

محمود الراعي

Hussien and Mahmoud Final Project -1 V17.docx

 Midocean University

Document Details

Submission ID

trn:oid::2945:378918691

Submission Date

Apr 22, 2026, 9:30 AM GMT+2

Download Date

Apr 22, 2026, 10:33 AM GMT+2

File Name

Hussien and Mahmoud Final Project -1 V17.docx

File Size

306.1 KB

14 Pages

3,854 Words

22,838 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.

Frequently Asked Questions

How should I interpret Turnitin's AI writing percentage and false positives?

The percentage shown in the AI writing report is the amount of qualifying text within the submission that Turnitin's AI writing detection model determines was either likely AI-generated text from a large-language model or likely AI-generated text that was likely revised using an AI paraphrase tool or word spinner.

False positives (incorrectly flagging human-written text as AI-generated) are a possibility in AI models.

AI detection scores under 20%, which we do not surface in new reports, have a higher likelihood of false positives. To reduce the likelihood of misinterpretation, no score or highlights are attributed and are indicated with an asterisk in the report (*%).

The AI writing percentage should not be the sole basis to determine whether misconduct has occurred. The reviewer/instructor should use the percentage as a means to start a formative conversation with their student and/or use it to examine the submitted assignment in accordance with their school's policies.

What does 'qualifying text' mean?

Our model only processes qualifying text in the form of long-form writing. Long-form writing means individual sentences contained in paragraphs that make up a longer piece of written work, such as an essay, a dissertation, or an article, etc. Qualifying text that has been determined to be likely AI-generated will be highlighted in cyan in the submission, and likely AI-generated and then likely AI-paraphrased will be highlighted purple.

Non-qualifying text, such as bullet points, annotated bibliographies, etc., will not be processed and can create disparity between the submission highlights and the percentage shown.





Prepared By

Name Mahmoud Al Ra'e

ID No.: 253001641

Name Hussein Mudeheg

ID No.: 253001186

Submitted in fulfillment of the requirements for the Master degree of
Science in Artificial Intelligence

At

College of Informatics Midocean University, UAE.

Supervisor Name: Dr. Hager Saleh

College of Informatics, Midocean University, Comoros

Contents

1.	Chapter 1: Introduction.....	3
1.1.	Problem Definition	3
1.2.	Thesis Objectives.....	3
1.3.	Thesis Contribution	3
1.4.	Thesis Organization.....	3
1.5.	Proposed Framework.....	4
2.	Chapter 2 Background and Related Work.....	5
2.1.	Background	5
2.1.1.	Artificial Intelligence (AI).....	5
2.1.2.	Deep Learning (DL)	5
2.1.3.	Convolutional Neural Networks (CNN)	6
2.1.4.	Transform Model.....	6
2.1.5.	Vision Transformer (ViT).....	6
2.1.6.	Internet of Things (IoT)	6
2.2.	Related Work.....	7
References	12	

1. Chapter 1: Introduction

1.1. Problem Definition

Plants diseases are among the major concerns affecting the quality and production levels of crops. The farmers often determine the occurrence of diseases in the plant through experience, which may be unreliable. Early symptoms of plant disease may be difficult to recognize leading to delay in diagnosis, thus more damage to the plant. Some diseases in plants look alike; therefore, it could lead to a false conclusion in the diagnosis. Detecting the diseases of plants in their natural environment becomes very difficult owing to many factors including illumination, occlusion, and leaf orientation of the plant. Normally, checking of the plants requires manual labor. The detection of the diseases from each individual plant becomes time-consuming particularly when dealing with many of them or a larger piece of farm. It is, thus, necessary to develop a system which will enable farmers to diagnose the diseases automatically. Artificial Intelligence and deep learning would help the farmers detect diseases faster and make informed decisions concerning the crops. Plant diseases account for about 40% annual losses globally. Plant diseases lead to huge financial losses annually, and they have been estimated to cost about 220 billion dollars annually globally.

1.2. Thesis Objectives


The general objective of this thesis is to develop an intelligent system for plant disease detection from leaf images using AI and DL techniques. The study aims to improve the accuracy and speed of disease diagnosis in order to support early detection and better management.

