

Crop Disease Detection from Transfer Adaptation in Smallholder Farm Images: Graph Analysis

Sora Endo*

Graduate School of Informatics, Kyoto University, Kyoto, Japan

Corresponding author: sora.endo@kyoto-u.ac.jp

Abstract

The early and accurate detection of crop diseases is paramount for ensuring global food security, particularly in the context of smallholder farming, which contributes significantly to the global food supply. Smallholder farm environments present unique challenges for automated visual disease detection, including heterogeneous crop varieties, complex backgrounds, variable illumination, and high degrees of occlusion. While deep convolutional neural networks have shown promise in controlled laboratory settings, their performance degrades substantially when deployed in real-world agricultural environments due to the domain shift between training and testing data. This paper proposes a novel framework that integrates transfer adaptation with graph analysis to predict and detect crop diseases in smallholder farm images. By conceptualizing the agricultural image as a structured graph where superpixels act as nodes and spatial relationships define the edges, the proposed method captures complex topological features that are robust to environmental noise. Transfer adaptation techniques are employed to align the feature distributions between source domains comprising high-quality, controlled images and target domains consisting of noisy, real-world smallholder images. Comprehensive evaluations demonstrate that the integration of graph-based relational reasoning with domain adaptation significantly improves classification accuracy and generalizability. The findings suggest that modeling the structural context of plant diseases can overcome the limitations of traditional grid-based feature extraction paradigms.

Keywords: Crop Disease Detection; Transfer Adaptation; Graph Analysis; Smallholder Farms; Domain Shift

1. Introduction

1.1 Background and Motivation

Global agricultural productivity faces continuous threats from a diverse array of plant pathogens, pests, and environmental stressors. Among the most vulnerable stakeholders in the agricultural sector are smallholder farmers, who manage small plots of land and often rely on their harvests for both sustenance and economic survival. Smallholder farms account for a substantial proportion of the food produced in developing regions, yet these farmers typically lack access to advanced agricultural extension services, specialized agronomic expertise, and expensive diagnostic tools. Consequently, crop diseases often go undetected until they reach an advanced stage, leading to devastating yield losses, economic hardship, and localized food insecurity. The advent of digital agriculture and artificial intelligence has introduced the possibility of deploying automated, smartphone-based disease detection systems directly into the hands of these farmers. Such systems have the potential to democratize access to expert-level diagnostic capabilities, enabling timely interventions and precision agriculture practices that were previously out of reach.

However, the translation of artificial intelligence research into viable solutions for smallholder farmers is fraught with technical difficulties. The vast majority of early research in automated plant disease detection relied on datasets collected under highly controlled laboratory conditions. In these datasets, leaves are typically detached, flattened, and photographed against uniform backgrounds with consistent lighting. Algorithms trained on such pristine data achieve exceptional accuracy rates but fail spectacularly when exposed to the chaotic visual environment of a real farm. Smallholder farm images are characterized by intense sunlight or deep shadows, overlapping foliage, non-target plant species, soil background clutter,

and symptoms of multiple concurrent diseases or nutrient deficiencies. This fundamental discrepancy between the training data and the operational data is known as domain shift, and it represents one of the most significant barriers to the successful deployment of agricultural machine learning models. Addressing this domain shift requires moving beyond simple feature extraction and embracing methodologies that can understand the structural and contextual essence of disease symptoms, irrespective of the environmental noise surrounding them.

1.2 Problem Statement

The central problem addressed in this research is the pronounced degradation in the performance of computer vision models when tasked with identifying crop diseases in unconstrained, real-world images from smallholder farms, following training on controlled, idealized datasets. Standard deep learning architectures process images as regular grids of pixels, prioritizing local textures and colors. While these features are highly discriminative in a laboratory setting, they are easily confounded by the extreme variability of field conditions. For instance, a shadow cast by an upper canopy leaf might be misinterpreted as a necrotic lesion, or the texture of dry soil in the background might be confused with a fungal infection. Furthermore, conventional domain adaptation techniques, which attempt to align the overall statistical distributions of the source and target domains, often struggle to preserve the fine-grained, localized details necessary to distinguish between visually similar diseases.

