

OPTIMUS: Observing Persistent Transformations in Multi-temporal Unlabeled Satellite-data

Raymond Yu^{*1}, Paul Han^{*1}, Piper Wolters², and Favyen Bastani²

¹University of Washington

²Allen Institute for AI

{ryu5, paulh27}@cs.washington.edu, {piperw, favyenb}@allenai.org

Abstract

In the face of pressing environmental issues in the 21st century, monitoring surface changes on Earth is more important than ever. Large-scale remote sensing, such as satellite imagery, is an important tool for this task. However, using supervised methods to detect changes is difficult because of the lack of satellite data annotated with change labels, especially for rare categories of change. Annotation proves challenging due to the sparse occurrence of changes in satellite images. Even within a vast collection of images, only a small fraction may exhibit persistent changes of interest. To address this challenge, we introduce OPTIMUS, a self-supervised learning method based on an intuitive principle: if a model can recover information about the relative order of images in the time series, then that implies that there are long-lasting changes in the images. OPTIMUS demonstrates this principle by using change point detection methods on model outputs in a time series. We demonstrate that OPTIMUS can directly detect interesting changes in satellite images, achieving an improvement in AUROC score from 56.3% to 87.6% at distinguishing changed time series from unchanged ones compared to baselines. Our code and dataset are available at <https://huggingface.co/datasets/optimus-change/optimus-dataset/>.

1. Introduction

Globally detecting long-lasting changes to the Earth's surface is critical for informing decisions around tackling looming environmental, climate, and conservation challenges [4, 7, 10, 18]. Monitoring deforestation helps to understand where and why forest loss is happening; tracking urban expansion helps to quantify the environmental effects of urban sprawl; and identifying areas impacted by natural

disasters like wildfires and earthquakes helps to target disaster relief efforts. Satellite imagery offers frequent views of any location on Earth; for example, the European Space Agency's Sentinel-2 mission provides images of land and coastal locations across the globe every five days.

Numerous computer vision datasets [3, 5, 23] and methods have been proposed for detecting changes in the tens of millions of square kilometers of the Earth that are imaged daily. However, the vast majority of these datasets focus narrowly on urban changes like new buildings and roads. This focus arises because annotating other types of long-lasting changes is expensive due to their rarity: most terrestrial locales exhibit minimal change over time, so even within a sizable dataset of satellite image time series, only a handful of time series may show certain changes like forest loss or the impact of wildfires. Furthermore, annotation demands specialized knowledge, as interpreting satellite images (especially outside urban areas) is not always straightforward. Additionally, the sheer number of potential categories of change makes it infeasible to develop a truly comprehensive dataset.

Because of these challenges—high annotation costs, the rarity of substantial changes, and the specialized knowledge required—many researchers have resorted to unsupervised methods for detecting changes in satellite image time series. A few unsupervised methods have been proposed, such as CaCo [20], which learns representations that diverge when a location undergoes change. However, these methods struggle to distinguish seasonal changes, such as deciduous trees changing color or crop fields being harvested, from long-lasting changes that permanently alter the Earth's surface.

To address these challenges, we propose OPTIMUS (Observing Persistent Transformations in Multi-temporal Unlabeled Satellite-data), a self-supervised learning method for classifying the presence of long-lasting change in satellite image time series. We focus on persistent changes, i.e., changes that have a visible impact on a

^{*}Equal contribution.