The specific objectives are to apply and evaluate different DL models for plant disease classification, particularly Convolutional Neural Networks (CNNs), Vision Transformers (ViT), and Swin Transformers. In addition, this thesis aims to compare these models in terms of classification accuracy, computational efficiency, and suitability for real-world agricultural environments. Another specific objective is to examine the ability of these models to work under practical conditions such as varying lighting, background complexity, and leaf position. The study also seeks to explore how these models can support smart agriculture applications when combined with modern technologies such as the Internet of Things (IoT).

1.3. Thesis Contribution

1.4. Thesis Organization

1.5. Proposed Framework



2. Chapter 2 Background and Related Work

2.1. Background

Technology has started to play a major role in many areas including the agriculture and plants, especially in the health of Crop and productivity. One of the most challenges faced by the farmers is the early detection of the plant diseases. This process earlier has a large dependency on the expertise of the human which is time consuming, could have a biased and sometimes are inaccurate. With the advancement of the technology and Ai, newly automated approaches have emerged to automate the process disease detection that making it faster and more reliable

2.1.1. Artificial Intelligence (AI)

AI refers to the ability of machines to perform tasks that normally require human intelligence, such as learning, reasoning and making a decision. In agriculture, AI is being increasingly used to analyze the plant conditions, detect the diseases based on the image and ultimately support in taking the right decision in by processing large amount of data and identifying patterns which enable early diagnostic for better crop management [1].

2.1.2. Deep Learning (DL)

Deep learning (DL) is a subset of the Artificial Intelligence (AI) that focuses on learning from large scale datasets using multi-layered neural network, enabling systems to automatically extract, meaningful and hierarchical features without the need for manual feature engineering or extensive intervention of Human. In plant disease detection, DL has demonstrated high efficiency due to its ability to analyze complex visual patterns in leaf images and identifying plant diseases under controlled conditions [2]

2.1.3. Convolutional Neural Networks (CNN)

Convolutional Neural Network (CNN) is one of the most popular deep learning models for image analysis. This is due to its ability to learn spatial hierarchies of features from input data. Many different types of CNN architecture exist, including those like LeNet, AlexNet, and their more recent versions, namely VGG, GoogLeNet, ResNet, DenseNet, and MobileNet. CNN works excellently in processing visual data and identifying patterns within an image, including edges, textures, and shapes. Within the agricultural sector, CNN is commonly used for disease diagnosis of plants through analysis of leaf images, with architectures such as VGG, ResNet, and DenseNet achieving an accuracy above 95%. [3]

2.1.4. Transform Model

The transformer architecture is a approach to deep learning, based around the concept of attention instead of the conventional form of convolution. Unlike the CNN, which cuts up the picture into smaller portions, the Transformer analyzes the entire image in its entirety, understanding the connections between various segments. It allows it to understand the broader picture and dependencies across long distances. Due to these properties, architectures using transformers such as Vision Transformers (ViTs) have become increasingly popular for image analysis compared to CNNs and other methods. [4]

2.1.5. Vision Transformer (ViT)

ViT is a model that applies transformer architecture to image data, and instead of analyzing the complete image, it divides it into smaller patches and then processes them in a similar way to words in a sentence. This approach allows the model to understand both local and global features. Many recent research studies show that ViT models could outperform traditional CNNs in plant disease classification tasks. Some studies have compared ViT with models like ResNet and DenseNet, which have reported accuracy levels approximately close to 99% that demonstrate strong performance and capabilities. However, the side effect of the ViT models often is that they require large datasets and higher computational resources which could limit the accessibility, especially for smaller research teams or applications with limited data availability. [5]

2.1.6. Internet of Things (IoT)

IoT is a network of devices that are connected and can share and gather data in real time. IoT is very important for smart farming systems in agriculture. Cameras, sensors, and drones are some of the devices that can keep an eye on crops all the time and send data for analysis. IoT makes it possible to have real-time plant disease detection systems when used with AI and deep learning. For instance, cameras with IoT can take pictures of leaves in the field, and deep learning models can quickly look at these pictures to find diseases. [6]