To overcome these limitations, there is a critical need for an analytical approach that models the inherent structure of the plant and the spatial distribution of disease symptoms. Diseases do not manifest as random collections of pixels; they spread across leaf veins, concentrate at the margins, or form specific clustering patterns. Traditional convolutional networks implicitly capture some of this spatial information but lack the explicit mechanisms to model irregular geometric relationships and long-range dependencies across the image. Therefore, the challenge lies in developing a framework capable of extracting robust, structurally aware representations of plant diseases and seamlessly transferring this knowledge across disparate visual domains.

1.3 Research Objectives

This study aims to formulate, implement, and evaluate a comprehensive framework for crop disease detection that synergizes transfer adaptation methodologies with advanced graph analysis. The primary objective is to demonstrate that representing farm images as topological graphs enables the extraction of relational features that are inherently more resilient to the environmental variations typical of smallholder agriculture. By framing disease detection as a graph classification problem, the research seeks to explicitly model the interactions between different regions of a leaf, thereby distinguishing genuine pathological symptoms from background clutter and lighting artifacts.

A secondary objective is to bridge the domain gap between controlled source datasets and unconstrained target datasets using adversarial adaptation strategies applied directly to the graph-based representations. The research investigates how aligning the latent structural features of disparate domains can facilitate the effective transfer of diagnostic knowledge. Ultimately, this paper strives to provide a robust algorithmic foundation for the next generation of mobile agricultural diagnostic tools, ensuring that they perform reliably in the hands of the smallholder farmers who need them most.

2. Literature Review

2.1 Traditional Crop Disease Detection

The endeavor to automate the detection of crop diseases spans several decades, initially relying on classical computer vision and image processing paradigms. Early approaches typically involved a sequential pipeline of image segmentation, feature extraction, and classification using traditional machine learning algorithms. Researchers utilized hand-crafted features based on color spaces, such as Hue Saturation Value and Lab color spaces, to isolate diseased lesions from healthy leaf tissue. Texture analysis played a crucial role, with techniques like Gray Level Co-occurrence Matrices and Local Binary Patterns being widely employed to capture the microscopic variations indicative of fungal infections or viral blights. Morphological features, including area, perimeter, and shape descriptors of the identified lesions, were also calculated to differentiate between diseases that produce similarly colored spots but with distinct geometric characteristics.

Once these features were extracted, algorithms such as Support Vector Machines, Random Forests, and k-Nearest Neighbors were trained to perform the final classification [1]. While these traditional methods provided valuable proof-of-concept demonstrations and required relatively little computational power, they were fundamentally limited by their reliance on manual feature engineering. The effectiveness of the system was entirely dependent on the researcher's ability to mathematically define what a disease looks like, a task that becomes impossibly complex when accounting for the vast morphological diversity of pathogens and host plants. Moreover, the segmentation techniques utilized in these early systems were highly sensitive to illumination changes and background complexity, restricting their applicability to laboratory environments where these variables could be strictly controlled.

2.2 Deep Learning and Transfer Adaptation

The introduction of deep convolutional neural networks revolutionized the field of agricultural image analysis, shifting the paradigm from manual feature engineering to automated feature learning. Architectures capable of learning hierarchical representations directly from raw pixel data demonstrated unprecedented accuracy on large-scale, standardized datasets. However, as researchers began to deploy these models into field environments, the limitations of standard convolutional approaches became apparent. The models were essentially overfitting to the specific environmental conditions of the training set, learning to associate certain background textures or lighting angles with specific disease classes [2]. This phenomenon highlighted the urgent need for techniques capable of generalizing across different visual domains.

Transfer adaptation, a specialized subfield of transfer learning, emerged as a promising solution to this domain shift problem. The goal of transfer adaptation is to minimize the discrepancy between the feature distributions of a source domain, where labeled data is abundant, and a target domain, where labeled data is scarce or nonexistent. Various strategies have been proposed, including discrepancy-based methods that seek to minimize metrics like Maximum Mean Discrepancy between the domains. More recently, adversarial adaptation techniques have gained prominence [3]. These methods typically employ a feature generator network that attempts to learn domain-invariant representations, while a domain discriminator network simultaneously attempts to distinguish between source and target features. Through this adversarial training process, the generator is forced to produce features that are highly discriminative for disease classification but completely devoid of domain-specific information. While effective, these approaches still operate on the regular grid structure of convolutional feature maps, which may not be optimal for capturing the irregular, organic shapes of plant diseases.

