

MULTI-TEMPORAL HYPERSPECTRAL ANALYSIS FOR EARLY FOREST DEGRADATION DETECTION USING DEEP LEARNING

**Thesis Submitted in Partial Fulfillment of the Requirements for the
Degree of**

**MASTER OF TECHNOLOGY
in
ARTIFICIAL INTELLIGENCE**

**by
RISHIKA GUPTA
(Roll No.: 24/AFI/27)**

Under the Supervision of

DR. KAVINDER SINGH
Assistant professor, Department of Computer Science and Engineering
Delhi Technological University



**Department Of Computer Science and Engineering
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi-110042, India**

May, 2026

DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi-110042, India

ACKNOWLEDGEMENT


I express my sincere gratitude to my supervisor, **Dr. Kavinder Singh**, for his guidance, constructive suggestions, and academic support throughout the development of this thesis. His feedback helped me refine the research problem, implementation design, experimental protocol, and final interpretation of results.

I am thankful to the Department of Computer Science and Engineering, Delhi Technological University, for providing an academic environment suitable for completing this work. I also acknowledge the availability of open hyperspectral benchmark datasets and NEON airborne observation products that made the experimental part of this thesis possible.

I am grateful to my family, friends, and classmates for their encouragement during the implementation, result checking, and report preparation stages of this work.

Place: Delhi

Date: May, 2026



Rishika Gupta


DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daultpur, Main Bawana Road, Delhi-110042, India

CANDIDATE'S DECLARATION

I, **Rishika Gupta** (Roll No. 24/AFI/27), hereby declare that the thesis report titled "**Multi-Temporal Hyperspectral Analysis for Early Forest Degradation Detection Using Deep Learning**" submitted to Delhi Technological University, Delhi, in partial fulfillment of the requirements for the award of the degree of Master of Technology in Artificial Intelligence, is my original work. The work reported in this thesis has been carried out under the supervision of Dr. Kavinder Singh, Department of Computer Science and Engineering, Delhi Technological University. The thesis has not been submitted, either in part or in full, to any other institution for the award of any degree, diploma, fellowship, or academic title.

Place: Delhi

Date: 29-5-26



Candidate's Signature

This is to certify that the student has incorporated all the corrections suggested by the examiners in the thesis and the statement made by the candidate is correct to the best of our knowledge.

Signature of Supervisor

Signature of External Examiner

DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daultpur, Main Bawana Road, Delhi-110042, India

CERTIFICATE

This is to certify that the thesis report titled "**Multi-Temporal Hyperspectral Analysis for Early Forest Degradation Detection Using Deep Learning**" submitted by **Rishika Gupta**, Roll No. 24/AFI/27, Department of Computer Science and Engineering, Delhi Technological University, Delhi, in partial fulfillment of the requirements for the award of the degree of Master of Technology in Artificial Intelligence, is a record of research work carried out by the student under my supervision.

To the best of my knowledge and belief, the work presented in this thesis is original and has not been submitted in part or in full for any degree, diploma, or academic award to this University or elsewhere.

Place: Delhi

Date: 29-5-26

Supervisor: _____

Dr. Kavinder Singh

Assistant Professor, Department of CSE

ABSTRACT

Forest degradation can begin before complete canopy removal. A forest area may still appear visually green while vegetation moisture, canopy structure, pigment response, or disturbance condition has already changed. Hyperspectral imagery is suitable for this problem because it records hundreds of narrow spectral bands and can therefore preserve subtle vegetation information that may not be visible in RGB or broad-band multispectral imagery.

This thesis presents ST-HyperDegradeNet, a spectral-temporal hyperspectral deep learning framework for early forest degradation detection and hyperspectral change analysis. The proposed framework operates on two hyperspectral images acquired over the same region at two different times. It performs invalid value handling, reflectance normalization, robust clipping, PCA-based dimensionality reduction, patch extraction, class-imbalance-aware training, full-scene inference, and vegetation-constrained refinement for the forest case study.

The main architectural idea is to use a true bi-temporal three-dimensional input representation. Patches from Time 1 and Time 2 are stacked along an explicit temporal dimension so that 3D convolution can learn spectral, spatial, and temporal interactions jointly. Squeeze-and-Excitation attention is used to recalibrate informative channels, while focal loss and class weighting support changed-class learning under imbalance.

Two labelled hyperspectral change-detection benchmarks, Bay Area AVIRIS and Hermiston Hyperion, are used for quantitative validation. NEON SOAP Forest is used as a weakly supervised forest-domain case study because manual pixel-level degradation labels are not available. In the final test-split results, the proposed model achieved 98.46% accuracy and 96.53% changed-class F1 on Bay Area, and 98.98% accuracy and 97.97% changed-class F1 on Hermiston. For NEON, the proposed model achieved 87.10% pseudo-reference agreement, and the vegetation-constrained final map identified 1.62% early degradation and 1.91% strong degradation within the selected crop.

The results show that ST-HyperDegradeNet is a competitive and interpretable framework for benchmark hyperspectral change detection and a practical candidate-screening approach for forest degradation mapping under limited-label conditions. The

thesis emphasizes cautious interpretation of pseudo-reference results and identifies future work involving field labels, multi-site validation, LiDAR integration, and longer multi-temporal time series.

TABLE OF CONTENTS

ACKNOWLEDGEMENT	ii
CANDIDATE’S DECLARATION	iii
CERTIFICATE	iv
ABSTRACT	v
LIST OF TABLES	ix
LIST OF FIGURES	x
LIST OF ABBREVIATIONS	xi
CHAPTER 1 INTRODUCTION	1
1.1 Background and Motivation.....	1
1.2 Problem Statement	2
1.3 Objectives and Scope	3
1.4 Contributions.....	3
1.5 Organization of the Thesis	4
CHAPTER 2 LITERATURE REVIEW	5
2.1 Hyperspectral Remote Sensing and Vegetation Monitoring.....	5
2.2 Evolution of Hyperspectral Classification and Change Detection Methods.....	6
2.3 Deep Learning for Hyperspectral Change Detection	8
2.4 Attention, Class Imbalance, and Weak Supervision	9
2.5 Research Gap and Comparative study	10
CHAPTER 3 PROPOSED WORK METHODOLOGY	13
3.1 Overview of the Proposed Framework	13
3.2 Preprocessing and Patch Formation	15
3.3 ST-HyperDegradeNet Architecture	16

3.4 Loss Function and Vegetation-constrained Refinement	18
3.5 Baseline Methods and Ablation Variants.....	19
CHAPTER 4 RESULTS AND ANALYSIS.....	21
4.1 Evaluation Metrics	21
4.2 Experimental Setup and Evaluation Protocol	21
4.3 Overall Performance Summary.....	23
4.4 Benchmark Results: Bay Area and Hermiston.....	24
4.5 NEON SOAP Forest Case Study	27
4.6 Ablation and Discussion	29
CHAPTER 5 CONCLUSION AND FUTURE SCOPE	31
5.1 Conclusion	31
5.2 Future Scope	32
5.3 Social Impact.....	33

LIST OF TABLES

Table 2.1: Comparative study of method families relevant to hyperspectral change and degradation detection.	10
Table 3.1: Dataset summary used in the thesis.	15
Table 3.2: Preprocessing settings used in the final implementation.	16
Table 3.3: Main architecture stages of ST-HyperDegradeNet.	17
Table 3.4: Methods compared in the final implementation.	19
Table 3.5: Ablation variants used for component analysis.	20
Table 4.1: Training settings used in the final implementation.	22
Table 4.2: Proposed ST-HyperDegradeNet summary across datasets.	23
Table 4.3: Bay Area AVIRIS comparison on the held-out test split.	24
Table 4.4: Hermiston Hyperion comparison on the held-out test split.	26
Table 4.5: NEON SOAP Forest pseudo-reference comparison on the held-out test split.	28
Table 4.6: Final NEON vegetation-constrained degradation area summary.	29
Table 4.7: Ablation results across datasets.	29