2.2. Related Work

Too et al. applied multiple types of pertained models such as VGG16, InceptionV4, ResNet50, and DenseNet121 which classify 38 different classes of plant diseases. The public dataset so called PlantVillage was used from KAGGLE, which has 54,306 images with 38 classes. The preprocessing steps included image resizing in order to match the required input size of the CNN architecture, and rotation, flipping, and scaling were applied to increase dataset variability and reduce overfitting. The result has shown that DenseNet121 achieved the highest classification accuracy among the evaluated models. Transfer learning significantly improved performance compared to training from scratch. The models demonstrated very high accuracy under laboratory conditions. [7]

The study examines the use of Vision Transformers (ViT) for plant disease classification to determine whether attention-based architectures can perform better than the traditional convolutional neural networks (CNNs) or not. A pre-trained ViT model on ImageNet was improved for plant disease detection and compared with CNN models such as ResNet50 and DenseNet121. The model utilizes five public plant datasets that include the famous PlantVillage dataset, which contains 54,306 images from 38 classes that were collected under controlled conditions. The preprocessing involved resizing the images to 244X244 pixels, patch embedding preparation, pixel normalization, and data augmentation such as rotation and flipping. The final results demonstrate that ViT achieved approximately 99% accuracy, which outperforms several models. [8]

The research proposed a lightweight deep learning model, so called Dise-Efficient which is built on EfficientNetV2 to identify plant diseases and pest identification. The study focused on improving model efficiency while ensuring strong recognition performance. The model was trained and evaluated using plant image datasets, using steps such as image resizing, normalization, and data augmentation for preprocessing, which were applied to enhance performance and generalization. The results show that the proposed model has achieved high accuracy of classification as well as reduced the computational complexity. [9]

The authors have proposed an integrated Internet of Things (IoT) solution that has combined deep learning techniques to detect plant diseases. Leaf images were captured through cameras along with environmental data collected from IoT sensors such as temperature and humidity. The purpose of the preprocessing techniques is to improve image quality prior to applying deep learning models for disease classification. The system includes training and evaluation stages that measure the performance of the model, which resulted in high accuracy. In addition, the framework supports real-time monitoring and feedback, allowing continuous tracking of crop conditions and early detection of diseases. [10]

Aldakheel et al. proposed a real time plant leaf disease detection system based on the YOLOv4 object detection framework. In contrast to the traditional classification models that only predict disease categories, YOLOv4 enables simultaneous detection of diseased regions within leaf images. The methodology involved training the YOLOv4 model to identify and classify multiple plant diseases directly from leaf images by bounding box coordinates and class probabilities in a single detection process. The model was trained on annotated plant leaf disease datasets that contain multiple disease categories in several environments. The preprocessing steps included image resizing to match the YOLOv4 input dimensions, image normalization, and data augmentation such as rotation, scaling, and flipping to improve robustness and at the same time to reduce the overfitting. The result of the experimentation demonstrated that YOLOv4 achieved higher accuracy with fast inference speed which making it suitable for real time agricultural applications. [11]

Duhan et al. have designed a real-time plant disease diagnosis based on mobile and edge devices. The authors have combined the RTR_Lite framework with MobileNetV2, a compact convolutional neural network known for its use of separable convolutions, which significantly reduce computational cost and still maintain strong performance. The goal was to develop a model that balances between accuracy and speed, which makes it very practical for field deployment where the computational resources are always limited. The model was trained on available public plant disease datasets that include various crops and disease categories which were captured under different environmental conditions. Prior to commencing the training the images were resized to match the input size required by MobileNetV2, and pixel values were normalized. Data augmentation techniques such as rotation, flipping, and scaling were applied to improve the ability of the models to generalize. The results showed that RTR_Lite_MobileNetV2 achieved high classification accuracy and at the same time has maintained a small model size and fast inference speed. [12]