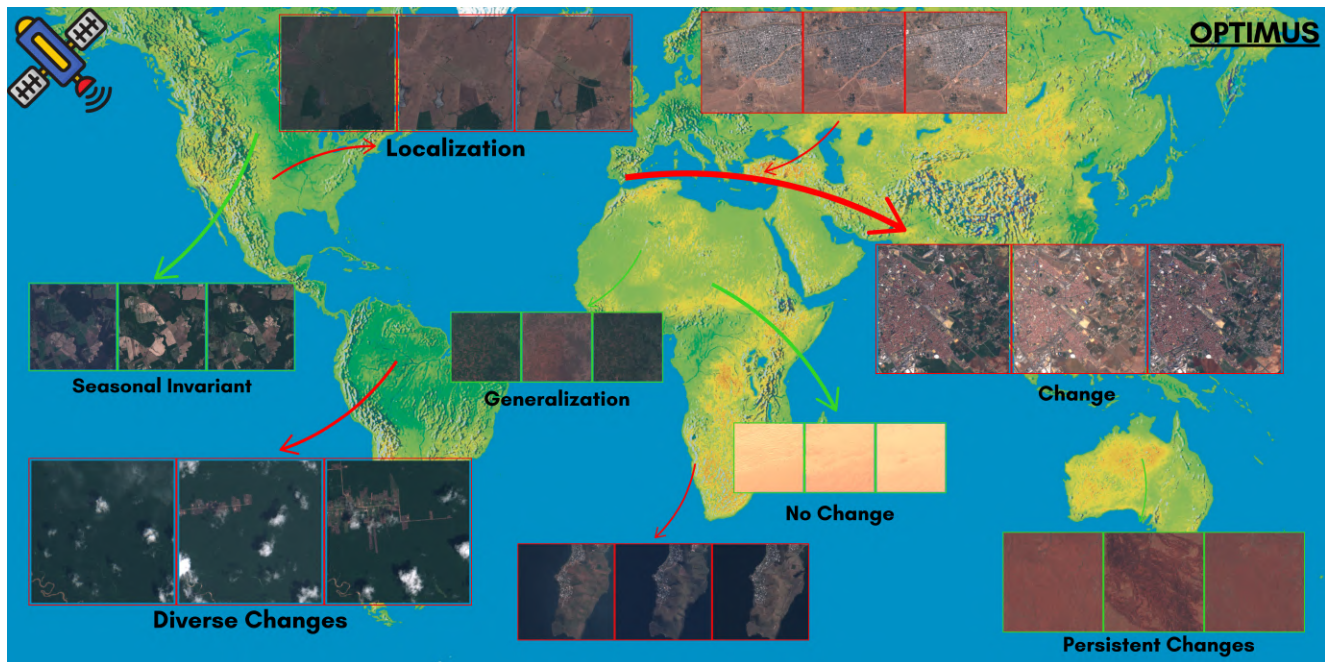


Figure 1. Overview of binary outputs (red for change, green for no change) generated by OPTIMUS across global regions. The figure showcases key capabilities: localization of changes into 128x128 images (top left) for precise spatial detection, identification of diverse changes such as deforestation in the Amazon Rainforest (red, label = 1) and shifting sand dunes in Chad (green, label = 0), and generalization across varied contexts (middle), including urban expansion in a city (red, label = 1) and lake recession due to drought (red, label = 1). OPTIMUS also demonstrates robustness to seasonal invariance, correctly identifying stable conditions such as vegetation changes in agricultural regions (green, label = 0).

location lasting longer than one year. OPTIMUS is based on the intuition that if a model can determine the correct long-term ordering of images in a time series, even when the sequence is hidden, it indicates the presence of significant, persistent changes in the images. Otherwise, if the location remained constant, the long-term ordering should be indecipherable. OPTIMUS only requires a collection of unlabeled satellite image time series for training.

Given a collection of image time series, OPTIMUS produces the subset of time series that underwent some long-lasting change. This enables users interested in analyzing changes in certain geographies to hone in on the small subset of patches within that geography that changed. In addition to classifying the presence of long-lasting changes, we show that simple extensions enable OPTIMUS to localize the changes both in space and in time within the image time series. Although OPTIMUS does not categorize the changes that it detects, category-specific annotation is substantially cheaper over the filtered subset, making it feasible to develop models that identify rare types of changes. Thus, OPTIMUS can be used to equip decision-makers with reliable data on the locations of changes like wildfires, flooding, deforestation, and so on, to help address and mitigate

the impacts of global environmental challenges.

We developed a dataset consisting of one million image time series for training, along with a set of time series that we annotated with binary classification labels (has change or no change). Our zero-shot results on the test set show that OPTIMUS significantly outperforms other unsupervised methods, improving the AUROC score at distinguishing changed time series from 56.3% to 87.6%. OPTIMUS helps users hone in on locations with persistent changes, allowing for a smaller, more manageable subset of images that can then be manually categorized by experts. This allows OPTIMUS to be a solution for wide-reaching applications in environmental monitoring, urban planning, disaster response, and conservation efforts.

2. Related Works

2.1. Existing Datasets

Several change detection datasets have been released, combining bitemporal satellite images or image time series with patchwise or pixelwise change labels. OSCD [5], Cropland-CD [15], and SYSU-CD [23] provide binary pixelwise change masks for bitemporal images, with most labels corresponding to new buildings, new roads, and con-

