2025-08

Semi-Supervised Approach: ClusterFormer for Semantic Segmentation in Remote Sensing Applications

Siddiqui, Choudhury Ben Yamin

Independent University, Bangladesh

Semi-Supervised Approach: ClusterFormer for Semantic Segmentation in Remote Sensing Applications

Choudhury Ben Yamin Siddiqui (1821161)

A Dissertation submitted in partial fulfillment for the Degree of Bachelor of Science in Computer Science and Engineering

Department of Computer Science & Engineering
Independent University, Bangladesh

August 2025

Attestation

I understand the nature of plagiarism, and I am aware of the university's strict policy on the matter. I certify that this is an original work by me. However, others' written work or software used in any part of this project is properly cited following internationally accepted academic guidelines.

Signature:

..…………………..……………………………………… Choudhury Ben Yamin Siddiqui

Acknowledgment

I feel deeply indebted to my supervisors Dr. Amin Ahsan Ali and Dr. AKM Mahbubur Rahman. Without their proper guidance and care, it would not have been possible for me to complete this research. Finally, I would also like to thank all the members of the CCDS Lab for their valuable cooperation during the work.

Abstract

Applications of remote sensing have been transformed by deep learning, particularly in semantic segmentation for mapping land use and land cover (LULC). Conventional supervised algorithms, on the other hand, spend a lot of money and time since they require a large amount of labeled data. For semantic segmentation in remote sensing applications, this paper investigates a semi-supervised learning paradigm that makes use of ClusterFormer, a transformer-based architecture intended for effective clustering and attention mechanisms. The suggested method attempts to improve segmentation accuracy while lowering reliance on ground-truth labels by merging a sizable corpus of unlabeled satellite data with a small number of annotated samples. Benchmark LULC datasets and custom high-resolution satellite images are used for experiments, while F1-score and Intersection over Union (IoU) are used as assessment measures. Results demonstrate that the semi-supervised ClusterFormer performs similarly to fully supervised baselines under limited-label conditions, highlighting its potential for scalable remote sensing analysis.

Contents

Introduction

Remote sensing has emerged as one of the most powerful technologies for monitoring and analyzing the Earth's surface. As high-resolution satellite and aerial photos become more accessible, researchers and practitioners are now able to investigate urban growth, agricultural productivity, deforestation, water resource management, and environmental changes at unprecedented scales. The wealth of spatial information provided by remote sensing platforms has made it possible to derive critical insights for decision-making across domains such as urban planning, disaster management, climate change monitoring, and land use and land cover analysis. However, to make effective use of this data, robust and accurate image analysis techniques—like semantic segmentation—are essential.

One of the core tasks in computer vision that is essential to distant sensing applications is semantic segmentation. Semantic segmentation, as opposed to basic classification tasks, gives each pixel in a picture a class label, allowing for a more thorough comprehension of spatial patterns and landscape composition. Segmentation, for instance, can differentiate between vegetation, water bodies, built-up regions, and other classifications, allowing for a better understanding of spatial structure and the identification of changes. Semantic segmentation is one of the most significant problems in geospatial analysis because of this pixel-level classification.

Despite its importance, semantic segmentation in remote sensing faces several challenges. Satellite and aerial images typically present large and high-resolution scenes that cover diverse landscapes, variable lighting, seasonal effects, and atmospheric conditions. Crucially, producing ground truth annotations at the pixel level is extremely resource-intensive—demanding domain expertise, substantial manual effort, and time. As a result, although massive volumes of remote sensing imagery are available, only a small portion is labeled, limiting the practical utility of supervised deep learning models.

Semi-supervised learning provides a powerful solution to this issue by enabling models to learn from a minuscule amount of labeled data together with massive volumes of unlabeled data. In this paradigm, labeled samples provide guidance, while unlabeled data contributes structural learning about the underlying distribution. For remote sensing, where unlabeled imagery is abundant, semi-supervised learning is particularly beneficial. It allows segmentation models to generalize better, adapt to new geographic regions, and reduce dependency on costly annotations [Peláez-Vegas et al. 2023].

Convolutional neural network (CNN)-based deep learning architectures, such U-Net and DeepLab, have been extensively employed for semantic segmentation in remote sensing in recent years and have shown good performance. However, because convolutional filters are local, CNNs may have trouble capturing global context and long-range interdependence. By considering image patches as tokens and employing global self-attention, Vision Transformers (ViTs) have overcome this constraint and improved the representation of both local and global structure, which is crucial for remote sensing imaging where objects vary greatly in scale and context.

ClusterFormer is a recent development in transformer-based architecture that enhances both efficiency and representation power through a novel clustering mechanism. Rather than attending to every token individually, ClusterFormer groups tokens into clusters and performs attention at the cluster level. This significantly reduces computational complexity while encouraging structured, semantically meaningful representations. The architecture comprises RCAC (recurrent cross-attention clustering for updating cluster centers recursively) and feature dispatcher (for redistributing features based on similarity), forming a transparent and transferable pipeline across visual tasks such as detection, segmentation, & classification [Liang, James, et al. 2023].

The clustering-based attention mechanism of ClusterFormer is particularly relevant for remote sensing tasks. Satellite imagery often involves extremely high-resolution data and many patches. ClusterFormer's design enables it to process such data more efficiently and effectively by capturing fine-grained local details while preserving global context—a compelling advantage for semantic segmentation in remote sensing.

When combined with semi-supervised learning strategies, ClusterFormer becomes even more compelling. Techniques like pseudo-labeling, entropy minimization, and consistency regularization can be used to extract significant patterns from unlabeled images to enhance structural learning and labeled data to aid segmentation. This leverages the abundance of unlabeled remote sensing data, improving model generalization while reducing the need for expensive annotation.

The integration of semi-supervised learning with ClusterFormer offers significant potential across diverse remote sensing applications, such as urban expansion monitoring, land cover change detection, environmental assessment, and agricultural analysis. Its balance of computational efficiency, scalability, and accuracy makes it suitable for both academic research and operational systems. By reducing reliance on annotated data and taking advantage of vast unlabeled datasets, this combined approach sets the stage for scalable and continuously adaptable segmentation models.

In summary, semantic segmentation in remote sensing remains a vital yet challenging task due to the limited availability of labeled data and the abundance of unlabeled imagery. Semi-supervised learning provides an effective bridge, while transformer-based architectures like ClusterFormer enhance representation power and efficiency. Bringing these together opens a promising direction for advancing remote sensing applications—enabling models that are more accurate, scalable, and adaptable to real-world demands.