Bhola and Kumar have designed a CNN model that automatically learns the important visual features from leaf images which eliminates the need for manual feature extraction. The main objective was to support farmers in detecting diseases at an early stage so that timely treatment can reduce crop damage and yield loss. Pre-trained CNN models such as VGG16, ResNet-50, and Inception-v3 were used. The model was trained and tested on a dataset of vegetable leaf images that included several disease categories, collected under both controlled and natural environmental conditions. Before starting the training, the images were resized to fit the input requirements of the network and normalized to standardize pixel values. Data augmentation techniques such as rotation, flipping, and brightness adjustment were also applied to increase dataset diversity and accordingly improve the ability of the models to generalize. The results showed that the proposed model achieved high accuracy and performed well in recognizing early symptoms of the disease. [13]

Kheir et al. have proposed a hybrid model combined of Vision Transformer (ViT) and MobilePlantViT, the lightweight design generalized plant disease image classification and balance both accuracy and efficiency for mobile and edge deployment. Traditional deep learning models for plant disease recognition often struggle on resource limited devices caused by high computational costs. MobilePlantViT combines efficient convolutional operations with a Vision Transformer backbone to capture local and global features while maintain the architecture compact. The model

was evaluated across multiple diverse plant disease datasets that has shown a strong generalization and test accuracies ranging from around 80% to over 99% depending on the dataset. Preprocessing consisted of the usual image resizing and normalization, along with data augmentation to improve its reliability and prevent overfitting. The outcomes from the experiments highlight the potential that MobilePlantViT holds for practical disease detection in plants. [14]

Thanjaivadivel developed EnConv which is an improved CNN, meant to accurately classify plant leaves and their diseases using more efficient feature extraction techniques. The objective is to use ML and DL to detect diseases in plants which does not need extensive knowledge in the field. EnConv uses deep convolutional layers, which optimize calculations and retain crucial properties such as leaf colors and shapes. Instead of having a flattening layer as in most CNNs, EnConv has global pooling layers that can reduce overfitting and improve performance. After testing EnConv on data from a total of 39 diseases from different crops, it had a maximum accuracy rate of 99.87%. This clearly indicates that EnConv is more accurate than other machine learning and deep learning models. [15]

In this research, we have proposed a deep learning and IoT-based framework for detecting plant leaf diseases in real-time. The proposed framework relies on a CNN algorithm to analyze leaf images and determine whether the plant is healthy or unhealthy. For improving the efficiency of pre-trained models such as MobileNetV2, ResNet50, and VGG16, transfer learning technology is used, along with datasets including PlantVillage. An IoT camera is used for taking images from the leaves and performing operations such as removing noise, adjusting size, and applying data augmentation techniques to enhance reliability irrespective of the weather conditions. Furthermore, the developed model is capable of executing at edge computing devices including Raspberry Pi or Jetson Nano. The model has an accuracy rate of more than 90%. [16]

The paper explains a framework that utilizes DL techniques to identify and classify plant leaf diseases is introduced. A custom model named CropLeafNet is constructed using reliable backbone architectures ResNet50, MobileNetV2, and VGG16 to enhance feature detection and classification processes. Using the concept of transfer learning, the proposed framework identifies not only the local characteristics such as the leaf textures but also the global features including the leaf colors and patterns found all around the world. This makes it possible to diagnose a variety of diseases through one single model. For training purposes, datasets such as PlantVillage as well as photographs from the field are employed to guarantee a wide range of leaf environments. Image preprocessing consists of resizing and normalizing operations along with some augmentation operations including rotation, flipping, and contrast changing. From the experimental results, it is clear that the developed model named CropLeafNet (CNN architecture for agricultural problems) is able to classify the diseases with high accuracy and performs better than other standard CNN models. [17]

In this study, an ML approach is proposed to detect plant disease in pigeon pea leaves, intending to provide assistance for early diagnosis for farmers. This paper explicitly uses no deep neural network nor other recent large-scale models called black boxes. The researchers propose a novel ML solution to the problem by combining classical computer vision techniques and classical classifiers. Dense SIFT (DSIFT) and Bag of Visual Words (BoVW) techniques are used as feature extractors that have proven themselves to be robust and invariant to noise and rotations in previous works. This framework tests several classifiers, including SVM, KNN, Random Forest, Decision Tree, Gradient Boosting, and a CNN classifier, in order to classify the pigeon pea leaf disease categories using their extracted descriptors. The dataset consists of 900 wild images of pigeon pea leaves (diseased and healthy). The preprocessing stage involved image size adjustment and enhancement of quality. It was discovered that the CNN classifier and the LightGBM classifier were efficient in classifying pigeon pea leaves into categories. [18]