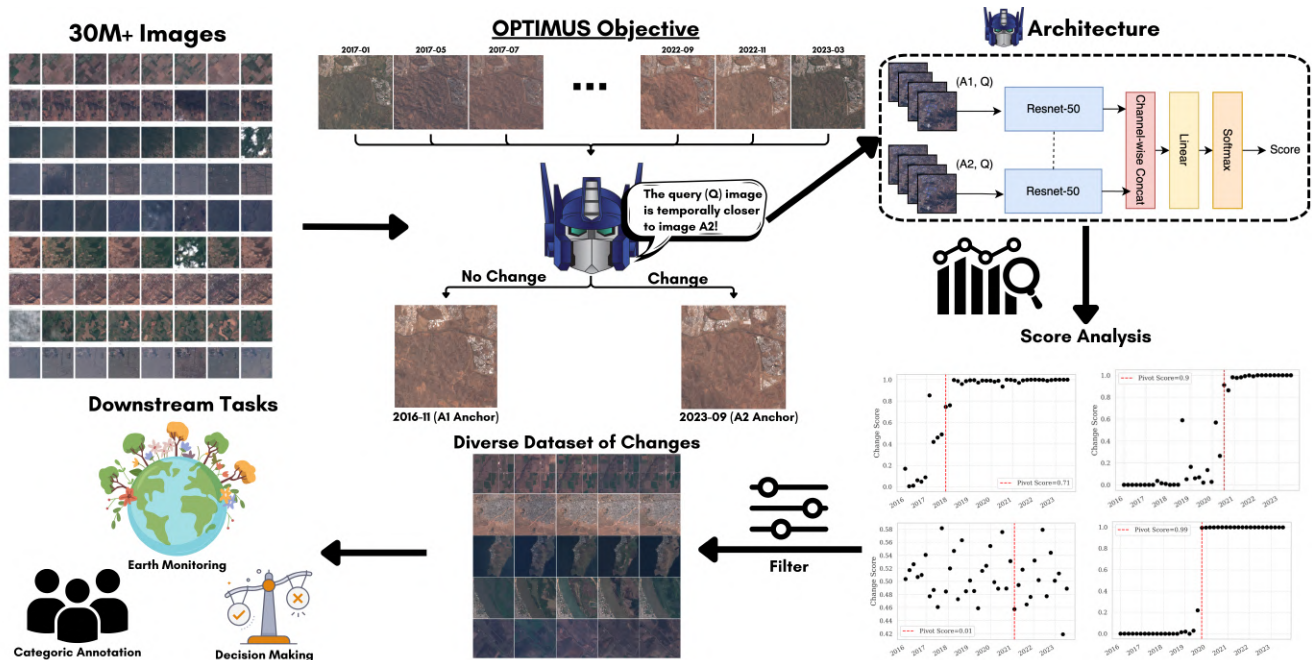


Figure 2. Overview of the OPTIMUS framework for detecting persistent changes in multi-temporal satellite imagery. The framework processes over 30 million satellite images, leveraging a Siamese network architecture to classify whether a query image (Q) is temporally closer to one of two anchor images (A1 or A2). This self-supervised approach generates a diverse dataset of annotated changes, distinguishing "Change" and "No Change" events while filtering out seasonal variations. Downstream tasks include environmental monitoring and decision-making, enabling actionable insights from large-scale satellite data.

struction work. Some datasets focus on individual categories. EGY-BCD [9], SI-BU [14], and S2Looking [22] have multiclass pixelwise change labels for constructed and demolished buildings. Hi-UCD [24] consists of labels for nine categories of change in images of Tallinn, including newly constructed buildings, greenhouses, and roads.

While these datasets are invaluable for change detection, they often focus on urban areas due to the relative ease of detecting changes in these regions. Acquiring labeled data for more complex environmental changes, like forest loss, wildfires, or desertification, is both costly and labor-intensive. These changes are subtle, long-term transformations that require specialized knowledge to detect and annotate accurately, unlike the more visually distinct changes in urban areas.

Some datasets provide multi-temporal data for land cover types. DynamicEarthNet [25] provides monthly pixelwise labels for seven land use and land cover (LULC) classes, focusing on environmental changes like vegetation shifts and urbanization. DFC2021 Track MSD [13] targets land cover transitions using multisensor data fusion. However, these datasets are often limited due to the cost and are primarily centered around urban areas, making scalability and broader applicability an issue.

2.2. Supervised Methods for Change Detection

Supervised methods for change detection have demonstrated significant success in identifying changes in satellite imagery, but they rely on large amounts of labeled data, which presents scalability challenges. For instance, ChangeNet [11] leverages deep neural networks with pre-trained weights, achieving high accuracy by fine-tuning on specific tasks. Despite its strong performance, ChangeNet requires extensive annotated datasets to generalize effectively, which is a notable limitation for broader applications. Similarly, the framework proposed by Wu et al [27] focuses on analyzing image pairs, restricting its utility in detecting long-term changes across extended time series.

Additionally, many supervised methods, including Zheng et al's. [30] framework, are targeted towards urban change detection where data is more abundant. This urban-centric focus further restricts their broader environmental applicability of the models.

While pre-trained models [1] can reduce the amount of labeled data needed for downstream tasks like change detection, supervised methods still face significant challenges in terms of generalizability and scalability due to the lack of labeled data, especially for rare categories of change like desertification, shrub regrowth, and logging.

2.3. Unsupervised Methods for Change Detection

Given the scarcity of labeled satellite data, unsupervised methods have gained traction in recent years. These approaches aim to detect changes without the need for annotated datasets, making them more scalable and applicable to various scenarios.

SeCo [21] is an unsupervised method pre-training method for remote sensing data that leverages temporal and positional invariance to learn transferable representations. Similarly, SSL4EO [26] utilizes self-supervised learning techniques to learn representations of satellite images that can be used to detect changes. Both SeCo and SSL4EO have shown effectiveness in identifying changes across diverse environments; however, they often require fine-tuning to optimize performance for specific tasks and datasets.

On the other hand, methods like CaCo [20] do not require fine-tuning. CaCo uses contrastive learning to differentiate between unchanged and changed areas in satellite imagery by obtaining changes using features extracted from the images. Specifically, CaCo addresses the issue of varying seasonal and location-specific changes by normalizing the distance between feature representations of images. It calculates a ratio of distances between long-term images and short-term seasonal images to normalize scaling differences across different locations. This ratio helps distinguish actual changes from seasonal variations. Despite this sophisticated approach, CaCo struggles with effectively separating long-lasting changes from environments with pronounced seasonal changes due to its reliance on feature distance ratios, which may not always capture the true magnitude of persistent changes.

Recent work by Mall et al. [19] introduces an event-driven change detection approach, which focuses on detecting meaningful change events from spatio-temporal satellite imagery in Cairo and California. Their method addresses the challenge of discovering significant changes from vast amounts of time-series data, similar to our objective of capturing persistent transformations. By focusing on events rather than pixelwise differences, this approach aligns with the need for methods that generalize across broader contexts.