LIST OF FIGURES

Figure 1.1: Forest degradation continuum addressed in the thesis.	2
Figure 2.1: Evolution of methods relevant to hyperspectral classification and change detection.	8
Figure 3.1: Representative bi-temporal false-colour scenes from the datasets used in the thesis.	13
Figure 3.2: Overall workflow of the proposed ST-HyperDegradeNet framework...	14
Figure 3.3: Architecture of the proposed true bi-temporal 3D CNN with SE attention.	17
Figure 4.1: Bay Area AVIRIS false-colour input, ground-truth map, proposed prediction, and HCDNet-Lite prediction.	25
Figure 4.2: Bay Area changed-class F1 comparison.....	25
Figure 4.3: Hermiston Hyperion false-colour input, ground-truth map, proposed prediction, and HCDNet-Lite prediction.	26
Figure 4.4: Hermiston accuracy comparison.	27
Figure 4.5: NEON SOAP Forest input, pseudo-reference map, proposed prediction, and final degradation map.....	28
Figure 4.6: NEON pseudo-reference changed-class F1 comparison.	29
Figure 4.7: Changed-class F1 comparison for ablation variants across datasets.	30

LIST OF ABBREVIATIONS

Abbreviation	Meaning
AVIRIS	Airborne Visible/Infrared Imaging Spectrometer
CNN	Convolutional Neural Network
CVA	Change Vector Analysis
F1	Harmonic mean of precision and recall
HSI	Hyperspectral Image
Kappa	Cohen kappa coefficient
NDVI	Normalized Difference Vegetation Index
NEON	National Ecological Observatory Network
OA	Overall Accuracy
PCA	Principal Component Analysis
RF	Random Forest
SAM	Spectral Angle Mapper
SE	Squeeze-and-Excitation
ST	Spectral-Temporal
T1/T2	First and second acquisition times

CHAPTER 1

INTRODUCTION

1.1 Background and Motivation

Remote sensing provides repeated observations of the Earth surface and is widely used for environmental monitoring. Hyperspectral remote sensing extends this capability by measuring reflected energy in many narrow and contiguous spectral bands. Instead of representing a location only through a few broad colour channels, a hyperspectral image stores a detailed spectral signature for every pixel. This property has made imaging spectroscopy useful for vegetation analysis, material identification, land-cover mapping, and ecological monitoring [\[1\]](#).

Forest degradation is a gradual process and is often more difficult to detect than complete deforestation. Deforestation usually produces a clear loss of canopy cover, while degradation may involve drought stress, disease, fire impact, selective extraction, pest disturbance, or slow canopy weakening. Such changes may affect pigment absorption, near-infrared canopy response, and short-wave infrared water-related bands before the land cover changes completely [\[2\]](#), [\[3\]](#). For this reason, hyperspectral data are suitable for early degradation screening.

The main challenge is that hyperspectral images are high dimensional and sensitive to atmospheric, phenological, illumination, and sensor-related variation. A simple difference image may capture change magnitude, but it may not distinguish real vegetation degradation from unrelated spectral variation. At the same time, deep learning models require careful design because changed or degraded pixels are often fewer than unchanged pixels. The motivation of this thesis is therefore to develop a compact spectral-spatial-temporal model that can learn from bi-temporal hyperspectral patches and can be evaluated under a transparent protocol.

Figure 1.1 illustrates the forest-condition continuum that motivates this thesis. It is placed in the background section because the thesis focuses on detecting early and strong degradation candidates before complete forest loss becomes visually obvious.




Figure 1.1:Forest degradation continuum addressed in the thesis.

1.2 Problem Statement

Given two hyperspectral images of the same region acquired at two different times, the problem is to identify pixels that represent change or degradation while preserving spectral, spatial, and temporal information. For labelled benchmark datasets, the task is binary change detection. For the NEON forest case study, the task is to produce a conservative degradation map using weak pseudo-reference labels and vegetation-constrained refinement because field-verified pixel-level degradation labels are not available.

This problem requires a model that can address four linked issues. First, hyperspectral data contain many correlated bands and therefore require dimensionality control. Second, degradation and changed pixels may be minority classes. Third, spatial context around a pixel is important because vegetation stress is rarely isolated to a single pixel. Fourth, forest degradation is an ecological condition, so pseudo-reference results must be interpreted more cautiously than benchmark ground-truth results.

1.3 Objectives and Scope

The objectives of this thesis are as follows:

- To design a lightweight bi-temporal hyperspectral deep learning framework for early forest degradation and change detection.
- To preprocess hyperspectral image pairs using invalid-value handling, reflectance scaling, clipping, PCA-based dimensionality reduction, and vegetation-index support where required.
- To compare the proposed method with traditional change detection, classical machine learning, and lightweight deep learning baselines under a common held-out test-split protocol.
- To evaluate the proposed framework on two labelled hyperspectral benchmark datasets and one weakly supervised NEON forest case study.
- To analyze the effect of the main architectural components through an ablation study.

The scope of the work is limited to the datasets and labels used in the final implementation. The NEON result is treated as a candidate degradation screening output, not as field-verified ecological accuracy.

1.4 Contributions

The main contributions of the thesis are summarized below.

- A true bi-temporal 3D input formulation that stacks Time 1 and Time 2 patches along an explicit temporal dimension.
- A compact ST-HyperDegradeNet architecture using 3D convolution, temporal aggregation, SE attention, 2D spatial refinement, and class-imbalance-aware training.
- A consistent evaluation design using Bay Area AVIRIS and Hermiston Hyperion as labelled benchmarks, and NEON SOAP Forest as a weakly supervised forest case study.
- A final vegetation-constrained degradation refinement for NEON that separates unchanged or healthy forest, early degradation, and strong degradation.

- A component-level ablation analysis comparing the full model with no-SE, no-focal, no-SE plus no-focal, and Siamese 2D variants.

1.5 Organization of the Thesis

The motivations, problem statement, research objectives, scope, and contribution have been introduced in Chapter 1. In Chapter 2, the literature review is carried out on Hyperspectral Image Classification, Hyperspectral Image Change Detection, Deep Learning Methods, Attention Mechanisms, Weak Supervision, and Forest Degradation Detection. The third chapter is dedicated to introducing the methodology used for this research, which includes the preprocessing stage, ST-HyperDegradeNet network structure, formulae involved, and the baseline algorithms used.

CHAPTER 2

LITERATURE REVIEW

2.1 Hyperspectral Remote Sensing and Vegetation Monitoring

Hyperspectral imaging techniques were designed to generate high-resolution spectral data useful in earth observations and material classification [1]. Visible, near-infrared and short-wave infrared bands are related to pigment absorption, internal leaf structure and canopy structure, respectively. Short-wave infrared band is related to water content and dry matter [2], [3]. Such relationships make the hyperspectral images valuable in detecting vegetation stress not obvious in regular images.

Vegetation index, e.g., NDVI, which represents forest greenness by means of red and near-infrared reflectances, is commonly used to describe the state of a forest canopy [4]. Still, one index alone does not characterize all spectral variations in forest canopies. Therefore, index-derived information is good for providing hints but should not be the basis of a final model.

Further, a hyperspectral pixel is not only a vector but a real signal subject to factors like material on the surface, illumination conditions, atmospheric correction, the canopy architecture, the presence of moisture and mixture of background. In case the studied scene contains trees, the signal can consist of a reflection produced by leaves, branches, bushes, shady zones of tree crowns, and soil. Therefore, the problem becomes even more complicated in comparison with identifying changes in a specific landscape. Here, one needs not only spectral information but also spatial and temporal data in addition to the latter.

Firstly, as the number of features increases significantly, there arises the widely known issue associated with hyperspectral images that an increase in the number of bands does not necessarily result in better classifier performance. The neighboring bands are strongly correlated. However, noise in bands and absence of sufficient training samples make classifier performance worse. This kind of problem was mentioned by other scientists due to the problem of small sample sizes in the hyperspectral literature.

For early forest degradation, the important signal may be subtle. A small decline in near-infrared response, a change in moisture-sensitive wavelengths or an increase in

exposed-soil contribution can indicate stress before full canopy removal. This is the reason why hyperspectral data are relevant to the thesis title. However, the same sensitivity can also capture illumination changes, seasonal differences and sensor effects. The literature therefore supports the need for careful preprocessing, cautious interpretation and a clear separation between benchmark accuracy and forest-case-study inference.