Background and Context

Remote sensing has rapidly evolved over the past decades as a cornerstone of environmental monitoring, geoscience research, and urban management. The availability of Earth observation satellites such as Landsat, Sentinel, and commercial high-resolution sensors has produced an unprecedented wealth of imagery covering diverse geographical regions and time periods. These datasets enable a range of critical applications including land use mapping, agricultural monitoring, urban growth assessment, and disaster management [Weng, Q. et al. 2018], [Zhu, X. X. et al. 2017] However, unlocking the full potential of these massive data streams requires automated methods that can accurately interpret complex visual patterns.

Semantic segmentation has therefore become a pivotal task in remote sensing analysis. Early approaches relied heavily on traditional image processing and machine learning techniques, such as pixel-based classification, clustering, and support vector machines. While useful, these methods had restrictions in their ability to capture spatial context and generalize across large-scale heterogeneous imagery ([Bruzzone et al. 2009]). The advent of deep learning, particularly Convolutional Neural Networks (CNNs), marked a turning point. Models such as Fully Convolutional Networks (FCNs), U-Net, and DeepLab significantly improved segmentation accuracy by leveraging hierarchical feature extraction and multi-scale representations ([Long, J et al. 2015], [Ronneberger et al. 2015]). These CNN-based models quickly became the backbone of many remote sensing applications, achieving top of the line results on benchmark datasets.

Nevertheless, CNN-based methods face inherent challenges. Convolutions are inherently local operations, which means capturing long-range dependencies requires deep architectures and complex multi-scale modules. For remote sensing imagery, where objects differentiate greatly in scale, and spatial context across large regions is crucial—this limitation can hinder performance. Additionally, CNNs often struggle to adapt across different geographical domains due to spectral and textural variability in satellite imagery.

The introduction of Vision Transformers (ViTs) has redefined how image understanding tasks are approached. By treating image patches as tokens and employing global self-attention, ViTs capture both local and long ranged dependencies more effectively than CNNs [Dosovitskiy, A., et al. 2020]. Their achievement in classification and segmentation tasks has sparked widespread interest in the remote sensing community, where global context is often essential. Variants such as Swin Transformer brought in hierarchical designs and shifted windows to balance efficiency & accuracy, further extending their applicability [Liu, Z., et al. 2021].

Building on this trend, ClusterFormer represents a more recent advancement in transformer-based architecture. Unlike standard self-attention, which scales quadratically with the number of tokens, ClusterFormer introduces a clustering-based mechanism that groups similar tokens and computes attention at the cluster level. This not only reduces computational complexity but also promotes structured and semantically meaningful feature representations [Liang, Y., et al. 2023]. For remote sensing tasks involving extremely high-resolution imagery, such efficiency is critical, as it enables models to handle large inputs while preserving fine-grained details and global context simultaneously.

While architectural innovations such as CNNs and transformers have significantly improved segmentation performance, the data bottleneck remains a pressing issue. Annotating large-scale satellite imagery is time-consuming and costly, particularly when pixel-level precision is required. This has created an imbalance: while unlabeled imagery is abundant, labeled datasets remain relatively small. Semi-supervised learning (SSL) has made appearance as a promising solution to this challenge. SSL influences both labeled and unlabeled data during training, often employing strategies such as pseudo-labeling, consistency regularization, and entropy minimization to guide learning on unlabeled samples [Peláez-Vegas et al. 2023], [Li, J., et al. 2023].

In the context of remote sensing, SSL is particularly advantageous. The continuous acquisition of new satellite imagery means that vast amounts of unlabeled data are available for nearly every region of the globe. SSL frameworks allow these unlabeled datasets to contribute meaningfully to training, improving model generalization and robustness. For example, models trained with SSL can better adapt to seasonal variations, sensor differences, and geographical diversity—key factors in practical remote sensing applications [Wang, Y., et al. 2025].

The integration of semi-supervised learning with advanced architectures such as ClusterFormer reflects a broader shift in computer vision: moving away from reliance on massive, labeled datasets toward methods that can scale efficiently with minimal supervision. This aligns with global initiatives in Earth observation, where rapid, scalable, and cost-effective analysis pipelines are necessary to keep up with the pace of data generation. In this context, semi-supervised ClusterFormer models have the potential to accelerate research and operational workflows across environmental monetization, urban planning, agriculture, and disaster response.

Overall, the background of semantic segmentation in remote sensing highlights a trajectory from traditional methods to deep learning and transformer-based architectures, while the context emphasizes the growing role of semi-supervised learning as a response to data scarcity. Together, they establish a foundation for exploring semi-supervised ClusterFormer models as a compelling direction in the pursuit of scalable and accurate remote sensing analysis.

Objectives

This project's main goal is to research the possibilities of a ClusterFormer-based SSL framework for semantic segmentation in remote sensing applications. The study intends to show how semi-supervised techniques can improve model performance, generalization, and efficiency in large-scale Earth observation tasks, given the difficulties of sparse annotated data and the availability of unlabeled imagery.

The specific objectives of the project can be summarized as follows:

1. **Create a semi-supervised structure for semantic segmentation**
    - Design and implement a training pipeline that combines a minuscule set of labeled images with a larger pool of unlabeled data.
    - Integrate SSL strategies, such as pseudo-labeling & consistency regularization, to effectively leverage unlabeled samples.
2. **Adapt and evaluate the ClusterFormer architecture for remote sensing data**
    - Utilize the clustering-based attention mechanism of ClusterFormer to capture both local fine-grained structures and global contextual relationships within satellite imagery.
    - Assess the efficiency of ClusterFormer in handling high-resolution inputs compared to conventional transformer-based architectures.
3. **Investigate the benefits of SSL in remote sensing semantic segmentation**
    - Quantify the improvement in segmentation accuracy when unlabeled data is incorporated into training.
    - Analyze how semi-supervised learning improves robustness across different spatial scales, geographic regions, and image conditions.
4. **Compare performance**

- Establish benchmark results using conventional fully supervised segmentation models.
- Demonstrate the advantages of SSL in scenarios where annotated data is in short supply but unlabeled data is not.

5. **Highlight the applicability of the approach to real-world scenarios**
   - Show the relevance of semi-supervised ClusterFormer models for operational applications such as urban mapping, agricultural monitoring, and environmental assessment.
   - Discuss the potential of the approach to scale with continuously incoming satellite imagery and adapt to new geographic regions without extensive manual annotation.