Our approach differs from existing methods [30] [2] [28] [29] by leveraging the full time series of satellite imagery in a fully unsupervised manner. This allows our model to better identify persistent changes, as it can learn from the temporal context provided by the entire series. By focusing on persistent changes, our methods aim to filter out large datasets, enabling experts to concentrate on images with significant and lasting alterations. This approach addresses some of the limitations of both supervised and unsupervised techniques, providing a novel method for identifying and monitoring significant environmental changes over time.

3. Dataset

Most existing change detection datasets, such as OSCD and ChangeNet, focus predominantly on urban areas. This urban-centric sampling is adopted to achieve a higher rate of detectable changes, making annotation more feasible. However, this approach limits the variety of changes captured, excluding many non-urban changes such as wildfires, desertification, and agricultural expansion.

To address these limitations, we compile a new dataset from publicly accessible sources, Sentinel-2 and NAIP, with the goal of detecting arbitrary persistent changes across diverse geographical locations. Unlike existing datasets, our approach involves randomly sampling image patches from across the entire globe, including non-urban areas, to ensure a broader spectrum of changes is captured. This strategy enables us to annotate a more diverse set of changes, avoiding the bias toward urban areas.

To make this global sampling feasible, we focus on classifying time series of image patches for change rather than segmenting individual pixels. This is particularly advantageous for capturing non-urban changes, where drawing precise segmentation labels can be challenging. For instance, changes such as wildfires and desertification can have diffuse boundaries that are difficult to delineate precisely.

For model training, we retain only the RGB channels (B2, B3, B4) and applied a cloud cover filter to exclude images with greater than 20% cloud cover. Each image in our dataset measures 512x512 pixels at a 10-meter/pixel resolution, allowing our models to detect fine-grained changes that may occupy only a few pixels. The dataset comprises one million time series, each representing the spatio-temporal context of a geographic location. Each time series includes 30-48 satellite images captured between January 2016 and December 2023, with a minimum interval of 2 months between any two images. This extensive temporal span of 8 years ensures the dataset captures a wide range of long-term changes, including deforestation and desertification, providing a robust basis for training models to detect persistent changes.

3.1. Evaluation Set

To evaluate our change detection models, we constructed an evaluation set of 300 satellite time series, each consisting of 512x512 images, randomly sampled from our dataset. A random subset of 100 of these images were further divided into 1600 128x128 patches, with each patch receiving a binary label. Each time series in this set is assigned a binary label: positive labels indicate the presence of a persistent and non-cyclic change (i.e., not seasonal changes), while negative labels indicate no change. Approximately one-third of the time series are positively labeled.

Given the complexity and diversity of changes we aim to detect, we prioritized the quality of annotations over quan-

tity. To ensure high-quality annotations, we provided annotators with detailed instructions on how to identify and label persistent changes, emphasizing consistency across different types of changes, especially those with ambiguous boundaries, such as desertification or gradual urban expansion. Annotators were instructed to focus on clear indicators of non-cyclic change, disregarding seasonal variations or temporary phenomena.

Appendix A provides qualitative examples of the annotated changes, illustrating the types of changes our models are designed to detect. These examples are accompanied by the specific instructions given to annotators, offering insight into the criteria used for labeling.

The evaluation task involves the models generating a change score for each time series in the evaluation set and computing the task accuracy. This evaluation closely mirrors practical downstream applications, where the objective is to identify locations with persistent changes over many years. Although using a binary label is a relatively simple approach—lacking the granularity of change categories or segmentation masks—it serves as an effective filtering mechanism and remains effective for types of changes like long-term droughts that have nebulous boundaries. This simplicity makes it cost-efficient to extract a diverse set of changes in future applications.

We will release the unsupervised training data, labeled test set, and code (which is also included in the supplementary material) under an open license.

4. Methods

OPTIMUS determines whether remote sensing satellite time series contain non-seasonal changes by hiding the ordering of the images in the series, and then training a model to attempt to recover the long-term ordering. If there are only seasonal changes but no persistent changes, then the long-term ordering should be indecipherable from the images alone. An example is lake levels changing due to seasonal precipitation. As this change is cyclic, while the images can be grouped into seasons, the ordering of the images across years cannot be distinguished. On the other hand, if a location exhibits persistent changes, it should be possible to distinguish all of the images captured before a change from those captured after a change. Road construction is such an example, with distinct stages like laying the subgrade, base, and asphalt.

In this section, we first describe a basic implementation of this intuition. We then identify flaws in the basic implementation, and detail how OPTIMUS addresses those flaws.

In the basic implementation, given a time series $\langle I_1, I_2, \dots, I_n \rangle$, we train a binary classifier (denoted as b) to predict whether an arbitrary intermediate image Q is closer in time to I_1 or I_n . The classifier inputs a tuple (I_1, I_n, Q) of the images only, with the timestamp of Q hidden, and

outputs a confidence score that Q is closer to I_n than I_1 . During training, examples are constructed by (1) sampling a time series, (2) picking an arbitrary Q between I_1 and I_n , and (3) computing the label based on whether Q is before or after $I_{n/2}$.

After training, to determine if a new time series contains change, we compute the confidence score from the model for every image in the time series, i.e. we compute a time series of scores $S = \{b(I_1, I_n, I_j) | j = 1, 2, \dots, n\}$. An oracle classifier would output a step function where the score switches from 0 to 1 halfway through the time series. If the time series contains change, then an effective model should perform similarly to the oracle (Figure 3, top left). However, if the time series contains no persistent changes, then the scores should fluctuate arbitrarily, since the model does not have sufficient information (Figure 3, bottom left). Then, we apply measures on S that broadly assess the degree to which it is monotonically increasing.