2.3 Graph Analysis in Agricultural Imaging

In recent years, Graph Neural Networks have emerged as a powerful tool for analyzing non-Euclidean data structures, prompting researchers to explore their applicability in computer vision tasks where spatial relationships are critical. Unlike standard convolutional networks that process pixels in a fixed neighborhood, graph-based models operate on nodes and edges, allowing them to capture long-range dependencies and complex topological patterns [4]. In the context of agricultural imaging, an image can be transformed into a graph by segmenting it into visually homogeneous regions, known as superpixels. Each superpixel becomes a node in the graph, with its color and texture features serving as node attributes. Edges are established between adjacent or visually similar nodes, creating a network that represents the structural composition of the scene.

The application of graph analysis to crop disease detection offers several distinct theoretical advantages. First, by aggregating information across nodes, the network can contextualize localized symptoms within the broader structure of the leaf, helping to differentiate true lesions from isolated artifacts. Second, the graph structure explicitly models the boundaries between the foreground plant and the background environment, potentially improving robustness against clutter. Third, the topological features extracted by graph networks are inherently more invariant to geometric transformations, such as scaling and rotation, which are ubiquitous in hand-held smartphone photography [5]. Despite these theoretical benefits, the integration of graph analysis with transfer adaptation for the specific challenge of smallholder farm imagery remains an underexplored area of research, presenting a significant opportunity for methodological innovation.

3. Proposed Methodology

3.1 Data Acquisition and Preprocessing

The foundation of the proposed framework relies on the careful preparation of the image data to facilitate subsequent graph construction and domain adaptation. The input data comprises two distinct sets: a source domain consisting of high-resolution, controlled laboratory images with detailed disease annotations, and a target domain encompassing unconstrained, noisy images captured in smallholder farming environments. The initial preprocessing pipeline is designed to standardize these disparate inputs without obliterating the essential diagnostic features. All images are first subjected to a normalization process to ensure a consistent scale and color distribution. Given the extreme variations in illumination prevalent in the target domain, a localized histogram equalization technique is applied to enhance contrast and reveal subtle textural details that might otherwise be obscured by shadows or overexposure.

Following normalization, the images undergo a superpixel segmentation process, which is a critical prerequisite for the subsequent graph construction phase. The Simple Linear Iterative Clustering algorithm is employed to partition each image into a set of compact, perceptually uniform regions. This algorithm clusters pixels based on their color similarity and spatial proximity, effectively reducing the dimensionality of the image from hundreds of thousands of pixels to a few hundred superpixels [6]. The parameters of the segmentation algorithm are carefully tuned to strike a balance between preserving the fine-grained boundaries of small disease lesions and maintaining a manageable computational complexity for the graph processing stages. Each resulting superpixel serves as a foundational building block for the structural representation of the agricultural scene.




Figure 1: Comprehensive Schematic of the Transfer Adaptation Framework via Graph Analysis

3.2 Graph Construction and Feature Extraction

Once the superpixel segmentation is complete, the image is transformed into a topological graph. Let the graph be defined by a set of nodes and a set of edges. Each node corresponds to a specific superpixel generated in the previous step. The initial feature vector for each node is constructed by extracting a comprehensive set of descriptors from the constituent pixels of the corresponding superpixel. These descriptors encompass statistical measures of color distribution, such as mean and variance across different color channels, as well as textural features computed using local gradient aggregations. By encapsulating these low-level visual characteristics into the node attributes, the graph retains the essential diagnostic information necessary for disease identification.

The edges of the graph are established to model the spatial and feature-based relationships between the superpixels. An adjacency matrix is formulated where an edge is instantiated between two nodes if their corresponding superpixels are spatially contiguous in the original image. To capture non-local similarities, additional edges may be introduced based on the Euclidean distance between the feature vectors of the nodes, allowing the network to recognize disjointed disease symptoms that share visual characteristics [7].