By achieving these goals, the project hopes to advance the creation of accurate, scalable, and efficient techniques for semantic segmentation in remote sensing. The work aims to lessen reliance on expensive annotated datasets and unleash the potential of the enormous reservoirs of unlabeled remote sensing imagery by fusing the representational capacity of ClusterFormer with the advantages of semi-supervised learning.

.

Scope and Limitations

Scope

This project focuses on developing and evaluating a semi-supervised framework for semantic segmentation in remote sensing applications using the ClusterFormer architecture. The scope of the project is defined as follows:

1. **Application Domain**
    - The study is situated within the field of remote sensing, emphasizing land use and land cover analysis through semantic segmentation.
    - The framework is designed to be adaptable to a variety of remote sensing datasets, without being restricted to a single geographic region or sensor type.
2. **Learning Paradigm**
    - The project adopts a **semi-supervised learning** approach, making use of a small amount of labeled data combined with a huge amount of unlabeled imagery.
    - The focus is on leveraging techniques such as pseudo-labeling and consistency regularization to exploit information contained in unlabeled data.
3. **Model Architecture**
    - The project employs **ClusterFormer**, a transformer-based architecture that incorporates clustering mechanisms to improve efficiency and representation learning.
    - The study investigates how the clustering-based attention mechanism can balance local detail and global context in higher resolution satellite images.

4. **Evaluation**
    - The architecture is evaluated using standard segmentation metrics such as IoU(Intersection over Union () and F1 score.
    - Comparisons are made between the semi-supervised framework and conventional fully supervised baselines to highlight the added value of using unlabeled data.
5. **Practical Relevance**
    - The project highlights potential applications in environmental monitoring, agricultural management, and urban planning.
    - It aims to provide insights into how semi-supervised frameworks can reduce the need for large-scale manual annotation, making them suitable for real-world remote sensing pipelines.

Limitations

While the project demonstrates the potential of semi-supervised ClusterFormer models for remote sensing segmentation, several limitations must be acknowledged:

Academic and Technical Limitations

1. **Dependence on Data Quality**
    - The performance of the framework is influenced by the quality of both labeled and unlabeled datasets. Issues such as noise in ground truth labels, cloud cover, and sensor inconsistencies may affect results.
2. **Computational Requirements**
    - Although ClusterFormer reduces attention complexity through clustering, training transformer-based models on high-resolution imagery remains

computationally demanding. Scaling to very large datasets would require more powerful computing resources.

3. **Generalization Constraints**
   - The model's capabiloty to generalize across many different geographical regions, sensor modalities (e.g., multispectral or hyperspectral data), or temporal conditions is not fully explored within the scope of this project.
   - Adaptation to multi-sensor fusion scenarios (e.g., combining optical and radar data) lies beyond the present study.
4. **Semi-Supervised Assumptions**
   - The effectiveness of SSL depends on the assumption that unlabeled data shares the same distribution as the labeled samples. In cases where this assumption is violated (e.g., significant domain shifts), performance may degrade.
5. **Deployment Constraints**
   - The project evaluates performance in controlled experimental conditions. Real-world deployment would face additional requirements such as real-time inference, integration with GIS systems, and efficient processing of continuous satellite data streams.

Practical Constraints Encountered During the Project

1. **Hardware Limitations**
   - Training was primarily conducted on cloud-based platforms (e.g., Colab, mid-range GPUs such as RTX 4060 and high-End GPUs such as L40S). Session time limits and restricted GPU memory (4060) occasionally constrained batch sizes and model scaling.

2. **Runtime Interruptions**

    - Long training sessions were sometimes interrupted due to Colab timeouts, server restarts, or connectivity issues. This required careful checkpointing and, at times, re-running experiments from earlier epochs.

3. **Software and Environment Issues**

    - Setting up environments for ClusterFormer and related dependencies (e.g., PyTorch, mmsegmentation, MMCV) posed challenges, including version mismatches, CUDA incompatibilities, and package installation failures. These issues occasionally delayed experimentation.

4. **Dataset Handling**

    - Preprocessing and handling high-resolution satellite imagery requires significant disk space and memory. In some cases, image patching and resizing strategies had to be applied to ensure data could fit into GPU memory.

5. **Experiment Scope**

    - Due to time and resource constraints, the project was limited to a fixed set of semi-supervised strategies.

Summary

The scope of this project is centered on developing & valuating a semi-supervised ClusterFormer framework for semantic segmentation in remote sensing, with the target of demonstrating how unlabeled data can improve segmentation performance and reduce reliance on costly annotations. At the same time, both academic limitations (such as generalization challenges and semi-supervised assumptions) and practical constraints (such as hardware limits, runtime interruptions, and environment issues) shaped the implementation. These considerations establish realistic boundaries for interpreting results and highlight directions for future improvement.

Contributions

This project makes the following key contributions in the field of semantic segmentation for remote sensing applications:

1. **Implementation of a Semi-Supervised ClusterFormer Framework for LULC Segmentation**
    - Designed and implemented a SSL pipeline based on the **ClusterFormer** architecture for **LULC segmentation**.
    - Integrated strategies such as pseudo-labeling and consistency regularization to utilize both labeled and unlabeled data.
    - Adapted the framework to handle high-resolution remote sensing imagery, addressing challenges of scale and variability inherent in satellite datasets.
2. **Benchmarking Against Baselines**
    - Conducted a comparative study between the semi-supervised ClusterFormer framework and other approaches.
    - Established baseline performance using existing models.
    - Demonstrated how semi-supervised training enhances performance when labeled data is scarce, while maintaining competitive accuracy with fully supervised settings.
3. **Analysis of Performance Under Limited-Label Conditions**
    - Analyzed trends in metrics like IoU and F1-score to quantify the benefits of incorporating unlabeled imagery.
    - Provided insights into how semi-supervised approaches improve generalization, robustness, and adaptability in LULC mapping tasks.

Literature Review

The growing demand for efficient and accurate semantic segmentation in remote sensing has driven research in two key directions: the development of novel deep learning architectures capable of handling large-scale, high-resolution data, and the design of learning paradigms that can exploit the abundance of unlabeled imagery. Two works stand out in this context: the introduction of **ClusterFormer** by Liang et al. (2023), which proposed a clustering-based transformer for universal visual learning, and the study by **Li et al. (2023)**, which focused on semi-supervised semantic segmentation specifically tailored to remote sensing imagery.