2.2 Evolution of Hyperspectral Classification and Change Detection Methods

Examples of techniques used in change detection include differencing, ratioing, change vector analysis, and spectral angle mapper. These techniques have several benefits, including being intuitive and easy to understand [5]. Change Vector Analysis measures the difference in the magnitude of spectra between two points in time [6], whereas Spectral Angle Mapper measures spectral differences using angular separation and is less sensitive to variations in illumination levels [7]. Otsu thresholding can be employed to create maps based on change measurements [8].

Classical machine learning improved upon threshold rules by learning decision boundaries from labelled samples. Random Forest is a strong ensemble method that handles nonlinear relationships and noisy features [9]. In remote sensing, machine-learning classifiers have been used widely for land-cover mapping and accuracy assessment [10]. However, pixel-level classifiers do not automatically model local spatial context, which is important for hyperspectral images.

Deep learning introduced automatic feature learning for hyperspectral image analysis. CNN-based methods learn spatial and spectral patterns directly from data [11], and 3D CNN models can process local spectral-spatial cubes more naturally than purely 2D models [12]. Recurrent and sequence-based approaches further treat spectral bands as ordered signals [13]. Reviews of hyperspectral classification show a clear evolution from spectral-only methods to spectral-spatial deep models, attention models, and transfer or limited-label methods [14], [15]. Figure 2.1 depicts this evolution and shows the position of the present work in the spectral-spatial-temporal deep learning stage.

Conventional spectral classifiers are still important because they establish a baseline for interpretability. Methods such as spectral angle comparison, change vector analysis and vegetation-index thresholding can be implemented without a large training set and their decision logic is easy to explain. Their weakness is that they generally operate at pixel level and depend heavily on thresholds. When the same spectral change score is used across different scenes, it may respond to true change, sensor variation, phenology or atmospheric differences. This limitation becomes more serious in degradation detection because degradation is not always a strong or abrupt spectral event.

Classical machine learning techniques such as SVMs and Random Forest performed better by learning the boundary lines between classes. They are able to detect nonlinear relationships among bands and tend to be robust in case the training data are not large. Nonetheless, they typically demand the inputs to be engineered, and cannot model spatial layout of pixels automatically. As the adjacent pixels within hyperspectral data may contain valuable information such as field boundaries, canopy texture and small degradation patches, researchers tended to switch from spectral-vector classifier to spectral-spatial classifiers [27].

Deep learning algorithms have decreased reliance on manual feature extraction. While the two-dimensional CNN is capable of extracting spatial information, the three-dimensional CNN can learn filters in both spectral and spatial domains simultaneously. 3D convolutions prove to be an effective technique for analysing hyperspectral data, since adjacent wavelengths are correlated in nature. On the other hand, deep models tend to overfit in the presence of few labelled examples. Thus, the existing literature provides guidance in finding an equilibrium between expressiveness and efficiency, which is true for this thesis as well.

Moreover, current research works have also seen a move towards the application of transfer learning and self-supervised pre-training and generalization across different datasets. This is vital due to the costly nature of generating dense annotations for each individual hyperspectral scene. Although the thesis does not make any effort to train a huge foundation model, the same concept will be applied in the sense that a viable hyperspectral model should not rely solely on abundant labeling.

Evolution of hyperspectral classification and change detection methods




Figure 2.1: Evolution of methods relevant to hyperspectral classification and change detection.

2.3 Deep Learning for Hyperspectral Change Detection

Hyperspectral change detection is different from single date classification in that the network needs to analyze two date observations at the same spatial location. Unsupervised Deep Slow Feature Analysis is able to suppress unchanged parts and amplify the changes in multi-temporal remote sensing images [16]. Binary Change Guided hyperspectral multiclass change detection proves that binary change can help guide multiclass changes [17]. From the above research, we see that temporal reasoning is necessary when dealing with two-date hyperspectral images.

Recently, many deep networks employ siamese, attention, transformer-like, and double stream CNN for change detection [18], [19]. Although those approaches have been shown to work well, their full versions may need lots of computational resource and specific training environments. In this thesis, in order to provide light weight baselines for further improvement, HCDNet-Lite and STT-Lite are provided in the same protocol, while not claiming that they are official full versions of those models.

The temporal interaction plays an important part in the difference between hyperspectral classification and hyperspectral change detection. In single-date classification, the prediction is made based on the observation from one date. In contrast, in change detection, the classifier needs to figure out whether there is any transition from one date to another date. In some cases, subtraction or independent

feature extraction might lose the relationship between dates or fail to take into account nonlinear spectral relationships. Therefore, the real challenge of using bi-temporal representation in change detection lies in the joint processing of these two dates together.

Unsupervised solutions are promising due to the lack of change ground truth. With deep slow feature analysis, binary change guidance or pseudo label method, effective change representation can still be obtained even without too much labeled training samples. Nonetheless, one should be careful when interpreting the unsupervised solution or pseudo-supervised solution since their predicted label is influenced by the assumption used to generate the targets [28]. This issue is especially prominent in the case study of NEON forests.

The literature also shows that visual change maps and numerical scores should be examined together. A method may obtain a high numerical score if the label generation procedure favours its assumptions, but the resulting map may still be spatially noisy or ecologically difficult to interpret. For this reason, the present thesis reports metrics, prediction maps and area summaries together rather than relying on a single accuracy number.

2.4 Attention, Class Imbalance, and Weak Supervision

Attention mechanisms are useful because not all learned channels are equally informative. Squeeze-and-Excitation networks recalibrate channels by learning channel-wise importance weights [20]. In hyperspectral analysis, this idea is relevant because some spectral-temporal responses may contribute more strongly to degradation or change detection than others. ST-HyperDegradeNet therefore includes SE attention after deep feature extraction.

Class imbalance is another important issue. In change detection, unchanged pixels often dominate the image, and a model can obtain high overall accuracy while missing minority changed pixels. Focal loss was introduced to focus learning on difficult samples and reduce the dominance of easy examples [21]. In this thesis, focal loss and automatic class weighting are used in the full model, and their effect is tested through ablation.

Weak supervision is particularly relevant for the NEON forest case study. Manual pixel-level degradation labels are not available for the selected scene, so pseudo-reference labels and vegetation constraints are used. Current trends in self-supervised, cross-dataset, and foundation-model style remote-sensing learning also indicate that reducing dependence on dense manual labels is an important research direction [22], [23].

2.5 Research Gap and Comparative study

The literature reviewed above shows a gradual movement from simple spectral comparison methods to machine learning, deep spectral-spatial learning, attention-based models, and label-efficient approaches. Table 2.1 gives a compact comparative view of the main method families that are relevant to this thesis. The table is included to clarify how the proposed work is positioned rather than to claim that one family is universally better than all others.

Table 2.1: Comparative study of method families relevant to hyperspectral change and degradation detection.

Method family	Representative approaches	Strengths	Limitations	Relevance to this thesis
Index and threshold methods	NDVI differencing, image differencing, Otsu thresholding	Simple, interpretable, low computation	Sensitive to thresholds, illumination, phenology, and non-degradation spectral variation	Used as baseline logic and NEON vegetation constraint
Spectral change measures	CVA, SAM	Directly compare two spectral observations; useful for unsupervised screening	Usually pixel-wise and weak in spatial-context modelling	Included as traditional baselines for benchmark comparison
Classical machine learning	Random Forest, SVM-style pixel classifiers	Can learn nonlinear class boundaries from labelled samples	Depend on engineered features and usually do not learn local spatial structure directly	Random Forest included as classical supervised baseline
Spectral-spatial CNN models	2D CNN, 3D CNN, patch-based HSI classifiers	Learn local spatial and spectral features from data	May overfit when labels are limited; single-date models do	Forms the foundation for the proposed

Method family	Representative approaches	Strengths	Limitations	Relevance to this thesis
Bi-temporal deep change models	Siamese CNN, DSFA, BCG-style change detection	Designed for two-date comparison and richer change representation	not fully solve temporal comparison Often require careful training, matched data, or may be heavier than needed for thesis-scale deployment	patch-based deep model Motivates direct T1-T2 learning and comparison with Siamese 2D variant
Attention and transformer models	SE attention, CNN-attention, spectral-temporal token attention	Improve feature weighting and long-range interaction modelling	Performance gain can be dataset-dependent and computation can increase	SE attention and lightweight STT-Lite are used in the final study
Limited-label and weak-supervision approaches	Self-supervised learning, cross-dataset training, pseudo-reference learning	Reduce dependence on dense manual labels	Need careful interpretation because pseudo-labels are not field truth	Supports the NEON case-study framing as pseudo-reference agreement

The first research gap concerns the need for joint spectral, spatial, and temporal learning. Traditional change scores such as CVA and SAM compare two observations directly, but they do not learn local spatial context and cannot easily separate ecological degradation from illumination or phenological variation. Single-date hyperspectral classifiers learn useful spectral-spatial features, but they are not designed to preserve the two acquisition dates as a true temporal input. This creates the need for a compact bi-temporal model that can learn T1 and T2 jointly rather than relying only on a precomputed difference image.