However, there are several flaws with this basic implementation. First, a location may undergo a change at a single timestamp, e.g. trees are logged within the span of a month. Suppose this change occurs between I_1 and I_2 . Then, during training, we would be training the model to predict that $\langle I_2, \dots, I_{n/2-1} \rangle$ are closer in time to I_1 , even though they are closer in appearance to I_n . Second, if I_1 or I_n are low in quality, due to clouds, shadows, or imaging artifacts, then the quality of the scores in S would be lowered. Below, we address these flaws, and also detail the measures that we use to capture a final change score based on the confidence scores in S , along with the model architecture.

Training Example Selection. To train the classifier, we sample triplets (A_1, A_2, Q) for each time series in the dataset, and train the model to predict whether Q is closer to A_1 or A_2 . Selecting $A_1 = I_1$, $A_2 = I_n$, and Q as a random image from within the time series is the simplest approach, but has flaws as mentioned above.

The first concern is that a change may occur just after A_1 or just before A_2 , causing a randomly chosen Q to be temporally closer to one anchor but visually more similar to the other. To address this, during training, rather than selecting a query image Q between A_1 and A_2 , we select Q to be before A_1 or after A_2 with equal probability; this guarantees that, with respect to persistent changes, Q will always be visually more similar to the same anchor that it is temporally closer to.

Second, to address the concern of low quality images due to environmental or photometric distortions, rather than using a single image for the anchors A_1 and A_2 , we actually provide the model with multiple consecutive images. Specifically, each anchor consists of $c = 3$ consecutive satellite images from the time series. Using multiple consecutive images for A_1 and A_2 reduces the likelihood that

all images in a set will be affected by distortions, increasing the robustness of the model. However, there is a trade-off: if the anchor sets are too long, there is a risk that changes may occur within A_1 or A_2 themselves, potentially complicating the model’s learning process. Through ablation studies, we determined that using three consecutive images strikes an optimal balance, providing sufficient robustness while minimizing the risk of internal changes within the anchor sets. More details on the input construction can be found in the Appendix C.

Model Architecture. For each triplet (A_1, A_2, Q) , we construct two tensors that the model processes independently. The first tensor is formed by concatenating the query image Q with each image in the anchor A_1 , resulting in an $3 \times (c + 1) \times 512 \times 512$ tensor. Similarly, the second tensor is created by concatenating Q with the images from anchor A_2 . These tensors are denoted as (A_1, Q) and (A_2, Q) , respectively. They are then input to the model, which is tasked with predicting whether the query Q is temporally closer to A_1 or A_2 . The model is trained using binary cross-entropy loss, defined as:

$$\omega = b((A_1, Q), (A_2, Q))$$

$$L_{\text{OPTIMUS}} = -y \log(\omega) - (1 - y) \log(1 - \omega)$$

where ω is the output of the binary classifier and $y \in \{0, 1\}$ is the ground truth label indicating whether Q is closer to A_1 ($y = 0$) or A_2 ($y = 1$).

The model architecture uses a Siamese Neural Network [12] with a Resnet-50 backbone [8] to map the tensors into a shared embedding space. The embeddings are concatenated, passed through a linear layer, and then processed by a softmax function to yield a score between 0 and 1. The model is trained using the AdamW [17] optimizer with a learning rate of 3×10^{-4} , a batch size of 5, and for 5 epochs. The dataset is partitioned into training and validation sets with an 80/20 split. The training objective and model architecture are depicted in Figure 4.

Change Score Measures. Given a time series $\langle I_1, I_2, \dots, I_n \rangle$, suppose we calculated the series of scores $S = \{b(I_1, I_n, I_j) | j = 1, 2, \dots, n\}$. We use two measures to quantify persistent changes in the sequence. The first is the Spearman rank correlation coefficient:

$$\rho = 1 - \frac{6 \sum_{i=1}^{|S|} (\text{rank}(s_i) - i)^2}{|S|(|S|^2 - 1)}$$

This metric assesses the monotonicity of S by measuring how well the ranks of the scores correspond to their positions in the sequence. The second measure is the pivot score P :

$$P = \max_{i=1}^n \left| \frac{\sum_{j=1}^i s_j}{i} - \frac{\sum_{j=i+1}^n s_j}{n - i} \right|$$

This score identifies the index i in S that maximizes the absolute difference between the average values of the segments before and after i . A high pivot score indicates a significant and persistent change in the sequence, typically aligning with abrupt transitions. For a visualization of the scoring process, see Figure 3.

In practice, as shown by an ablation study in Appendix C, the pivot score outperformed the Spearman coefficient, so all subsequent results and experiments report scores based on the pivot score.

4.1. Spatial Localization of Changes

OPTIMUS operates effectively as a general framework for large 512×512 images, corresponding to a $5\text{km} \times 5\text{km}$ region at 10m resolution. However, this extensive spatial context complicates the precise localization of regions or pixels where changes occur. For instance, urbanization might be confined to a small section of the image due to the construction of a few buildings. Running OPTIMUS on smaller patches to localize changes in a time series significantly degrades performance because the model was trained on large spatial contexts and cannot generalize well to smaller patches.