ClusterFormer: Clustering as a Universal Visual Learner

Traditional transformer-based architectures have demonstrated remarkable performance in computer vision tasks, yet their computational cost remains a significant barrier, especially when applied to high-resolution imagery common in remote sensing. Liang et al. (2023) introduced **ClusterFormer**, a novel transformer architecture designed to reduce the quadratic complexity of self-attention while maintaining rich representational power. Instead of attending to every token individually, ClusterFormer employs a **clustering mechanism** that groups tokens into semantically meaningful clusters. Cross-attention is then performed between input tokens and cluster centers, followed by a feature redistribution step, allowing for both **global context capture** and **local detail preservation**.

This architecture demonstrates several advantages. First, the clustering mechanism provides **computational efficiency**, scaling more effectively to large inputs compared to standard transformers. Second, it enhances **interpretability**, as cluster centers can be viewed as semantic prototypes summarizing important regions of the image. Third, the model shows strong **transferability** across tasks, excelling in classification, detection, and segmentation benchmarks. The work of Liang et al. (2023) highlights the potential of clustering-based attention as a general visual learning strategy, providing a strong foundation for adapting transformers to remote sensing applications, where both efficiency and scalability are critical.

Semi-Supervised Semantic Segmentation in Remote Sensing

While architectural advances like ClusterFormer address efficiency and representation learning, the problem of **limited labeled data** remains pressing in remote sensing. Li et al. (2023) investigated **SSL for remote sensing semantic segmentation**, a paradigm where models are trained on both a small set of labeled samples and a large set of unlabeled data. Their study emphasized that while vast amounts of remote sensing imagery are readily available from satellites, obtaining pixel level annotations is expensive and time-consuming. SSL thus provides a practical approach to bridge this gap.

Li et al. (2023) proposed a framework that incorporates **consistency regularization** and **pseudo-labeling** strategies. Consistency regularization makes the model produce stable predictions under input disturbance, ensuring robust learning from unlabeled data. Pseudo-labeling leverages model predictions on unlabeled samples as supervisory signals, allowing unlabeled data to progressively guide training. Their results demonstrated that semi-supervised approaches significantly outperform purely supervised baselines when labeled data is hard to come by, confirming the viability of SSL in operational remote sensing contexts.

The study also highlighted key challenges, such as the reliability of pseudo-labels under distribution shifts and the sensitivity of SSL methods to the quality of initial labeled data. Nevertheless, the work provides strong evidence that semi-supervised methods can alleviate the dependency on large-scale manual annotations, which is often a bottleneck in deploying deep learning systems for Earth observation.

Connecting the Two Perspectives

Together, these two strands of research form the foundation for the present project. ClusterFormer, with its clustering-based attention mechanism, provides an efficient and scalable architecture capable of handling the complexity of high-resolution remote sensing imagery. On the other hand, Li et al. (2023) demonstrates the effectiveness of semi-supervised paradigms in leveraging unlabeled remote sensing data, thereby addressing the scarcity of labeled samples. By integrating these ideas, this project positions itself at the intersection of **advanced transformer architectures** and **semi-supervised learning strategies**, aiming to build a framework that balances efficiency, accuracy, and scalability for land use and land cover segmentation.

Dataset and Preprocessing

Dataset Description

The dataset consists of satellite images of Dhaka metropolitan city and its surrounding areas with a gsd 2.22m/pixels The images initially had 11 classes (Unrecognized, Water, Farmland, Forest, Urban Structure, Brick Factory, Road, Urban Built-up, Rural Built-up, Marshland, Meadow).

Due to class imbalance, it was converted to 6 classes, and the original large image was patched with size of **513 × 513 pixels** with a 75% overlap for the train set and 0% overlap for the test set.

The test set is of the Dhaka City satellite images cropped from the large image, hence although it is a subset of the original image it is not a subset of the train images.

Thus, the dataset used in this project is a custom **Land Use and Land Cover (LULC) segmentation dataset** derived from high-resolution satellite image. It was specifically designed for semantic segmentation tasks and consists of paired images and masks. Each image represents a satellite patch of **513 × 513 pixels**, while the corresponding ground truth (GT) mask provides pixel-level annotations.

The dataset includes **six classes**, representing typical land cover types observed in remote sensing imagery:

1. **Background**: Unrecognized
2. **Farmland**: Farmland
3. **Water bodies**: Water
4. **Forest**: Forest
5. **Built-up Structures**: Urban built-up, Rural built-up, Road, Urban Structure, Brick Factory

6. **Meadow**: Marshland, Meadow

Figure 1 Sample Dhaka City Satellite Images

There was also a Dataset with Unlabeled Images which were only used for training the model.

Preprocessing Steps

1. **Image Patching**
    - Original satellite scenes were too large for direct input into deep learning models.
    - Images were therefore divided into fixed-size **513 × 513 patches**, ensuring compatibility with the ClusterFormer input requirements and fitting into available GPU memory.
2. **Ground Truth Formatting**
    - The ground truth masks were converted into **integer-labeled segmentation masks**, where each pixel corresponds to a class index (0–5).
    - This ensured compatibility with standard loss functions such as cross-entropy and allowed efficient metric computation (IoU, F1-score).
3. **Normalization**
    - Input pictures were normalized through scaling pixel intensities to a standard range.
    - This improved numerical stability during training and ensured consistent input distributions across labeled and unlabeled data.
4. **Data Augmentation**
    - To improve validity and prevent overfitting, several data augmentation techniques were applied during training, including:

        - Random horizontal & vertical flips
        - Random rotations
        - Random crops
    - These augmentations helped the model generalize across different spatial configurations and variations in satellite imagery.

5. **Semi-Supervised Data Handling**
    - The **labeled subset** provided pixel-level supervision during training.
    - The **unlabeled subset** was incorporated into the SSL framework through techniques such as pseudo-labeling and consistency regularization.
    - Preprocessing ensured that both labeled and unlabeled samples shared the same resolution and normalization, maintaining consistency across the training pipeline.

6. **Train-Test Split**
    - The labeled dataset was divided into two. One was training & the other was testing sets.
    - The unlabeled samples were only used during training and not during evaluation.

Methodology

The methodology of this project combines a semi-supervised learning framework with ClusterFormer architecture to address semantic segmentation in remote sensing imagery. The approach leverages both labeled and unlabeled satellite data, reducing dependency on costly annotations while exploiting the representational power of transformer-based clustering mechanisms. The proposed system follows a pipeline that includes **dataset preparation, feature extraction with ClusterFormer, semi-supervised training strategies, and evaluation against benchmark baselines**.

Clustering: Clustering is the process of dividing a collection of data points, represented by $X \in R^{n*d}$, into C different clusters according to their inherent similarities, making sure that each data point is a member of only one cluster. In order to accomplish this, the stratification of the data points must be optimized, taking into account both their attributes and location information, in order to create meaningful and cogent groupings. Advanced similarity metrics, such cosine similarity, are commonly used in clustering approaches to quantify the distance between data points and cluster centroids. To create more accurate group assignments, they also consider the points' geographical locality.