The second limitation pertains to insufficient labels and unbalanced class. The pixels that have undergone changes or have been degraded are typically fewer than those that have remained unchanged, making the mere accuracy score misleading for evaluation purposes. It could happen that.

Another research limitation pertains to the connection between performance on benchmarks and applicability to forests. The majority of hyperspectral analysis studies demonstrate great results when working with benchmarks; however, the task may not capture the diversity associated with the forest ecosystem and its

degradation processes. Forest degradation tends to occur gradually, spatially heterogeneous, and affected by factors like moisture, topography, density of canopy cover, etc. For that reason, if an approach performed successfully in a benchmark environment, a corresponding case study should be conducted before declaring its relevance for environmental monitoring purposes.

CHAPTER 3

PROPOSED WORK METHODOLOGY

3.1 Overview of the Proposed Framework

The proposed framework is named ST-HyperDegradeNet, where ST denotes spectral-temporal learning. The framework uses two co-registered hyperspectral images of the same scene acquired at different times. Instead of subtracting the images at the beginning, the method preserves both temporal states and allows the model to learn before-and-after relationships from local patches.

Figure 3.1 illustrates some representative examples of bi-temporal false-color images from the data sets that have been utilized for this research. Figure 3.1 has been put forward in the methodology section since these particular images provide an idea as to why Bay Area and Hermiston are chosen as labeled benchmark regions whereas NEON SOAP Forest as weakly supervised forest example.




Figure 3.1: Representative bi-temporal false-colour scenes from the datasets used in the thesis.

Figure 3.2 depicts the complete workflow used in the thesis. The process starts from bi-temporal hyperspectral input, applies preprocessing and PCA, extracts local patches, trains the proposed model and baseline methods under a common split, and produces quantitative metrics and full-scene prediction maps. For NEON, an additional vegetation-constrained refinement is applied to obtain the final degradation map.




Figure 3.2: Overall workflow of the proposed ST-HyperDegradeNet framework.

Table 3.1 gives the dataset summary used in the methodology. The table is placed here because the label source directly affects the interpretation of the training and evaluation protocol.

The Bay Area AVIRIS benchmark is chosen as the first labelled dataset due to the presence of a well-defined change detection scenario in this dataset, where changes in the area are detected based on a binary ground truth image. In the context of the thesis, its purpose is to evaluate if the proposed model is able to learn relevant patterns of change from bi-temporal data, where there exist well-defined training labels. The image exhibits spatial and spectral variations, thus making it possible to check the capabilities of the model for separating changed and unchanged parts.

The Hermiston Hyperion dataset is used as the second labelled benchmark. It is important because it differs from Bay Area in sensor characteristics, scene appearance and spatial pattern. Using a second benchmark reduces the risk that the model is only adapted to one scene. In this thesis, Hermiston provides additional evidence that the proposed true bi-temporal 3D design can generalize across a different hyperspectral change-detection dataset. The high performance obtained on both Bay Area and Hermiston is therefore stronger evidence than performance on only one dataset.

The NEON SOAP Forest data are used for the forest-domain case study. This dataset is closest to the thesis title because it represents an actual forest monitoring context. However, it does not provide manual pixel-level labels for early and strong degradation in the selected crop. The NEON part of the thesis is therefore not treated as a conventional ground-truth benchmark. Instead, the data are used to demonstrate how the proposed pipeline can be adapted to forest degradation screening when only pseudo-reference labels and vegetation constraints are available [31].

The three-dataset design is intentional. Bay Area and Hermiston answer the technical question of whether the model can detect hyperspectral change under labelled conditions. NEON answers the application question of whether the same pipeline can

produce a cautious candidate degradation map in a real forest-domain setting. These two forms of evidence are complementary, but they are not the same. This distinction prevents the thesis from overclaiming forest accuracy while still demonstrating relevance to early degradation detection.

Data choice influences how class structure is treated too. Binary classifications changed/unchanged are used for both Bay Area and Hermiston data whereas NEON classification uses a three-class approach that treats the dataset as a pseudo reference with healthy/unchanged, early degradation and strong degradation. The difference is captured in how the classification outputs are treated, metrics used and conclusions drawn in the end. Direct comparison of results from Bay Area, Hermiston and NEON is thus avoided due to difference in labeling sources and ecological significance.

Table 3.1: Dataset summary used in the thesis.

Dataset	Sensor/product	Dates	Size used	Bands	Role
Bay Area AVIRIS	AVIRIS	2013/2015	600 x 500	224	Ground-truth benchmark
Hermiston Hyperion	EO-1 Hyperion	2004/2007	180 x 225	Selected bands	Ground-truth benchmark
NEON SOAP Forest	NEON AOP DP3.30006.002	2018- 06/2024-06	128 x 128 crop	426	Weakly supervised forest case study

3.2 Preprocessing and Patch Formation

Preprocessing is done by removing bad data, normalizing reflectance data, truncating data, and reducing dimensionality through PCA. PCA is performed as a means of compression only and not for classification. PCA helps to reduce redundancy between neighboring hyperspectral bands.

$$X' = W^T(X - \mu) \quad (3.1)$$

In Equation (3.1), X represents the original spectral vector, μ is the mean vector, W contains the selected principal components, and X' is the reduced representation. After PCA, each pixel is represented by fewer informative components while retaining most of the useful spectral variance.

For each valid pixel location, a local patch of size $P \times P$ is extracted from both acquisition dates. In the final implementation, P is set to 7. The two patches are stacked into a true bi-temporal tensor as shown in Equation (3.2).

$$Z_{ij} \in \mathbb{R}^{(B' \times 2 \times P \times P)} \quad (3.2)$$

Here, B' is the reduced spectral dimension, the dimension of length 2 represents Time 1 and Time 2, and $P \times P$ represents the local spatial neighbourhood. This representation is central to the thesis because it lets the network learn temporal change directly rather than forcing a hand-crafted difference at the input stage.

The final preprocessing choices are summarized in Table 3.2. The NEON configuration includes vegetation-index support because its final interpretation is forest degradation rather than generic binary change.

Table 3.2: Preprocessing settings used in the final implementation.

Dataset	PCA components	Vegetation indices	Clipping	Label setting
Bay Area AVIRIS	24	No	1-99 percentile	Binary benchmark
Hermiston Hyperion	30	No	1-99 percentile	Binary benchmark
NEON SOAP Forest	24	Yes	2-98 percentile	Three-class pseudo-reference

3.3 ST-HyperDegradeNet Architecture

This approach implements three-dimensional convolutions to jointly learn features in spectral space, time, aggregation of time, squeeze-and-excitation (SE) block, two-dimensional spatial refinement operations, and the classification head. This 3D convolution layer takes the permuted data tensor as an input and performs filter learning through the spectral space dimension, both times, and local spatial neighbourhood.




Figure 3.3: Architecture of the proposed true bi-temporal 3D CNN with SE attention.

The main architecture stages are summarized in Table 3.3. This table should be read together with Figure 3.2: the figure shows the flow of tensors, while the table explains the function of each block. . For a batch of patches, the tensor is represented as $[N, T, C, P, P]$, where N is batch size, T is temporal depth, C is the number of PCA or feature channels, and P is the patch size. Before 3D convolution, the tensor is permuted to $[N, C, T, P, P]$. This allows the convolution kernels to span the two acquisition dates and the local spatial neighbourhood jointly.

