, with predicted class labels and confidence scores. This ensures practical usability in dynamic remote sensing applications.

IX. RESULTS AND DISCUSSION

The experiments validate the effectiveness of the proposed OCL framework.

A. Baseline Model Performance


The baseline YOLO model achieved strong performance on the old domain but failed to generalize to cloudy scenes.

B. Performance After OCL


Incremental learning improved performance significantly:

The adapted model achieved near-perfect performance with precision and recall approaching 1.0 and mAP@50 reaching approximately 0.995.


C. Performance Visualization and Deployment Results



(a) Box Loss vs Epoch



(b) mAP50 Training Curve



Metrics Table:

	Metric	Before_OCL	After_OCL
0	Precision	0.0	1.000
1	Recall	0.0	1.000
2	mAP50	0.0	0.995

(c) Performance Comparison

(d) Metrics Table Before vs After OCL

Fig. 2. Training and performance evaluation results



(a) Streamlit Home Interface



(b) Upload Interface



(c) Cloudy Scene Detection



(d) Desert Scene Detection

Fig. 3. Deployment and detection results



Fig. 4. Additional detection example

REFERENCES

- [1] German I. Parisi et al.
"Continual Lifelong Learning with Neural Networks: A Review," Neural Networks, vol. 113, pp. 54–71, 2019.
- [2] K. Shmelkov, C. Schmid, and K. Alahari,
"Incremental Learning of Object Detectors Without Catastrophic Forgetting," Proc. IEEE ICCV, 2017.
- [3] J. Kirkpatrick et al.,
"Overcoming Catastrophic Forgetting in Neural Networks," PNAS, vol. 114, no. 13, pp. 3521–3526, 2017.
- [4] Ultralytics,
"Ultralytics YOLO Documentation." Available: <https://docs.ultralytics.com/>
- [5] S.-A. Rebuffi, A. Kolesnikov, G. Sperl, and C. H. Lampert,
"iCaRL: Incremental Classifier and Representation Learning," Proc. IEEE CVPR, 2017.