A straightforward solution is to train the model on 128×128 patches to adapt to the lower spatial context. However, this approach introduces the challenge of sparse changes when subdividing time series into smaller patches. Most regions in an image do not exhibit change, resulting in a low signal-to-noise ratio during model training.

To address this issue, we propose an iterative approach to extend OPTIMUS for spatial localization. Initially, we train OPTIMUS on the entire dataset. We then apply the trained model to filter the top 50% of time series that show the most persistent changes, thereby increasing the density of changes in the dataset and enhancing the model’s ability to learn from significant changes. Subsequently, we retrain OPTIMUS on this filtered dataset using 128×128 patches.

5. Experiments

5.1. Change Detection

We first evaluate OPTIMUS against baselines on distinguishing 512×512 time series exhibiting persistent changes from those with no changes in our evaluation set.

Baselines. We compare OPTIMUS to three baselines: CaCo, SeCo, and OSCD pre-training.

CaCo is a contrastive method specifically designed to detect long-term changes while being insensitive to seasonal variations. The authors apply CaCo for classifying changes by computing the distance between feature representations. To classify a time series using CaCo, we randomly sampled

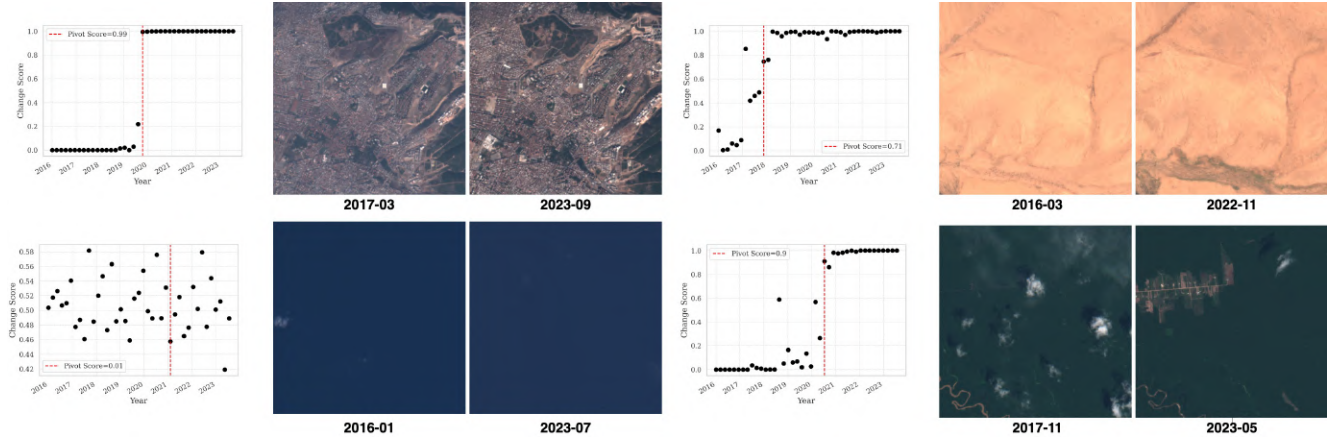


Figure 3. Analysis of pivot scores and corresponding satellite images for different environmental changes. The top left shows a high pivot score indicating significant urban expansion with a marked spike in change, suggesting a sudden development. The bottom left illustrates ocean images associated with an extremely low pivot score, indicating minimal change. The top right depicts a desert region experiencing reforestation, with shrubs growing back. The bottom right shows a region undergoing deforestation.

three image pairs from the first and last years, then computed the average distance between their representations. A larger distance indicated a greater change score, relative to the other time series. Following the authors’ recommendations, we also normalized the distances between feature vectors representing long-term and short-term images, as the feature distances can scale differently depending on the type of location.

CaCo [20] is the only recent method geared for change classification, since most methods focus on segmentation. Thus, we make adaptations to compare the other two baselines.

Unlike CaCo, SeCo [21] is a contrastive model intended for seasonally variant and invariant downstream applications. As such, it produces embeddings with three independent subspaces, including one variant and another invariant to seasonal changes. Since our task focuses on detecting non-seasonal changes, we explicitly remove the seasonally variant subspace and only use the other two subspaces to quantify changes. Following this, however, the evaluation method for SeCo is kept the same as CaCo.

The third baseline consists of pre-training on OSCD [5], a dataset consisting of bi-temporal images paired with change masks. After pre-training, we compare feature representations to classify change. Since CaCo reports the highest performance on OSCD, we use it for the model, fine-tuning it on OSCD following the authors’ training procedure. To evaluate on our dataset, change scores for each time series are calculated as before, using distances between representation vectors.

Metrics. For each baseline method, a change score was computed for each series in the evaluation set, and a threshold was applied to distinguish predicted changes from un-

Table 1. Comparative Evaluation versus Baseline Methods

Method	F1 Score	AUROC
SeCo	0.540	0.491
CaCo	0.551	0.494
CaCo (OSCD)	0.585	0.563
OPTIMUS	0.760	0.876

changed series. Performance was then measured in terms of accuracy against ground truth labels. Given the varying ranges of score values produced by different methods, the primary performance metric used is the non-thresholded Area Under the Receiver Operating Characteristic Curve (AUROC). In addition, we report the maximum F1 score across all thresholds.