Table 3.3: Main architecture stages of ST-HyperDegradeNet.

Stage	Representation	Function
Input patch	$[N, 2, C, 7, 7]$	Bi-temporal patch with PCA/index channels and local neighbourhood.
Permutation	$[N, C, 2, 7, 7]$	Reorders the tensor for 3D convolution.
3D Conv block 1	$C \rightarrow 64$	Learns initial temporal and spatial change features.
3D Conv block 2	$64 \rightarrow 128$	Builds deeper spectral-spatial-temporal features.
Temporal pooling + SE	128 channels	Aggregates temporal responses and recalibrates channels.
2D Conv block	$128 \rightarrow 256$	Refines spatial features after temporal aggregation.
Classifier	$256 \rightarrow \text{classes}$	Uses pooling, dropout, and fully connected layers for prediction.

$$F = \text{Conv3D}(Z; \theta) \quad (3.3)$$

Equation (3.3) represents the 3D convolutional feature extraction step, where Z is the bi-temporal input tensor and θ denotes the learnable convolutional parameters.

$$s = \sigma(W_2 \delta(W_1 z)), \quad \tilde{U}_c = s c U_c \quad (3.4)$$

Equation (3.4) represents the SE recalibration operation. The channel descriptor z is passed through two learnable transformations and nonlinearities to obtain channel weights s , which are used to rescale the feature channel U_c .

The architecture is intentionally compact. It first learns temporal and spatial interactions using 3D convolution, then reduces temporal depth through aggregation, and finally refines the learned spatial representation using 2D convolution. This design avoids a heavy transformer or very deep network while still allowing the two dates to interact inside the feature extractor. The compactness is important for thesis reproducibility and for possible future use on larger remote-sensing scenes.

The use of a true temporal dimension distinguishes the proposed architecture from a simple channel concatenation or late-fusion Siamese design. In late fusion, each date may be processed mostly independently before comparison. In ST-HyperDegradeNet, the temporal kernel in the first 3D convolution directly spans the two dates. This allows the model to learn patterns such as spectral reduction, spectral increase, or local structural change as joint temporal features.

SE attention is applied after feature extraction rather than before it. This is appropriate because the network first forms higher-level spectral-spatial-temporal features and then learns which of those channels are more useful for classification. In hyperspectral degradation detection, the informative response may not be located in one fixed band or index; it may appear as a learned combination of spectral and spatial cues. Channel recalibration therefore provides a lightweight mechanism for emphasizing useful learned responses.

3.4 Loss Function and Vegetation-constrained Refinement

Changed or degraded pixels usually form a smaller part of the scene than unchanged pixels. Therefore, overall accuracy alone is not sufficient. The proposed full model uses focal loss with class weighting to reduce the dominance of easy unchanged examples during training.

$$L = -\alpha(1 - p)^\gamma \log(p) \quad (3.5)$$

Equation (3.5) shows focal loss, where p is the model probability for the target class, α is the class weight, and γ controls the emphasis on difficult examples. In the final implementation, focal loss is tested through ablation rather than assumed to be universally beneficial.

For the NEON case study, the final map is constrained using vegetation-decline evidence. The goal is to avoid assigning degradation when the predicted class is not supported by an NDVI-drop cue. The final refinement rule is shown in Equation (3.6).

$$\hat{y}_{\text{final}}(i,j) = \{2 \text{ if strong, } 1 \text{ if early, } 0 \text{ otherwise}\} \quad (3.6)$$

In Equation (3.6), $\hat{y}(i,j)$ is the model prediction, $\Delta\text{NDVI}(i,j)$ is the vegetation decline indicator, τ_e is the early degradation threshold, and τ_s is the strong degradation threshold. This rule makes the NEON output more conservative and appropriate for candidate screening.

3.5 Baseline Methods and Ablation Variants

Table 3.4 lists the comparison methods used in the final implementation. The methods are selected to cover traditional, classical machine-learning, modern lightweight deep-learning, and proposed-model categories.

Table 3.4: Methods compared in the final implementation.

Method	Category	Role in thesis
CVA binary	Traditional threshold method	Magnitude of spectral change.
SAM binary	Spectral-angle method	Angular difference between spectral vectors.
CVA 3-class	Threshold multiclass method	Separates unchanged, moderate, and strong change.
NDVI-drop 3-class	Vegetation rule	Vegetation-decline thresholds.
Random Forest	Classical ML	Pixel-level supervised baseline.
HCDNet-Lite	Deep baseline	Representative CNN-attention baseline.
STT-Lite	Deep baseline	Representative spectral-temporal attention baseline.
ST-HyperDegradeNet	Proposed method	True bi-temporal 3D CNN with SE attention.

The ablation variants are summarized in Table 3.5. These variants are used to separate the contribution of the true bi-temporal 3D design from the effects of attention and loss design.

Table 3.5: Ablation variants used for component analysis.

Variant	Description
Full reference	True bi-temporal 3D input with SE attention and focal loss.
No SE	Removes channel recalibration.
No focal	Uses weighted cross-entropy instead of focal loss.
No SE + no focal	Removes both SE attention and focal loss.
Siamese 2D	Processes each date separately.

CHAPTER 4

RESULTS AND ANALYSIS

4.1 Evaluation Metrics

Accuracy, macro-F1, weighted-F1, changed-class F1, and Cohen's kappa are used.

Accuracy is

$$OA = \frac{TP + TN}{TP + TN + FP + FN}, \quad (4.1)$$

Precision, recall, and F1-score are

$$Precision = \frac{TP}{TP + FP}, Recall = \frac{TP}{TP + FN}, \quad (4.2)$$

$$F1 = \frac{2 \cdot Precision \cdot Recall}{Precision + Recall}. \quad (4.3)$$

Kappa is

$$\kappa = \frac{P_o - P_e}{1 - p_e}. \quad (4.4)$$

Changed-class F1 is emphasized because changed or degraded pixels may be minority samples. A method may obtain high accuracy by predicting the majority unchanged class, but such a result is not sufficient for change detection. Macro-F1 and changed-class F1 provide more balanced evidence.

4.2 Experimental Setup and Evaluation Protocol

The implementation is organized as a Python project with dataset-specific configuration files. Each configuration controls preprocessing, patch size, model settings, training limits, evaluation, prediction-map generation, and figure output. The final experiments use a common held-out protocol so that numerical comparisons are not affected by evaluating methods on different pixel locations.

Table 4.1 reports the final training settings used for the three datasets. The number of classes differs because Bay Area and Hermiston are binary benchmark tasks, while NEON is a three-class pseudo-reference forest case study.

Table 4.1: Training settings used in the final implementation.

Dataset	Classes	Patch	Batch	Epochs	Learning rate	Weight decay
Bay Area AVIRIS	2	7	128	25	0.001	0.0001
Hermiston Hyperion	2	7	128	60	0.001	0.0001
NEON SOAP Forest	3	7	96	25	0.001	0.0001

All learning algorithms and the deep learning algorithms are trained on the training split chosen by the validation split and then tested on the unseen test split. Traditional algorithms like CVA and SAM produce the complete map but their scores are calculated based on the same set of test-pixel locations which the learning algorithms use to calculate their results. Scores like Accuracy, macro-F1 score, Changed-class F1 score, and Cohen Kappa are reported because Accuracy alone cannot detect the presence of minority class errors [24].

Accuracy is the most straightforward way of calculating performance because it directly calculates the percentage of correctly labeled test pixels. Although Accuracy may be helpful for a broad view of algorithmic performance, it is not enough for Change and Degradation Detection because, in case of imbalance, the classifier can get high accuracy by simply guessing the majority class.

Precision and recall offer a more fine-grained analysis of class behavior. Precision refers to the number of correctly identified pixels that have changed or degraded based on the available reference, whereas recall refers to the number of pixels that have changed or degraded according to the reference and have been properly detected. Recall is an important factor in the forest screening task since missing the degraded pixels could be problematic, but precision is also important since it is undesirable to have too many false positives, thereby making the map useless.