Once the graph topology is established, a multi-layer Graph Convolutional Network is deployed to perform feature extraction. During each layer of processing, a node updates its feature representation by aggregating information from its immediate neighbors, guided by the established edge connections. This message-passing mechanism enables the network to learn progressively higher-level, context-aware representations, culminating in a global graph feature vector that encapsulates the structural essence of the entire image.

3.3 Transfer Adaptation Framework

The core innovation of the proposed methodology lies in the integration of the graph-based feature extraction process with an adversarial transfer adaptation framework. The objective is to ensure that the global graph features generated by the network are discriminative for the crop disease classification task while simultaneously being invariant to the domain origin of the input image. To achieve this, the architecture incorporates two distinct prediction heads that operate on the extracted graph features: a disease classifier and a domain discriminator. The disease classifier is trained exclusively on the labeled data from the source domain, optimizing a standard classification loss function to accurately categorize the different types of crop diseases based on their structural representations.

Simultaneously, the domain discriminator is trained to identify whether a given graph feature vector originated from a source domain image or a target domain image. The critical component of this adversarial setup is the incorporation of a gradient reversal mechanism during the backpropagation process [8]. As the network learns, the gradients derived from the domain discriminator's classification error are multiplied by a negative constant before being passed back to the graph feature extractor. This adversarial dynamic forces the graph convolutional layers to learn representations that actively confuse the domain discriminator. Consequently, the network converges on a state where the extracted structural features are highly relevant for diagnosing diseases but lack any domain-specific artifacts, thereby enabling robust prediction and detection in the challenging conditions of smallholder farm images.

4. Experimental Setup and Evaluation Metrics

4.1 Smallholder Farm Image Dataset

The empirical evaluation of the proposed methodology necessitates a robust dataset that accurately reflects the dichotomy between controlled and field conditions. The source domain utilizes a well-established, publicly available repository containing tens of thousands of leaf images captured under strictly controlled laboratory settings. These images represent a diverse array of crop species and disease classifications, providing a rich foundation for initial feature learning. The images in this source dataset are characterized by uniform backgrounds, consistent artificial lighting, and centralized positioning of the detached leaves, making them ideal for establishing baseline diagnostic patterns.

The target domain dataset, conversely, is constructed from images collected directly from smallholder agricultural plots across diverse geographic regions. This dataset is specifically curated to embody the severe visual challenges previously discussed. The field images contain high levels of background complexity, including soil, weeds, and human artifacts. Illumination varies drastically, ranging from direct, harsh sunlight causing specular highlights on the leaf surfaces, to deep, uneven shadows caused by canopy occlusion. Furthermore, the leaves are often photographed in situ, meaning they may be curled, overlapping, or partially obscured by stems and other plant structures [9]. The discrepancy between these two datasets provides a rigorous testbed for evaluating the efficacy of the transfer adaptation and graph analysis techniques.

Table 1: Domain Dataset Characteristics and Visual Complexity Metrics

Dataset Domain	Environment Type	Background Complexity	Illumination Variance	Total Image Samples
Source Domain	Controlled Laboratory	Minimal (Uniform)	Low (Consistent)	25,000
Target Domain	Smallholder Field	High (Cluttered)	Extreme (Variable)	8,500

4.2 Training Configuration and Metrics

The experimental configuration is designed to ensure a fair and comprehensive assessment of the model's capabilities. The graph-based feature extractor and the adversarial adaptation modules are trained end-to-end using an optimization algorithm that adaptively adjusts the learning rate based on the gradients. A batch

size is selected to balance computational efficiency with the stability of the adversarial training process. The learning rate schedules are carefully calibrated to prevent the domain discriminator from overpowering the feature extractor early in the training phase, which is a common failure mode in adversarial adaptation. Regularization techniques, including dropout applied to the fully connected layers and weight decay, are implemented to mitigate the risk of overfitting to the source domain data.