Quantitative Results. Table 1 shows the performance metrics for all evaluated methods on the evaluation set. OPTIMUS achieves a substantial performance advantage over all other methods. We attribute this to OPTIMUS’s approach of explicit classification of images based on temporal locality, which directly encodes changes into the feature representations. However, contrastive methods such as SeCo and CaCo only indirectly capture changes through differences in feature representations, resulting in less reliable outcomes. Lastly, supervised pretraining on OSCD increases the performance for CaCo, but it still falls far short of OPTIMUS. This gain can be attributed to CaCo’s direct training for urban change detection, which helps detect urban images in our evaluation set. However, the overall effectiveness is limited due to the broader range of changes in our evaluation dataset beyond urban environments.



Figure 4. Illustration of the proposed pipeline for fine-grained change detection. Satellite image time series are divided into 128x128 patches to isolate localized changes, enabling the OPTIMUS model to classify each patch as exhibiting meaningful changes (e.g., urban expansion) or no significant change (e.g., seasonal variations). This approach addresses challenges in detecting changes across diverse components within larger images by focusing on spatially localized transformations.

5.2. Localization

In Table 2, we present the results of the localized evaluation for OPTIMUS. Various versions of OPTIMUS are compared to demonstrate the necessity of the iterative training procedure outlined in Section 4.1. The original OPTIMUS model, trained on 512×512 images, performs poorly with an AUROC score when applied to smaller patches. Training OPTIMUS directly on 128×128 images without the iterative procedure results in an improvement in AUROC. However, the best performance is achieved with OPTIMUS trained using the iterative method, which increases AUROC significantly. This improvement is expected, given that the iterative OPTIMUS trains on a dataset with a much higher density of changes. The fairly high AUROC for the iterative OPTIMUS method indicates its effectiveness in localizing regions of interest for potential persistent changes. While the overall performance decreases compared to 512×512 images, this is anticipated due to the reduced spatial context.

Table 2. Comparative Localized Evaluation

Method	F1 Score	AUROC
OPTIMUS (512x512)	0.587	-
OPTIMUS (128x128)	0.615	0.8714
OPTIMUS (512x512 → 128x128 iterative)	0.658	0.9415

5.3. Qualitative Examples

Figure 3 provides examples of change score series generated by OPTIMUS on the evaluation dataset, along with the outputs of the pivot scores. It also shows the exact time at which the pivot score detected the greatest degree

of change. The results illustrate that the OPTIMUS accurately measures the degree of change and can often pinpoint precisely the location of the greatest change, where the partitioned sets of images before and after the pivot exhibit the greatest disparity. Qualitative examples are located in Appendix B.

6. Conclusion

Change detection in remote sensing is essential for applications such as monitoring deforestation, tracking urban expansion, and assessing the impact of natural disasters. However, supervised methods struggle with the scarcity and high cost of obtaining labeled data, particularly for non-urban areas. In this paper, we introduced OPTIMUS, an unsupervised approach that identifies persistent, long-term changes in satellite image time series by learning temporal patterns. OPTIMUS improves the AUROC score from 56.3% to 87.6%, significantly outperforming previous methods in distinguishing significant changes from seasonal variations.

References

- [1] Favyen Bastani, Piper Wolters, Ritwik Gupta, Joe Ferdinando, and Aniruddha Kembhavi. SatlasPretrain: A Large-Scale Dataset for Remote Sensing Image Understanding, 2023. 3, 11
- [2] Csaba Benedek and Tamás Sziranyi. Change Detection in Optical Aerial Images by a Multilayer Conditional Mixed Markov Model. *IEEE Transactions on Geoscience and Remote Sensing*, 47(10):3416–3430, 2009. 4
- [3] Hao Chen and Zhenwei Shi. A spatial-temporal attention-based method and a new dataset for remote sensing image change detection. *Remote Sensing*, 12(10):1662, 2020. 1
- [4] Rodrigo Caye Daudt, Bertr Le Saux, Alexandre Boulch, and Yann Gousseau. Urban Change Detection for Multispectral