Macro-F1 is reported because it gives equal importance to all classes. This is useful when the minority changed class is important even though it occupies fewer pixels. Weighted-F1 is also useful in some contexts because it considers class support, but it

can still be influenced by the majority class. For this reason, the thesis gives special attention to changed-class F1, which directly indicates how well the changed or degraded class is recognized.

Cohen kappa is included because it measures agreement after accounting for chance agreement. In remote-sensing accuracy assessment, kappa has often been used as an additional indicator beyond overall accuracy [33]. The thesis does not treat kappa as perfect or sufficient by itself, but it provides a helpful check when the class distribution is imbalanced. A model with high accuracy but weak kappa may be benefiting from majority-class dominance.

For Bay Area and Hermiston, the metrics are interpreted against labelled benchmark references. For NEON, the same formulas are used, but the interpretation is different because the labels are pseudo-reference labels. Therefore, NEON accuracy, F1 and kappa are described as pseudo-reference agreement values. This distinction is important throughout the results chapter and prevents the NEON case study from being overstated as field-verified ecological accuracy.

4.3 Overall Performance Summary

Table 4.2 summarizes the final proposed-model performance across the three datasets. The table separates ground-truth benchmark evidence from pseudo-reference forest case-study evidence, which is essential for correct interpretation.

Table 4.2: Proposed ST-HyperDegradeNet summary across datasets.

Dataset	Evaluation type	Label type	Accuracy (%)	Macro-F1 (%)	Changed-F1 (%)	Kappa
Bay Area AVIRIS	Ground-truth benchmark	Binary	98.46	97.77	96.53	0.9553
Hermiston Hyperion	Ground-truth benchmark	Binary	98.98	98.64	97.97	0.9729
NEON SOAP Forest	Pseudo- reference case study	3-class	87.10	44.87	38.51	0.2972

The proposed model performs strongly on the two labelled benchmarks and provides a usable weakly supervised output for NEON. The lower NEON macro-F1 and kappa should not be interpreted as a failure of the model alone; they reflect the difficulty of

pseudo-reference learning and the absence of manually verified forest degradation labels.

4.4 Benchmark Results: Bay Area and Hermiston

Table 4.3 presents the Bay Area AVIRIS comparison. The proposed model achieves 98.46% accuracy and 96.53% changed-class F1, showing a clear improvement over traditional methods and Random Forest. HCDNet-Lite is also competitive, so the result is best described as a strong benchmark performance rather than an absolute universal dominance claim.

Table 4.3: Bay Area AVIRIS comparison on the held-out test split.

Method	Label type	Accuracy (%)	Macro-F1 (%)	Changed-F1 (%)	Kappa
CVA binary	Binary	67.58	55.65	32.65	0.1159
SAM binary	Binary	64.22	54.78	34.12	0.1093
CVA 3-class	3-class	69.12	35.37	30.27	0.0938
NDVI-drop 3-class	3-class	63.10	31.52	28.26	0.0136
Random Forest	Binary	94.83	92.45	88.20	0.8490
Proposed ST-HyperDegradeNet	Binary	98.46	97.77	96.53	0.9553
HCDNet-Lite	Binary	98.18	97.39	95.95	0.9478
STT-Lite	Binary	93.46	91.00	86.30	0.8205

Figure 4.1 shows the Bay Area input and prediction maps, while Figure 4.2 plots the changed-class F1 comparison. The spatial maps confirm that the learned deep models produce more coherent predictions than threshold-based change scores.

Bay Area AVIRIS visual results




Figure 4.1: Bay Area AVIRIS false-colour input, ground-truth map, proposed prediction, and HCDNet-Lite prediction.




Figure 4.2: Bay Area changed-class F1 comparison.

Table 4.4 reports the Hermiston Hyperion comparison. The proposed model achieves 98.98% accuracy and 97.97% changed-class F1. This second benchmark supports the stability of the proposed approach across a different sensor and scene structure.

Table 4.4: Hermiston Hyperion comparison on the held-out test split.

Method	Label type	Accuracy (%)	Macro-F1 (%)	Changed-F1 (%)	Kappa
CVA binary	Binary	68.64	65.61	55.40	0.3420
SAM binary	Binary	64.38	61.71	51.59	0.2774
CVA 3-class	3-class	73.84	41.30	50.56	0.2826
NDVI-drop 3-class	3-class	72.00	40.13	57.32	0.2819
Random Forest	Binary	97.88	97.12	95.64	0.9424
Proposed ST-HyperDegradeNet	Binary	98.98	98.64	97.97	0.9729
HCDNet-Lite	Binary	98.86	98.49	97.74	0.9698
STT-Lite	Binary	97.73	97.00	95.51	0.9399

Figure 4.3 shows the Hermiston input and prediction maps, and Figure 4.4 summarizes the accuracy comparison. The results again show that patch-based deep learning is more effective than pure spectral-angle or change-magnitude thresholding.

Hermiston Hyperion visual results




Figure 4.3: Hermiston Hyperion false-colour input, ground-truth map, proposed prediction, and HCDNet-Lite prediction.




Figure 4.4: Hermiston accuracy comparison.

4.5 NEON SOAP Forest Case Study

NEON SOAP Forest is included to connect the thesis to the title objective of early forest degradation detection. Unlike the benchmark datasets, NEON does not provide manual pixel-level degradation labels for the selected crop. Therefore, its reported numerical values are pseudo-reference agreement values, and the final map is interpreted as a candidate screening output rather than field-verified degradation accuracy.

Table 4.5 presents the NEON pseudo-reference comparison. CVA 3-class obtains high pseudo-reference agreement because the pseudo-reference itself is partly derived from threshold-style change magnitude and vegetation logic. For this reason, NEON results must be interpreted with caution.

Table 4.5: NEON SOAP Forest pseudo-reference comparison on the held-out test split.

Method	Label type	Accuracy (%)	Macro-F1 (%)	Changed-F1 (%)	Kappa
CVA binary	Binary	71.85	52.31	21.78	0.1590
SAM binary	Binary	77.82	56.53	26.10	0.2070
CVA 3-class	3-class	98.01	69.78	89.04	0.7660
NDVI-drop 3-class	3-class	69.20	30.01	4.89	- 0.0073
Random Forest	3-class	97.41	56.58	52.27	0.5023
Proposed ST-HyperDegradeNet	3-class	87.10	44.87	38.51	0.2972
HCDNet-Lite	3-class	85.59	45.71	36.68	0.2753
STT-Lite	3-class	90.54	47.49	46.15	0.3040

Figure 4.5 shows the NEON false-colour views, pseudo-reference map, proposed prediction, and final vegetation-constrained degradation map. Figure 4.6 summarizes the changed-class F1 comparison for NEON.





Figure 4.5: NEON SOAP Forest input, pseudo-reference map, proposed prediction, and final degradation map.

Table 4.6 gives the final vegetation-constrained area summary. The final map is conservative: most of the selected crop remains unchanged or healthy, while 1.62% is marked as early degradation and 1.91% as strong degradation.

Table 4.6: Final NEON vegetation-constrained degradation area summary.

Class	Class name	Pixels	Area (m2)	Percentage (%)
0	Unchanged / healthy	15,806	15,806	96.47
1	Early degradation	265	265	1.62
2	Strong degradation	313	313	1.91

**Figure 4.6:** NEON pseudo-reference changed-class F1 comparison.

4.6 Ablation and Discussion

Table 4.7 reports the ablation results. The strongest and most consistent conclusion is that the true bi-temporal 3D formulation is better than the Siamese 2D variant on the labelled benchmarks. SE attention and focal loss have dataset-dependent effects, especially on NEON where pseudo-reference label quality influences the training signal.

Table 4.7: Ablation results across datasets.