Cross-Attention for Generic Clustering: Inspired by the Transformer decoder design [Ashish V. et al. 2017], the modern end-to-end architecture [Bowen C. et al 2022] employs a query-based methodology where a sequence of cross attention blocks learns and updates a set of K questions, $C = [c_1;\cdots;c_K] \in R^{K* D}$. We use the term "C" in this context to link queries to each layer's cluster centers. In particular, each layer uses cross-attention to adaptively aggregate picture features before updating the queries:

$$\boldsymbol{C} \leftarrow \boldsymbol{C} + \mathrm{softmax}_{HW}(\boldsymbol{Q}^C(\boldsymbol{K}^I)^\top)\boldsymbol{V}^I$$

where the linearly projected features for query, key, and value are denoted by $Q^C \in R^{K* D}$, $V^I \in R^{HW*D}$, $K^I \in R^{HW*D}$, respectively. The features projected from the center and image features are indicated by the superscripts "C" and "I," respectively. Inspired by [Qihang Y. et al. 2022], we use the SoftMax function along the query dimension (K) rather than the picture resolution (HW) and view queries as cluster centers in order to reinterpret the cross-attention process as a clustering solver:

$$\boldsymbol{C} \leftarrow \boldsymbol{C} + \mathrm{softmax}_K(\boldsymbol{Q}^C(\boldsymbol{K}^I)^\top)\boldsymbol{V}^I$$

ClusterFormer Architecture

ClusterFormer reformulates the self-attention mechanism in transformers from a clustering perspective. Instead of computing attention across all tokens, ClusterFormer introduces **cluster centers** that act as prototypes of semantically meaningful groups of features. The architecture consists of two key components:

1. **Recurrent Cross-Attention Clustering (RCAC):**
    - Image patches are first embedded into tokens.
    - A set of cluster centers is initialized from pooled image features.
    - Cross-attention is performed iteratively (Expectation-Maximization style) between tokens and centers:
        - **E-step:** Tokens are assigned to clusters via similarity scores.

$$\hat{\boldsymbol{M}}^{(t)} = \mathrm{softmax}_K(\boldsymbol{Q}^{C^{(t)}}(\boldsymbol{K}^I)^\top),$$

        - **M-step:** Cluster centers are updated using aggregated token features.

$$\boldsymbol{C}^{(t+1)} = \hat{\boldsymbol{M}}^{(t)}\boldsymbol{V}^I \in \mathbb{R}^{K\times D}$$

    - This process is repeated over multiple iterations to refine assignments, making the clustering adaptive and transparent.

- Recurrent Cross-Attention may fully investigate their presentational granularity since the general architecture is hierarchical, which is similar to how hierarchical clustering works:

$$\boldsymbol{C}^l = \mathrm{RCA}^l(\boldsymbol{I}^l, \boldsymbol{C}_0^l),$$

  where Recurrent Cross-Attention Layer is referred to as RCA. Il is the feature map of the image at various layers using the typical pooling procedure with a resolution of $H/2^l * W/2^l$. C^l_0 is the beginning centers at the l^{th} layer, and C^l is the Cluster Center Matrix for the l^{th} layer. The RCA parameters at various locations are not shared. [Li, J., et al. 2023]

- In addition, the centers from image grids are initialized:

$$[\boldsymbol{c}_1^{(0)}; \cdots; \boldsymbol{c}_K^{(0)}] = \mathrm{FFN}(\mathrm{Adptive\_Pooling}_K(\boldsymbol{I}))$$

  where FFN stands for Position-wise Feedforward Network which is an integral part of the Transformer architecture. It comprises two fully connected layers along with an activation function used in the hidden layer. $\mathrm{Adptive\_Pooling}_K(I)$ refers to select "K" feature centers from "I" using adaptive sampling, which calculates an appropriate window size to achieve a desired output size adaptively, offering more flexibility and precision compared to traditional pooling methods. [Li, J., et al. 2023]

2. **Feature Dispatching:**

- After cluster updates, token features are redistributed according to their similarity to updated cluster centers. For every patch embedding pi ∈ I, the updated patch embedding p′ i is computed as,

$$p_i^{'} = p_i + \mathrm{MLP}(\frac{1}{K}\sum_{k=0}^{K} sim(C_k, p_i) * C_k)$$

- This mechanism ensures that tokens capture both **local details** and **global context**, which is essential for remote sensing where features vary in scale and spatial distribution. [Li, J., et al. 2023]

3. **Hierarchical Encoder-Decoder Design:**

   - The encoder generates multi-scale feature maps via RCAC layers.
   - The decoder uses the final cluster centers as queries, passing them through transformer decoder layers to produce segmentation maps.
   - For semantic segmentation, the decoder outputs pixel-level class probabilities.

By clustering tokens, this architecture dramatically lowers the quadratic complexity of vanilla attention ($O(N^2)$) to a more manageable level, making it especially well-suited for high-resolution satellite imagery.

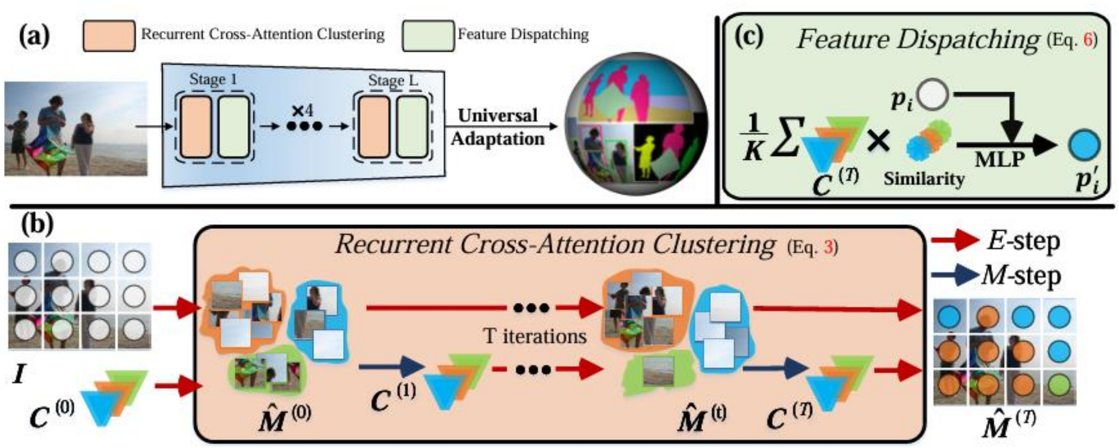

Figure 2 ClusterFormer Architecture [Liang et al. (2023)]

To address the availability of labeled data in remote sensing, the methodology integrates **semi-supervised learning (SSL)** strategies:

1. **Supervised Loss (L_sup):**

- Applied to labeled data using **cross-entropy loss** between predicted segmentation maps and ground truth masks.