PlantDoc was designed as a real-life dataset, which is supposed to aid in improving the abilities of artificial intelligence to recognize plant diseases under non-laboratory circumstances. As such, all the images used in PlantDoc come from a variety of farms, thus being quite messy, with lots of background information, inconsistent illumination, and pictures shot from a wide range of perspectives. A series of deep learning models were chosen by the authors to see how successfully they would be able to recognize plant diseases. Prior to the training phase, all the images were processed, having been resized, cropped using bounding boxes for disease localization, normalized, and subjected to various augmentations, like images' flip and adjustment of brightness levels. This yielded some positive results, as PlantDoc-trained models demonstrated higher robustness when applied to real farm conditions than other models, which were trained exclusively based on laboratory data. At the same time, it resulted in somewhat worse performance overall due to the nature of real-life data. [19]

In this work, it was examined how EfficientNet, which is known for scaling depth, width, and resolution together, performs regarding recognizing plant diseases. Researchers used different versions of EfficientNet, from its smaller version B0 to a larger one, B7. Transfer learning was performed based on existing plant disease recognition datasets such as PlantVillage. The data was processed thoroughly before entering the training phase. Specifically, the images were resized according to the requirements of the model, normalized, and augmented using techniques like rotation, flipping, zooming, and others. It easily beat the performance of the previous versions of CNNs, like ResNet and Inception, while being even more efficient due to using less number of parameters. The classification accuracy exceeded 99%. It should be noted that these tests were conducted using clean datasets. [20]

This paper describes ResNet, which is a deep convolutional neural network using residual learning to solve the problem of degradation in extremely deep models. Designed for generalized image recognition, ResNet has shown impressive performance in plant disease classification due to its efficient image feature extraction capabilities. One of the main concepts covered in the paper is the skip connection, which helps retain the continuity of gradient flow during training and enables ultra-deep neural nets, including ResNet50, ResNet101, and ResNet152.

ResNet is fine-tuned using transfer learning on platforms like PlantVillage. This process involves preprocessing the images which includes resizing, normalizing the data and using data augmentation to increase the size of the dataset. In various research studies, the model has consistently achieved impressive accuracy rates, as seen in reference [21].