Dataset	Variant	Accuracy (%)	Macro-F1 (%)	Changed-F1 (%)	Kappa
Bay Area	Full reference	98.46	97.77	96.53	0.9553
Bay Area	No SE	97.66	96.67	94.85	0.9334
Bay Area	No focal	97.98	97.11	95.52	0.9422
Bay Area	No SE + no focal	98.08	97.25	95.73	0.9449
Bay Area	Siamese 2D	95.45	93.69	90.36	0.8741
Hermiston	Full reference	98.98	98.64	97.97	0.9729
Hermiston	No SE	98.73	98.33	97.51	0.9666
Hermiston	No focal	98.72	98.30	97.47	0.9661
Hermiston	No SE + no focal	98.73	98.32	97.48	0.9663

Dataset	Variant	Accuracy (%)	Macro-F1 (%)	Changed-F1 (%)	Kappa
Hermiston	Siamese 2D	92.76	90.74	86.43	0.8152
NEON	Full reference	87.10	44.87	38.51	0.2972
NEON	No SE	82.40	42.90	31.60	0.2396
NEON	No focal	94.51	49.05	57.89	0.4977
NEON	No SE + no focal	95.72	51.85	64.74	0.5707
NEON	Siamese 2D	83.48	41.13	34.43	0.2142

Figure 4.7 visually compares changed-class F1 across ablation variants. The figure supports the interpretation that the true bi-temporal design is the main architectural contribution, while attention and focal loss should be described as useful but not universally superior components.




Figure 4.7: Changed-class F1 comparison for ablation variants across datasets.

Overall, the benchmark results support the proposed model as a strong and compact hyperspectral change-detection framework. The NEON case study demonstrates how the same pipeline can be applied to a forest-domain degradation problem, but it also shows the limitation of pseudo-reference evaluation. A scientifically careful reading is therefore that ST-HyperDegradeNet is benchmark-validated and practically useful for candidate forest-degradation screening, but field validation is required before operational ecological claims are made.

CHAPTER 5

CONCLUSION AND FUTURE SCOPE

5.1 Conclusion

This thesis presented ST-HyperDegradeNet, a spectral-temporal hyperspectral deep learning framework for early forest degradation detection and hyperspectral change analysis. The method uses a true bi-temporal 3D input representation, where Time 1 and Time 2 patches are preserved as an explicit temporal dimension. This design allows the model to learn temporal change patterns jointly with local spectral and spatial context.

The proposed framework was evaluated on Bay Area AVIRIS and Hermiston Hyperion as labelled hyperspectral change-detection benchmarks, and on NEON SOAP Forest as a weakly supervised forest-domain case study. On the labelled benchmarks, the proposed model achieved high accuracy and changed-class F1, showing that the method is effective under controlled ground-truth evaluation. On NEON, the method produced a conservative vegetation-constrained degradation map, identifying 1.62% early degradation and 1.91% strong degradation within the selected crop.

The ablation study showed that the true bi-temporal 3D design is the most consistent contribution of the model. SE attention and focal loss are useful design components, but their effect is dataset-dependent and must be interpreted in relation to label quality and class imbalance. The thesis therefore avoids overclaiming and presents NEON as pseudo-reference agreement rather than field-verified degradation accuracy.

In summary, the work demonstrates a reproducible, lightweight, and interpretable pipeline for connecting benchmark hyperspectral change detection with forest degradation screening. The thesis contributes a clear methodology, consistent evaluation protocol, and cautious interpretation framework suitable for further research in multi-temporal hyperspectral environmental monitoring.

5.2 Future Scope

The first direction for future work is field validated forest degradation mapping. The NEON case study should be extended with field plots, tree mortality records, canopy health observations, drought information, or even expert labelled degradation maps. In that way, the validation would let the model output be judged as ecological degradation accuracy, not only pseudo reference agreement.

The second direction is multi site and multi season evaluation. Forest degradation hinges on climate, vegetation type, disturbance process, and also seasonality. If the framework gets applied to multiple NEON sites, or other hyperspectral forest datasets, it would be a decent test of whether the model generalizes across ecological conditions. With longer time series, it might also separate persistent degradation from temporary phenological variation, which can be kinda tricky.

The third direction is multimodal learning. If we integrate LiDAR canopy height, multispectral time series, thermal data, rainfall, and field measurements, the resulting degradation map should become more ecologically meaningful. Hyperspectral data bring spectral sensitivity, while LiDAR and field measurements can deliver structural and biological context, in a more grounded way.

The fourth direction is semi supervised and domain adaptive training. Dense labelled hyperspectral data are hard to collect , so future models should rely on a small set of labelled samples together with unlabelled pixels, use self supervised pretraining, or even do cross dataset adaptation. These paths could make hyperspectral degradation detection more practical for operational environmental monitoring, and yes it should help the whole pipeline feel less fragile.

5.3 Social Impact

Early forest degradation detection has practical social and environmental relevance, it matters. Forests support biodiversity, carbon storage, water regulation, soil protection, and local livelihoods too. If we detect degradation before complete canopy loss, then decision-makers can prioritize field surveys, and also plan quicker intervention, basically sooner.

A hyperspectral degradation screening system can help environmental monitoring agencies by flagging possible stress zones, you know. It can also assist researchers who study drought stress, pest disturbance, fire effects, and longer-term ecosystem change. This method is not meant to replace field survey entirely; instead it can guide where field validation might be the most useful , in practice.

The social impact of this thesis depends on careful interpretation. Automated maps should not be used for punitive or high-stakes decisions without validation first. So this thesis stresses transparent limitations, careful pseudo-reference handling, and the need for field or ecological confirmation before any operational deployment, so we do not overclaim..

REFERENCES

- [1] A. F. H. Goetz, G. Vane, J. E. Solomon, and B. N. Rock, "Imaging spectrometry for earth remote sensing," *Science*, vol. 228, no. 4704, pp. 1147-1153, 1985.
- [2] P. S. Thenkabail, J. G. Lyon, and A. Huete, Eds., *Hyperspectral Remote Sensing of Vegetation*. Boca Raton, FL, USA: CRC Press, 2018.
- [3] D. Landgrebe, *Signal Theory Methods in Multispectral Remote Sensing*. Hoboken, NJ, USA: Wiley, 2003.
- [4] C. J. Tucker, "Red and photographic infrared linear combinations for monitoring vegetation," *Remote Sensing of Environment*, vol. 8, no. 2, pp. 127-150, 1979.
- [5] A. Singh, "Digital change detection techniques using remotely-sensed data," *International Journal of Remote Sensing*, vol. 10, no. 6, pp. 989-1003, 1989.
- [6] W. A. Malila, "Change vector analysis: An approach for detecting forest changes with Landsat," *Proceedings of the Machine Processing of Remotely Sensed Data Symposium*, 1980.
- [7] F. A. Kruse et al., "The spectral image processing system (SIPS)-interactive visualization and analysis of imaging spectrometer data," *Remote Sensing of Environment*, vol. 44, pp. 145-163, 1993.
- [8] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62-66, 1979.
- [9] L. Breiman, "Random forests," *Machine Learning*, vol. 45, pp. 5-32, 2001.
- [10] J. A. Richards and X. Jia, *Remote Sensing Digital Image Analysis*, 5th ed. Berlin, Germany: Springer, 2013.
- [11] Y. Chen, H. Jiang, C. Li, X. Jia, and P. Ghamisi, "Deep feature extraction and classification of hyperspectral images based on convolutional neural networks," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 10, pp. 6232-6251, 2016.
- [12] A. Ben Hamida, A. Benoit, P. Lambert, and C. Ben Amar, "3-D deep learning approach for remote sensing image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 8, pp. 4420-4434, 2018.

- [13] L. Mou, P. Ghamisi, and X. X. Zhu, "Deep recurrent neural networks for hyperspectral image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 55, no. 7, pp. 3639-3655, 2017.
- [14] M. E. Paoletti, J. M. Haut, J. Plaza, and A. Plaza, "Deep learning classifiers for hyperspectral imaging: A review," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 158, pp. 279-317, 2019.
- [15] Ganji Tejasree, "An extensive review of hyperspectral image classification and prediction: techniques and challenges," *Multimedia Tools and Applications*, 2024.
- [16] L. Zhang, M. Zhang, and H. Zhang, "Deep slow feature analysis for change detection in multi-temporal remote sensing images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 12, pp. 9878-9890, 2019.
- [17] M. Hu, C. Wu, L. Zhang, and B. Du, "Binary change guided hyperspectral multiclass change detection," *IEEE Transactions on Image Processing*, vol. 32, pp. 791-806, 2023.
- [18] Z. Li, X. Tang, W. Li, and R. Tao, "A hyperspectral change detection framework based on double-stream convolutional neural networks and attention," *Remote Sensing*, vol. 16, no. 5, 2024.
- [19] Y. Chen, Y. Wang, and L. Zhang, "Spectral-temporal transformer for hyperspectral image change detection," *Remote Sensing*, vol. 15, no. 14, 2023.
- [20] J. Hu, L. Shen, and G. Sun, "Squeeze-and-Excitation networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7132-7141, 2018.
- [21] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar, "Focal loss for dense object detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 2, pp. 318-327, 2020.
- [22] Cross-Dataset Model Training for Hyperspectral Image Classification Using Self-Supervised Learning, *IEEE Transactions on Geoscience and Remote Sensing*, 2024.
- [23] Towards Knowledge Guided Pretraining Approaches for Multimodal Foundation Models: Applications in Remote Sensing, *ACM*, 2026.