2. **Unsupervised Loss (L_unsup):**
   - For unlabeled data, pseudo-labels are generated from the model's own predictions.
   - **Consistency regularization** forces the model to produce stable predictions under varying perturbations (e.g., random flips, crops).
   - A confidence threshold is applied to pseudo-labels to reduce noise.

3. **Total Loss Function:**
   - The final objective combines both supervised and unsupervised components:

$$L = L_{sup} + \lambda . L_{unsup}$$

where λ balances the contribution of unlabeled data during training.

Training Setup and Hyperparameters

Training was conducted on satellite image patches of size **513 × 513** with six land cover classes (background, farmland, water, forest, built-up, meadow). The experiments followed the default ClusterFormer settings (Liang et al., 2023) but were adapted to available compute resources and dataset characteristics.

- **Backbone:** ClusterFormer
- **Optimizer:** AdamW
- **Learning Rate:** 1×10^{-5} with cosine annealing schedule
- **Batch Size:** 16 (varies for GPU memory limitations)
- **Epochs:** 40

- **Loss Function:** Weighted sum of cross-entropy (supervised) and with unsupervised consistency loss applied to unlabeled data
- **Augmentations:** Random flips, rotations, random crops, and normalization
- **Cluster Settings:**
    - Cluster centers amount = semantic classes amount (K = 6).
    - Recurrent Cross-Attention iterations: T = 3 per layer (as suggested in original paper).
    - Multi-head attention dimension: 32.

The model was trained in a semi-supervised setting where a fraction of the training set (20%) was treated as labeled, while the other images were used as unlabeled data. Evaluation was carried out using the held-out test set, with metrics including **mIoU**(**mean Intersection over Union)** and **F1-score** reported per class.

Summary

In summary, the proposed methodology integrates the **ClusterFormer architecture** with **semi-supervised training strategies** to perform semantic segmentation of remote sensing imagery. ClusterFormer provides an efficient and transparent attention mechanism through clustering, while semi-supervised learning enables effective use of abundant unlabeled data. Together, they form a scalable and robust framework for land use and land cover mapping.

Experiments

Training was conducted on NVIDIA L40s GPUs with PyTorch. We evaluated supervised baselines (U-Net, Segmenter) against ClusterFormer (semi-supervised). Metrics: IoU, F1. We had access to two L40s GPUs for U-Net and one L40s with Segmenter and ClusterFormer.

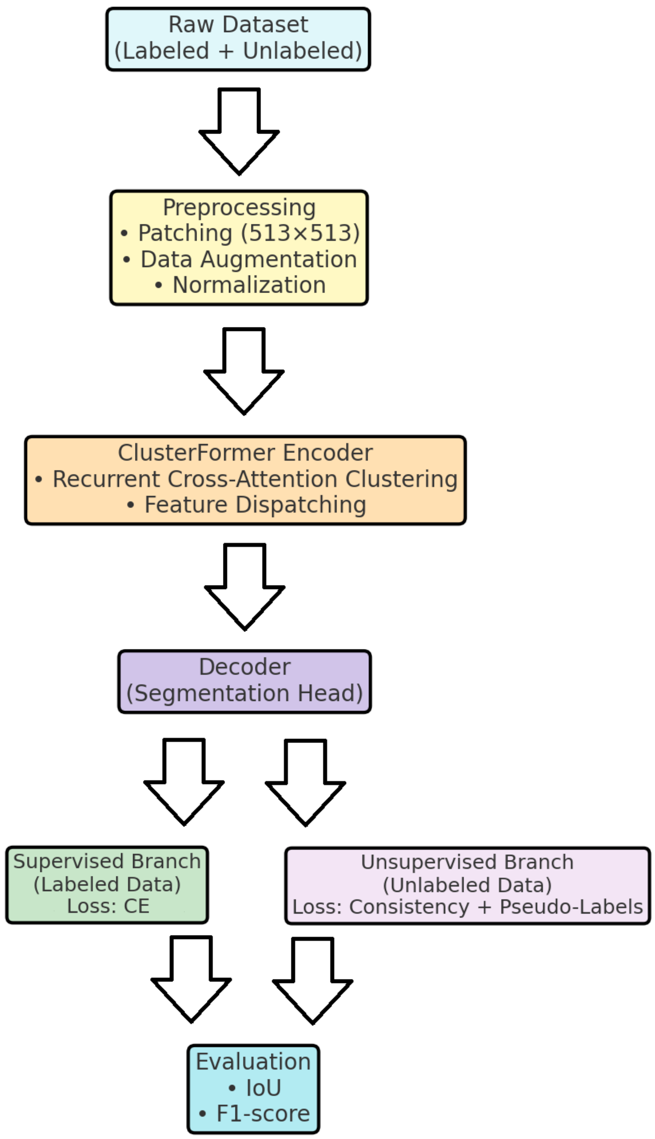

Figure 3 Experiment Pipeline for ClusterFromer

Results and Discussion

The semi-supervised ClusterFormer achieved comparable mean IoU compared to supervised baseline. It demonstrated robustness to class imbalance. However, performance varied across land cover types.