The evaluation of the system relies on standard classification metrics tailored for multi-class problems. The primary metric is overall accuracy, which provides a high-level view of the model's performance across all disease categories. However, given that disease datasets often exhibit class imbalances, precision, recall, and the harmonic mean of precision and recall are also calculated for each individual class [10]. Precision measures the proportion of positive identifications that are actually correct, while recall assesses the proportion of actual positives that were correctly identified. These metrics are crucial in the agricultural context, where the cost of a false negative missing a critical disease outbreak might be significantly higher than the cost of a false positive leading to unnecessary inspection.

5. Results and Discussion

5.1 Quantitative Performance Analysis

The empirical results demonstrate a substantial performance improvement when utilizing the proposed graph-based transfer adaptation framework compared to baseline methodologies. When traditional convolutional networks are trained on the source domain and evaluated directly on the target domain without any adaptation, the classification accuracy drops precipitously, validating the severity of the domain shift problem. The convolutional models struggle to separate the foreground disease symptoms from the chaotic background elements of the smallholder images. Even when standard domain adaptation techniques are applied to these grid-based convolutional features, the improvement is marginal, as the global statistical alignment fails to preserve the intricate structural details of the disease lesions.

In contrast, the proposed methodology, which leverages topological graph representations, exhibits remarkable resilience to the domain shift. By encoding the image as a network of superpixels, the model effectively isolates the structural morphology of the plant from the background noise. The adversarial alignment of these graph features ensures that the diagnostic criteria learned in the laboratory are successfully translated to the field environment. Quantitative comparisons reveal that the integrated approach significantly outperforms both the non-adapted graph models and the adapted convolutional models across all primary metrics [11]. The improvements are particularly pronounced in disease categories characterized by complex, irregular lesion shapes that are highly susceptible to visual distortion in field conditions.

Table 2: Performance Comparison of Detection Frameworks on Target Domain

Methodology Configuration	Overall Accuracy	Average Precision	Average Recall	F1-Score
Baseline Convolutional	62.4%	61.8%	63.1%	62.4%
Convolutional with Adaptation	71.5%	70.2%	71.9%	71.0%
Graph Analysis (No Adaptation)	68.3%	67.9%	68.5%	68.2%
Proposed Framework	84.7%	83.5%	85.1%	84.3%

5.2 Qualitative Analysis and Robustness

Beyond the aggregate quantitative metrics, a qualitative examination of the model's internal representations provides valuable insights into its operational mechanics. Visualizing the activation patterns of the graph convolutional layers reveals that the network successfully learns to prioritize nodes corresponding to diseased tissue while suppressing nodes associated with healthy tissue or background elements. The message-passing mechanism enables the model to connect disjointed symptomatic regions across the leaf surface, forming a holistic structural understanding of the infection pattern. This structural awareness is fundamentally missing in standard convolutional models, which tend to focus on highly localized texture patches that are easily confused by shadows or soil debris.




Figure 2: Visual Comparison of Node Activation Maps on Field Images

The robustness of the proposed framework is further evidenced by its performance under extreme occlusion and illumination variance. In scenarios where a significant portion of the diseased leaf is obscured by an overlapping healthy leaf, the graph model can still aggregate sufficient diagnostic information from the visible nodes to make an accurate prediction. Similarly, the localized normalization performed during the superpixel extraction phase, combined with the domain-invariant feature alignment, allows the model to process images with harsh shadows or intense glare without severe degradation in accuracy [12]. These qualitative observations confirm that transitioning from a rigid pixel grid to a flexible topological graph provides a more robust foundation for interpreting the chaotic visual environment of smallholder farming.

6. Conclusion and Future Directions

6.1 Summary of Contributions

This research has systematically addressed the critical challenge of predicting and detecting crop diseases in unconstrained smallholder farm images by proposing a novel integration of graph analysis and transfer adaptation. The investigation demonstrated that traditional convolutional approaches, while highly effective in controlled settings, lack the structural awareness necessary to navigate the severe domain shift encountered in real-world agricultural environments. By reformulating the image representation as a topological graph of interconnected superpixels, the proposed methodology successfully captured the complex spatial relationships and morphological characteristics of disease symptoms, effectively separating them from background clutter and illumination artifacts. Furthermore, the application of adversarial domain adaptation to these graph-based features facilitated the seamless transfer of diagnostic knowledge from pristine laboratory datasets to chaotic field imagery, resulting in a substantial and statistically significant improvement in classification accuracy and robustness.