- [24] G. Congalton and K. Green, *Assessing the Accuracy of Remotely Sensed Data: Principles and Practices*, 3rd ed. Boca Raton, FL, USA: CRC Press, 2019.
- [25] National Ecological Observatory Network (NEON), "Airborne Observation Platform hyperspectral reflectance data product DP3.30006.002," NEON Data Portal.
- [26] G. F. Hughes, "On the mean accuracy of statistical pattern recognizers," *IEEE Transactions on Information Theory*, vol. 14, no. 1, pp. 55-63, 1968.
- [27] R. Harikumar, G. Harish, V. Kumar, and K. Soman, "Classification of hyperspectral images using spectral-spatial method," *International Journal of Computer Applications*, 2016.
- [28] J. Li, P. Marpu, A. Plaza, J. Bioucas-Dias, and J. Benediktsson, "Generalized composite kernel framework for hyperspectral image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 51, no. 9, pp. 4816-4829, 2013.
- [29] *Conventional to Deep Ensemble Methods for Hyperspectral Image Classification: A Comprehensive Survey*, 2024.
- [30] *Recent Advances on Spectral-Spatial Hyperspectral Image Classification*, 2024.
- [31] National Ecological Observatory Network, "Soaproot Saddle (SOAP) field site and Airborne Observation Platform data products," NEON Data Portal.
- [32] D. M. W. Powers, "Evaluation: From precision, recall and F-measure to ROC, informedness, markedness and correlation," *Journal of Machine Learning Technologies*, vol. 2, no. 1, pp. 37-63, 2011.
- [33] M. Story and R. G. Congalton, "Accuracy assessment: A user's perspective," *Photogrammetric Engineering and Remote Sensing*, vol. 52, no. 3, pp. 397-399, 1986.

Rishika Gupta

hsi classification

 Quick Submit Quick Submit Delhi Technological University

Document Details

Submission ID

trn:oid::1:3581811386

Submission Date

May 29, 2026, 3:28 PM GMT+5:30

Download Date

May 29, 2026, 3:33 PM GMT+5:30

File Name

version2_hsi.pdf

File Size

1.1 MB

47 Pages

9,813 Words

60,127 Characters

*% detected as AI

AI detection includes the possibility of false positives. Although some text in this submission is likely AI generated, scores below the 20% threshold are not surfaced because they have a higher likelihood of false positives.

Caution: Review required.

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.

Disclaimer

Our AI writing assessment is designed to help educators identify text that might be prepared by a generative AI tool. Our AI writing assessment may not always be accurate (i.e., our AI models may produce either false positive results or false negative results), so it should not be used as the sole basis for adverse actions against a student. It takes further scrutiny and human judgment in conjunction with an organization's application of its specific academic policies to determine whether any academic misconduct has occurred.

hsi classification

by Rishika Gupta

Submission date: 29-May-2026 03:28PM (UTC+0530)

Submission ID: 2971929822

File name: version2_hsi.pdf (1.11M)

Word count: 9813

Character count: 60127

hsi classification

ORIGINALITY REPORT

8%

SIMILARITY INDEX

7%

INTERNET SOURCES

5%

PUBLICATIONS

3%

STUDENT PAPERS

PRIMARY SOURCES

1	Submitted to Delhi Technological University Student Paper	1%
2	www.dspace.dtu.ac.in:8080 Internet Source	1%
3	dspace.dtu.ac.in:8080 Internet Source	1%
4	uwspace.uwaterloo.ca Internet Source	1%
5	Submitted to dtusimilarity Student Paper	1%
6	Rodrigo Guedes Pereira Pinheiro, Claudia Lage Rebello da Motta. "A Patch-Based Computational Framework for the Analysis of Structurally Heterogeneous Bioelectrographic Images", Applied Sciences, 2026 Publication	<1%
7	export.arxiv.org Internet Source	<1%
8	www.cse.griet.ac.in Internet Source	<1%
9	www.mdpi.com Internet Source	<1%
10	Tao Lei, Shuxin Zhang, Shaoxiong Lin, Tongfei Liu, Zhiyong Lv, Tao Gao, Maoguo Gong, Asoke K. Nandi. "Remote sensing image change detection using deep learning	<1%

techniques: a comprehensive survey",
Artificial Intelligence Review, 2026

Publication

11	www.ecse.rpi.edu Internet Source	<1 %
12	htwk-leipzig.qucosa.de Internet Source	<1 %
13	ia802200.us.archive.org Internet Source	<1 %
14	irlib.ambou.edu.et Internet Source	<1 %
15	aranne5.bgu.ac.il Internet Source	<1 %
16	calhoun.nps.edu Internet Source	<1 %
17	123dok.net Internet Source	<1 %
18	Jan Askne. "Sensors and Environmental Applications of Remote Sensing", CRC Press, 2026 Publication	<1 %
19	Rashmi Agrawal, Marcin Paprzycki, Neha Gupta. "Big Data, IoT, and Machine Learning - Tools and Applications", CRC Press, 2020 Publication	<1 %
20	auetd.auburn.edu Internet Source	<1 %
21	essay.utwente.nl Internet Source	<1 %
22	idr.nitkkr.ac.in:8080 Internet Source	<1 %

23	Internet Source	<1 %
24	oaklandfocus.blogspot.com Internet Source	<1 %
25	Ping Jian, Yimin Ou, Keming Chen. "Uncertainty Aware Graph Self-Supervised Learning for Hyperspectral Image Change Detection", IEEE Transactions on Geoscience and Remote Sensing, 2024 Publication	<1 %
26	d.lib.msu.edu Internet Source	<1 %
27	iris.unitn.it Internet Source	<1 %
28	theses.hal.science Internet Source	<1 %
29	tudr.thapar.edu Internet Source	<1 %
30	Jiao Jiao, Yue Zhang, Hao Sun, Xue Yang, Xun Gao, Wen Hong, Kun Fu, Xian Sun. "A Densely Connected End-to-End Neural Network for Multiscale and Multiscene SAR Ship Detection", IEEE Access, 2018 Publication	<1 %
31	Yichu Xu, Di Wang, Lefei Zhang, Liangpei Zhang. "Dual selective fusion transformer network for hyperspectral image classification", Neural Networks, 2025 Publication	<1 %
32	Ziqiang Wang, Zhi Liu, Gongyang Li, Yang Wang, Tianhong Zhang, Lihua Xu, Jijun Wang. "Spatio-Temporal Self-Attention Network for	<1 %

Video Saliency Prediction", IEEE Transactions on Multimedia, 2021

Publication

33	apps.dtic.mil Internet Source	<1 %
34	d197for5662m48.cloudfront.net Internet Source	<1 %
35	dehesa.unex.es Internet Source	<1 %
36	repository.uantwerpen.be Internet Source	<1 %
37	technodocbox.com Internet Source	<1 %
38	umpir.ump.edu.my Internet Source	<1 %
39	Pedram Ghamisi, Naoto Yokoya, Jun Li, Wenzhi Liao, Sicong Liu, Javier Plaza, Behnood Rasti, Antonio Plaza. "Advances in Hyperspectral Image and Signal Processing: A Comprehensive Overview of the State of the Art", IEEE Geoscience and Remote Sensing Magazine, 2017 Publication	<1 %
40	Nicolas R. Dalezios. "Remote Sensing Applications in Environmental and Earth System Sciences", CRC Press, 2021 Publication	<1 %

Exclude quotes On

Exclude matches Off

Exclude bibliography On