- Earth Observation Using Convolutional Neural Networks. In *IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium*, pages 2115–2118, 2018. 1
- [5] Rodrigo Caye Daudt, Bertrand Le Saux, Alexandre Boulch, and Yann Gousseau. Urban Change Detection for Multi-spectral Earth Observation Using Convolutional Neural Networks, 2018. 1, 2, 7
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 11
- [7] Ritwik Gupta, Richard Hosfelt, Sandra Sajeev, Nirav Patel, Bryce Goodman, Jigar Doshi, Eric Heim, Howie Choset, and Matthew Gaston. xBD: A Dataset for Assessing Building Damage from Satellite Imagery, 2019. 1
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition, 2015. 6
- [9] Shima Holail, Tamer Saleh, Xiongwu Xiao, and Deren Li. AFDE-Net: Building Change Detection using Attention-Based Feature Differential Enhancement for Satellite Imagery. *IEEE Geoscience and Remote Sensing Letters*, 20:1–5, 2023. 3
- [10] Masroor Hussain, Dongmei Chen, Angela Cheng, Hui Wei, and David Stanley. Change detection from remotely sensed images: From pixel-based to object-based approaches. *ISPRS Journal of Photogrammetry and Remote Sensing*, 80:91–106, 2013. 1
- [11] Deyi Ji, Siqi Gao, Mingyuan Tao, Hongtao Lu, and Feng Zhao. ChangeNet: Multi-Temporal Asymmetric Change Detection Dataset, 2024. 3
- [12] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese Neural Networks for One-shot Image Recognition. In *ICML Deep Learning Workshop*, 2015. 6
- [13] Zhuohong Li, Fangxiao Lu, Hongyan Zhang, Lilin Tu, Jiayi Li, Xin Huang, Caleb Robinson, Nikolay Malkin, Nebojsa Jojic, Pedram Ghamisi, Ronny Hänsch, and Naoto Yokoya. The Outcome of the 2021 IEEE GRSS Data Fusion Contest—Track MSD: Multitemporal Semantic Change Detection. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15:1643–1655, 2022. 3
- [14] Cheng Liao, Han Hu, Xuekun Yuan, Haifeng Li, Chao Liu, Chunyang Liu, Gui Fu, Yulin Ding, and Qing Zhu. BCE-Net: Reliable Building Footprints Change Extraction based on Historical Map and Up-to-date Images using Contrastive Learning. *ISPRS Journal of Photogrammetry and Remote Sensing*, 201:138–152, 2023. 3
- [15] Mengxi Liu, Zhuoqun Chai, Haojun Deng, and Rong Liu. A CNN-transformer network with multiscale context aggregation for fine-grained cropland change detection. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15:4297–4306, 2022. 2
- [16] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, Furu Wei, and Baining Guo. Swin transformer v2: Scaling up capacity and resolution, 2022. 11
- [17] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization, 2019. 6
- [18] Sahel Mahdavi, Bahram Salehi, Weimin Huang, Meisam Amani, and Brian Brisco. A PolSAR Change Detection Index Based on Neighborhood Information for Flood Mapping. *Remote Sensing*, 11(16), 2019. 1
- [19] Utkarsh Mall, Bharath Hariharan, and Kavita Bala. Change Event Dataset for Discovery from Spatio-temporal Remote Sensing Imagery. In *NeurIPS Datasets and Benchmarks Track*, 2022. 4
- [20] Utkarsh Mall, Bharath Hariharan, and Kavita Bala. Change-Aware Sampling and Contrastive Learning for Satellite Images. In *CVPR*, 2023. 1, 4, 7
- [21] Oscar Mañas, Alexandre Lacoste, Xavier Giro i Nieto, David Vazquez, and Pau Rodriguez. Seasonal contrast: Unsupervised pre-training from uncurated remote sensing data, 2021. 4, 7
- [22] Li Shen, Yao Lu, Hao Chen, Hao Wei, Donghai Xie, Jiabao Yue, Rui Chen, Shouye Lv, and Bitao Jiang. S2Looking: A Satellite Side-Looking Dataset for Building Change Detection. *Remote Sensing*, 13(24):5094, 2021. 3
- [23] Qian Shi, Mengxi Liu, Shengchen Li, Xiaoping Liu, Fei Wang, and Liangpei Zhang. A deeply supervised attention metric-based network and an open aerial image dataset for remote sensing change detection. *IEEE transactions on geoscience and remote sensing*, 60:1–16, 2021. 1, 2
- [24] Shiqi Tian, Ailong Ma, Zhuo Zheng, and Yanfei Zhong. Hi-UCD: A Large-scale Dataset for Urban Semantic Change Detection in Remote Sensing Imagery. *arXiv preprint arXiv:2011.03247*, 2020. 3
- [25] Aysim Toker, Lukas Kondmann, Mark Weber, Marvin Eisenberger, Andrés Camero, Jingliang Hu, Ariadna Pregel Hoderlein, Çağlar Şenarar, Timothy Davis, Daniel Cremers, Giovanni Marchisio, Xiao Xiang Zhu, and Laura Leal-Taixé. DynamicEarthNet: Daily Multi-Spectral Satellite Dataset for Semantic Change Segmentation, 2022. 3
- [26] Yi Wang, Nassim Ait Ali Braham, Zhitong Xiong, Chenying Liu, Conrad M Albrecht, and Xiao Xiang Zhu. SSL4EO-S12: A Large-Scale Multi-Modal, Multi-Temporal Dataset for Self-Supervised Learning in Earth Observation, 2023. 4
- [27] Chen Wu, Bo Du, and L. Zhang. Fully Convolutional Change Detection Framework With Generative Adversarial Network for Unsupervised, Weakly Supervised and Regional Supervised Change Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45:9774–9788, 2022. 3
- [28] Kunping Yang, Gui-Song Xia, Zicheng Liu, Bo Du, Wen Yang, Marcello Pelillo, and Liangpei Zhang. Asymmetric Siamese Networks for Semantic Change Detection in Aerial Images. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–18, 2022. 4
- [29] Chenxiao Zhang, Peng Yue, Deodato Tapete, Liangcun Jiang, Boyi Shangguan, Li Huang, and Guangchao Liu. A deeply supervised image fusion network for change detection in high resolution bi-temporal remote sensing images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 166:183–200, 2020. 4
- [30] Zhuo Zheng, Ailong Ma, Liangpei Zhang, and Yanfei Zhong. Change is Everywhere: Single-Temporal Supervised Object Change Detection in Remote Sensing Imagery, 2023. 3, 4