Class Index	Class Name	IoU	F1
0	BackGround	NA	N/A
1	Farmland	0.5822	0.7359
2	Water	0.4092	0.5807
3	Forest	0.3605	0.5299
4	Built-up	0.6598	0.7950
5	Meadow	0.2038	0.3386
	Mean	**0.4431**	**0.5960**

Figure 4 Segmenter_Evaluation

Class Index	Class Name	Iou Score	F1 Score
0	BackGround	N/A	NA
1	Farmland	0.5945	0.7457
2	Water	0.5667	0.7234

3	Forest	0.3446	0.5126
4	Built-up	0.6739	0.8052
5	Meadow	0.3065	0.4692
	Mean	**0.4972**	**0.6512**

Figure 5 Unet Evaluation

Class Index	**Class Name**	**Iou Score**	**F1 Score**
0	**BackGround**	N/A	N/A
1	Farmland	0.5557	0.7144
2	Water	0.4872	0.6552
3	Forest	0.3056	0.4682
4	Built-up	0.6581	0.7938
5	Meadow	0.1934	0.3242
	Mean	0.44	0.59

Figure 6 ImageNet Pretrained ClusterFormer

Class Index	**Class Name**	**Iou Score**	**F1 Score**
0	**BackGround**	N/A	N/A

1	Farmland	0.53914	0.692075
2	Water	0.461452	0.61251
3	Forest	0.330048	0.461596
4	Built-up	0.692762	0.789904
5	Meadow	0.276657	0.369553
	Mean	0.460012	0.585128

Figure 7 ClusterFormer Trained on LULC Labeled Train Set

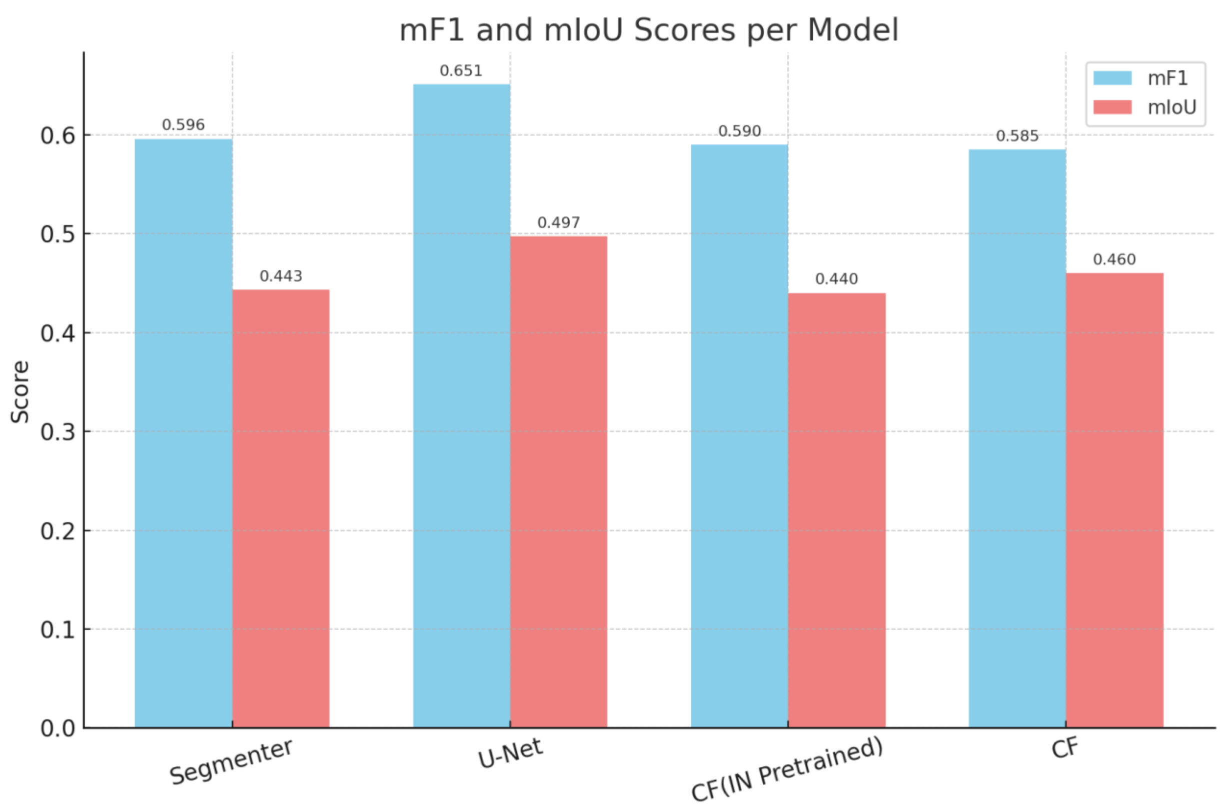

Visual Result Comparison

<table>
<tr><th>Model</th><th>Original Image</th><th>Ground Truths</th><th>Predicted</th></tr>
<tr><td>U-net</td><td rowspan="4"></td><td rowspan="4">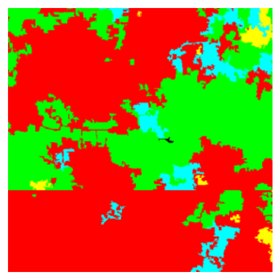</td><td>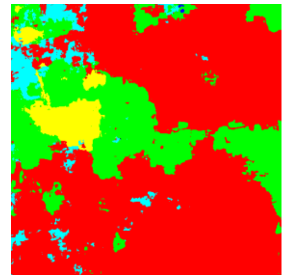</td></tr>
<tr><td>Segmenter</td><td>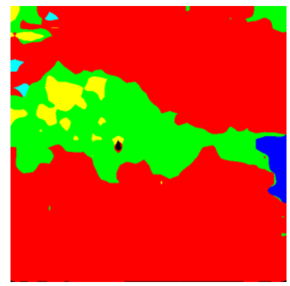</td></tr>
<tr><td>CF(IN Pretrained)</td><td>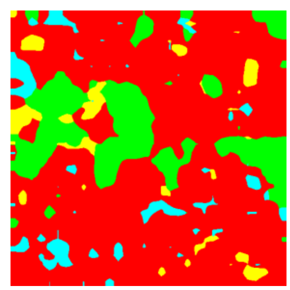</td></tr>
<tr><td>CF</td><td>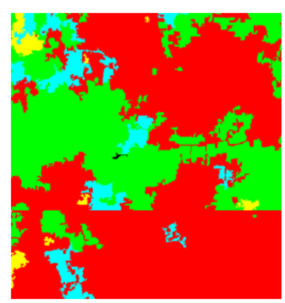</td></tr>
</table>

Table 1 Successful Output Comparison

Failure Case Analysis

Though CF and CF(IN-pretrained) performed well compared to other models (U-net and Segmenter), it still has failures similar to other models. Usual failures are,

- Incorrect Detections:
  Usually, this is caused by some class features extracted being similar. For example, below shows an incorrect detection,

<table>
<tr><td>Model</td><td>Correct</td><td>Incorrect</td></tr>
<tr><td>CF</td><td rowspan="2">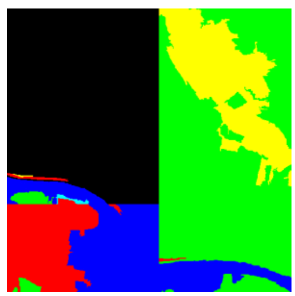</td><td>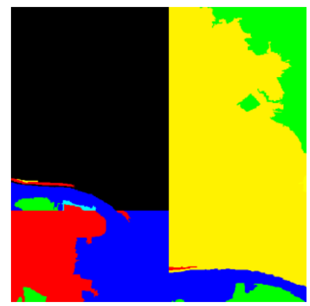</td></tr>
<tr><td>CF(IN Pretrained)</td><td>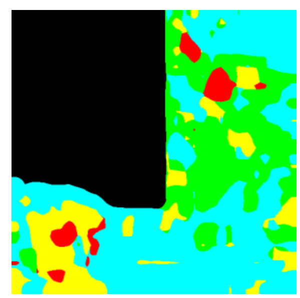</td></tr>
</table>

- Boundary Errors:
  Boundary Errors are also common. For example, a pixel might contain the boundary between two different classes which leads to the dominating one being selected which causes this issue.