Ref	Best Model	Dataset Source & Link	Result
Too et al., 2019 [7]	DenseNet (Fine-tuned)	PlantVillage Dataset – Kaggle: https://www.kaggle.com/datasets/emmarex/plantdisease	~99.7% accuracy
Chen et al., 2022 [8]	Vision Transformer (ViT-B/16)	PlantDoc: https://www.kaggle.com/datasets/abdulhasibuddin/plant-doc-dataset	~99% accuracy, strong generalization in the wild
Xie et al., 2023 [9]	Disc-Efficient (EfficientNetV2-based)	PlantVillage , IP102 and mixed plant disease & pest datasets https://www.frontiersin.org/journals/plant-science/articles/10.3389/fpls.2023.1227011/full?utm_source=chatgpt.com	98–99% accuracy, lightweight & efficient performance
Talaat et al., 2026 [10]	CNN-based IoT-integrated system	Field-collected dataset (custom) + PlantVillage	~97–99% accuracy, real-time detection with IoT integration
Aldakheel, E. A., Zakariah, M., & Alabdallal, A. H. (2024) [11]	YOLOv4	PlantVillage (Detection version)	mAP ≈ 95%, real-time multi-disease detection
Duhan et al., 2025 [12]	RTR_Lite + MobileNetV2	Multiple datasets: PlantVillage + Mendeley Data (Soybean, Coffee, Wheat, Sugarcane) https://data.mendeley.com/	~98% accuracy, low latency mobile deployment
Bhola & Kumar, 2025 [13]	Lightweight CNN-based model	Vegetable Leaf Dataset (Custom Field Dataset)	~92–93% accuracy, efficient lightweight model for early-stage disease detection
Tonmoy et al., 2025 [14]	MobilePlantViT (Hybrid CNN + ViT)	Multi-crop dataset (PlantVillage + Field images)	~80–99% accuracy, lightweight
Castillo Ossa & Corchado, 2025 [15]	EnConv (Enhanced Convolutional Neural Network)	https://link.springer.com/article/10.1007/s41348-024-01033-6 DataSet: PlantVillage	~98–99% accuracy, improved interpretability
Real-Time Monitoring System [16]	CNN + Deep Learning models (VGG16,VGG19,InceptionV3,MobileNet,DenseNet121,Xception)	https://www.sciencedirect.com/science/article/pii/S2772899424000417 PlantVillage + Real-time camera dataset	~95–97% real-time detection accuracy
Conditional Mutual Information DL [17]	CNN + CMI Regularization	PlantVillage (PDF) Conditional Mutual Information Constrained Deep Learning: Framework and Preliminary Results	2–4% improvement over baseline CNN
Rajput & Doddamani, 2024 [18]	BoVW + DSIFT + SVM	Pigeon Pea Leaf Dataset (Custom)	~92–95% accuracy
Singh et al., 2020 [19]	Faster R-CNN / SSD	PlantDoc Dataset – https://github.com/pratikkayal/PlantDoc-Dataset	Better real-world performance
Tan & Le, 2019 [20]	EfficientNet	PlantVillage – https://www.kaggle.com/datasets/emmarex/plantdisease	>99% accuracy
He et al., 2016 [21]	ResNet	ImageNet – https://www.image-net.org/	High accuracy baseline

References

- [1] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. Pearson, 2021.
- [2] S. Mohanty, D. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Frontiers in Plant Science*, vol. 7, p. 1419, 2016
- [3] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [4] A. Dosovitskiy et al., "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," *International Conference on Learning Representations (ICLR)*, 2021.
- [5] X. Zhang et al., "A Multitask Learning-Based Vision Transformer for Plant Disease Classification," *Springer*, 2024.
- [6] K. Rose, S. Eldridge, and L. Chapin, "The Internet of Things: An Overview – Understanding the Issues and Challenges of a More Connected World," *Internet Society*, 2015.
- [7] S. Too, L. Yujian, S. Njuki, and L. Yingchun, "A comparative study of fine-tuning deep learning models for plant disease identification," 2019.
- [8] X. Chen et al., "Vision Transformer for Plant Disease Classification in the Wild: A Multimodal Approach," 2022.
- [9] X. Xie et al., "A lightweight model for efficient identification of plant diseases and pests based on deep learning," 2023.
- [10] F. M. Talaat and M. Ibrahim, "IoT-Integrated Robotic System for Automated Plant Disease Detection and Environmental Monitoring," 2023.
- [11] M. Rahman et al., "Detection and Identification of Plant Leaf Diseases Using YOLOv4," 2024.
- [12] A. Duhan et al., "RTR_Lite_MobileNetV2 for Plant Disease Detection," 2025.
- [13] R. Bhola and A. Kumar, "Deep Learning Model for Early Vegetable Disease Detection," 2025.
- [14] M. Kheir et al., "MobilePlantViT: A Mobile-Friendly Hybrid Vision Transformer for Generalized Plant Disease Image Classification," 2025.
- [15] M. Thanjaivadivel, "Enhanced CNN for Leaf Disease Classification," 2025.
- [16] "A Real-Time Monitoring System for Accurate Plant Leaves Disease Detection Using Deep Learning," 2024.
- [17] "Conditional Mutual Information Constrained Deep Learning for Classification," 2023.
- [18] S. Rajput and S. Doddamani, "Pigeon Pea Leaf Disease Classification Using BoVW and DSIFT," 2024.
- [19] D. Singh et al., "PlantDoc: A Dataset for Visual Plant Disease Detection," 2020.
- [20] M. Tan and Q. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," 2019.
- [21] K. He et al., "Deep Residual Learning for Image Recognition," 2016.