6.2 Limitations and Future Work

Despite the promising results, the current framework exhibits certain limitations that warrant further investigation. The superpixel segmentation process, while effective for dimensionality reduction, introduces an additional computational overhead that may constrain real-time application on low-resource mobile devices commonly used by smallholder farmers. Future research must focus on optimizing the efficiency of the graph construction phase, potentially exploring dynamic or hierarchical segmentation algorithms that balance granularity with processing speed. Additionally, the current adversarial adaptation strategy assumes a closed-set scenario where the disease classes in the target domain exactly match those in the source domain. Subsequent studies should explore open-set domain adaptation techniques to handle the inevitable discovery of novel or uncharacterized pathogens in the field, ensuring that the detection systems remain adaptable and resilient in the face of evolving agricultural threats.

References

1. Aziz, A.; Mahmood, D.; Qureshi, M.S.; Qureshi, M.B.; Kim, K. AI-Based Peak Power Demand Forecasting Model Focusing on Economic and Climate Features. *Front. Energy Res.* 2024, 12, 1328891.
2. Wesoly, M.; Przewodowski, W.; Ciosek-Skibińska, P. Electronic noses and electronic tongues for the agricultural purposes. *TrAC Trends Anal. Chem.* 2023, 164, 117082.
3. Scillitoe, J.L. The role of alumni attachment in the university technology transfer process. In *Proceedings of the 2013 Proceedings of PICMET '13: Technology Management in the IT-Driven Services (PICMET)*, San Jose, CA, USA, 28 July–1 August 2013; pp. 2397–2406.
4. Maertens, A.; Barrett, C.B. Measuring Social Networks' Effects on Agricultural Technology Adoption. *Am. J. Agric. Econ.* 2013, 95, 353–359.
5. Vosooghizaj, M.; Taghipour, A.; Canel-Depitre, B. Supply chain coordination under information asymmetry: A review. *Int. J. Prod. Res.* 2020, 58, 1805–1834.
6. Regulation 2023/1543. Regulation (EU) 2023/1543 on European Production Orders and European Preservation Orders for Electronic Evidence in Criminal Proceedings and for the Execution of Custodial Sentences Following Criminal Proceedings. Available online: <https://eur-lex.europa.eu/eli/reg/2023/1543/oj/eng> (accessed on 3 October 2025).
7. Wang, J.; Zou, J.; Shi, Y.; Zeng, N.; Guo, D.; Wang, H.; Zhao, C.; Luan, F.; Zhang, X.; Sun, J. Traditional Chinese medicine and mitophagy: A novel approach for cardiovascular disease management. *Phytomedicine* 2024, 128, 155472.
8. Ivanov, D.; Dolgui, A. Viability of intertwined supply networks: Extending the supply chain resilience angles towards survivability. A position paper motivated by COVID-19 outbreak. *Int. J. Prod. Res.* 2020, 58, 2904–2915.
9. Fuller, A.; Fan, Z.; Day, C.; Barlow, C. Digital twin: Enabling technologies, challenges, and open research. *IEEE Access* 2020, 8, 108952–108971.
10. Shah, N.P.; Levin, D.Z.; Cross, R. Secondhand social capital: Boundary spanning, secondhand closure, and individual performance. *Soc. Netw.* 2018, 52, 18–27.
11. Belhadi, A.; Kamble, S.; Jabbour, C.J.C.; Gunasekaran, A.; Ndubisi, N.O.; Venkatesh, M. Manufacturing and service supply chain resilience to the COVID-19 outbreak: Lessons learned from the automobile and airline industries. *Technol. Forecast. Soc. Change* 2021, 163, 120447.
12. Vignes, V.M.; Harini, S.M.P.; Satheesh, R.; Das, V.; Padmanaban, S. AI-Driven Cybersecurity Framework for Anomaly Detection in Power Systems. *Sci. Rep.* 2025, 15, 35506.