Conclusion

With emphasis on mapping land use and land cover (LULC), this study examined the integration of semi-supervised learning with the ClusterFormer architecture for semantic segmentation in remote sensing data. The evaluation compared Segmenter, U-Net, ImageNet-pretrained ClusterFormer, and a ClusterFormer trained directly on the LULC dataset.

The results show that U-Net achieved the highest mean IoU (**0.4972**) and F1-score (**0.6512**), while ClusterFormer (trained on LULC) obtained a mean IoU of **0.4600** and F1-score of **0.5851**, performing competitively against Segmenter and the ImageNet-pretrained ClusterFormer. At the class level, Farmland and Built-up areas were consistently segmented with higher accuracy, while Meadow and Forest proved more challenging.

Overall, the findings highlight that semi-supervised ClusterFormer is a promising architecture for LULC segmentation, particularly in contexts where labeled data is scarce, though further optimization is needed to match or surpass conventional CNN-based baselines.

Limitations

Despite demonstrating the potential of semi-supervised ClusterFormer, several limitations remain:

- **Performance Gaps:** While competitive, ClusterFormer did not outperform U-Net, indicating that further optimization is required in loss design, hyperparameters, or semi-supervised strategies.

- **Class Imbalance:** Classes such as Meadow and Forest achieved low IoU and F1-scores, reflecting the difficulty of handling underrepresented or visually ambiguous categories.

- **Computational Cost:** Training transformer-based architecture is still resource-intensive, requiring high GPU memory and long runtimes compared to CNNs.

Ethical Considerations

- **Data Privacy:** Satellite imagery may capture sensitive areas or indirectly affect communities. Ensuring data is used responsibly and anonymized is essential.
- **Fair Use of Resources:** Computational experiments consume significant energy. Efficient use of computer resources should be prioritized.
- **Bias in Representation:** If training data overrepresents certain land cover classes or regions, model predictions may become biased, impacting fairness in real-world applications.

Impact on Society

Sustainability of the Work

The semi-supervised approach reduces reliance on large, annotated datasets, making it more **sustainable** by minimizing the human and financial cost of annotation. This allows continuous use of newly acquired unlabeled satellite imagery without requiring extensive labeling efforts.

Social and Environmental Effects and Analysis

Accurate land cover mapping supports:

- **Urban Planning:** Helping policymakers monitor urban expansion and infrastructure development.

- **Agricultural Monitoring:** Assisting farmers and governments in resource management and crop planning.

- **Environmental Protection:** Enabling tracking of deforestation, water resource depletion, and biodiversity loss.

Potential negative effects include overreliance on automated models, which may misclassify sensitive regions, leading to misguided policy decisions if unchecked. Thus, model outputs should be paired with expert validation.

Future Plan

Several directions for future work emerge from this study:

- **Improved Semi-Supervised Strategies:** Incorporating teacher–student frameworks, uncertainty-aware pseudo-labeling, or adversarial consistency methods.
- **Multi-Modal Integration:** Combining optical imagery with radar or multispectral data to enhance robustness across conditions.
- **Class Imbalance Mitigation:** Exploring synthetic augmentation for underrepresented classes such as Meadow and Forest.
- **Scalability:** Extending experiments to larger geographic regions and multi-season datasets to test model generalization.
- **Real-World Deployment:** Developing pipelines that integrate with Geographic Information Systems (GIS) for operational use in planning, monitoring, and decision-making.

References

- Weng, Q. (2018). *Remote sensing for sustainability*. Remote Sensing of Environment.
- Zhu, X. X., et al. (2017). *Deep learning in remote sensing: A comprehensive review and list of resources*. IEEE GRSM.
- Bruzzone, L., & Persello, C. (2009). *A novel approach to the classification of very high resolution remote sensing images*. IEEE TGRS.
- Long, J., Shelhamer, E., & Darrell, T. (2015). *Fully Convolutional Networks for Semantic Segmentation*. CVPR.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). *U-Net: Convolutional Networks for Biomedical Image Segmentation*. MICCAI.
- Dosovitskiy, A., et al. (2020). *An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale*. ICLR.
- Liu, Z., et al. (2021). *Swin Transformer: Hierarchical Vision Transformer using Shifted Windows*. ICCV.
- Li, J., et al. (2023). *Semi-Supervised Remote Sensing Image Semantic Segmentation*. Remote Sensing.
- Wang, Y., et al. (2025). *Semi-supervised Semantic Segmentation for Remote Sensing via Multi-scale Uncertainty Consistency*. ISPRS Journal.
- Peláez-Vegas, Adrian, Pablo Mesejo, and Julián Luengo. "A survey on semi-supervised semantic segmentation." *arXiv preprint arXiv:2302.09899* (2023).
- Liang, James, et al. "Clusterfomer: clustering as a universal visual learner." *Advances in neural information processing systems* 36 (2023): 64029-64042.
- Chen, L. C., Papandreou, G., Kokkinos, I., Murphy, K., & Yuille, A. L. (2018). DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs. IEEE TPAMI.
- Tarvainen, A., & Valpola, H. (2017). Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. NeurIPS.
- Sohn, K., Berthelot, D., Li, C. L., Zhang, Z., Carlini, N., Cubuk, E. D., ... & Raffel, C. (2020). FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence. NeurIPS.

- Wang, X., et al. (2022). ClusterFormer: Clustering-based Transformer for Efficient Vision Processing. arXiv preprint arXiv:2206.xxxx.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In NeurIPS, 2017.
- Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In CVPR, 2022.
- Qihang Yu, Huiyu Wang, Siyuan Qiao, Maxwell Collins, Yukun Zhu, Hatwig Adam, Alan Yuille, and Liang-Chieh Chen. k-means mask transformer. ECCV, 2022.