Multi-Domain Few-Shot Learning and Dataset for Agricultural Applications

Sai Vidyaranya Nuthalapati*
Sensyne Health PLC
vidyaranya.ns@gmail.com

Anirudh Tunga*
Purdue University
atunga@purdue.edu

Abstract

Automatic classification of pests and plants (both healthy and diseased) is of paramount importance in agriculture to improve yield. Conventional deep learning models based on convolutional neural networks require thousands of labeled examples per category. In this work we propose a method to learn from a few samples to automatically classify different pests, plants, and their diseases, using Few-Shot Learning (FSL). We learn a feature extractor to generate embeddings and then update the embeddings using Transformers. Using Mahalanobis distance, a class-covariance-based metric, we then calculate the similarity of the transformed embeddings with the embedding of the image to be classified. Using our proposed architecture, we conduct extensive experiments on multiple datasets showing the effectiveness of our proposed model. We conduct 42 experiments in total to comprehensively analyze the model and it achieves up to 14% and 24% performance gains on few-shot image classification benchmarks on two datasets.

We also compile a new FSL dataset containing images of healthy and diseased plants taken in real-world settings. Using our proposed architecture which has been shown to outperform several existing FSL architectures in agriculture, we provide strong baselines on our newly proposed dataset.

1. Introduction

In agriculture, correct classification of different pests and plants (both healthy and diseased) is a major issue due to the high similarity and shared characteristics between different species. Automatic classification of different categories using deep learning is an area of active research [28, 22, 6, 33]. However, in practice, people still rely on manual classification by experts for classification of different species. This can be partly attributed to the fact that using traditional deep learning networks based on Convolutional Neural Networks (CNNs) for classification require thousands of labelled ex-

 amples per target category for training [16] and labeling samples on such a large scale requires domain experts. Further, the number of categories a trained CNN-based model can recognize remains fixed after training. To expand the set of categories that the network can recognize, it has to be further trained with new samples from novel classes, a process called fine-tuning [4]. Moreover, during training it is important that there is enough data (thousands per class) to prevent the network from overfitting [9]. Humans, on the other hand learn new tasks with very little supervision - a child can generalize the concept of "horse" from a single picture. Also, humans can generalize to recognize novel categories from only one or a few examples [17]. To enable networks to learn from a few examples, recently a lot of focus has been placed on Few-Shot Learning (FSL). FSL aims to tackle the problem of classification with very few training examples, and it is becoming popular in many fields [5, 19, 27, 34]. Specifically, in FSL, we are given two sets of labelled image data: *meta-train* and *meta-test* such that the image classes in both sets are mutually exclusive. The aim is to use the data in the meta-train set and learn transferable knowledge (also called meta-training) to construct a classifier on the visual classes in meta-test set which can classify a given *query* sample even with very few labeled (*support*) examples. Further, if the domain of visual classes in the meta-train set is different from that of meta-test set it is called cross-domain FSL. Feature distribution discrepancies across domains is one of the main challenges in cross-domain FSL. Similarly, in mixed-domain FSL, both the meta-train and meta-test sets contain classes from multiple domains and single-domain FSL contains instances from a single domain.

Recent FSL approaches in agriculture [20, 21] learn an embedding function using the meta-train set which is then used to generate embeddings of the samples in the meta-test set. Finally, for a given query sample, a distance metric applied on the embeddings of the support samples gives the final prediction. A major limitation of this approach is the assumption that the transferable knowledge learned using meta-train classes can be directly applied to classification tasks generated using meta-test classes [35]. For

---

*equal contribution; order determined by a coin toss

example, this approach assumes that the discriminative features for differentiating two plant species is same as the discriminative features required for differentiating two pest species. This problem becomes even more pronounced in cross-domain settings. In this work, we adopt an adaptation method that modifies the representations derived from the embedding function. The modified representations are tailored to maximize the discriminative power of the visual representations for a task at hand. Following [35], we use a Transformer [31] as a set-to-set function approximator that adapts the generated embeddings to the current task.

Secondly, the existing few-shot architectures in agriculture [1, 21] focus on using the Euclidean distance to compute the similarity of instance representations. However, using the Euclidean distance assumes that the feature dimensions are un-correlated and the feature dimensions have uniform variance [2], which do not always hold. Hence, inspired by [2], we use Mahalanobis distance [8], a class-covariance-based metric, to compute the similarity of the embeddings.

Finally, the existing few-shot plant datasets for agriculture [21, 13] contain images of a single leaf in laboratory settings as shown in Fig 3. However, in real-world applications of agriculture, we rarely get such images of a single leaf for classification. For instance, an image captured by a person using the ubiquitous smartphone camera is very different and contains challenging scenarios of lighting, orientation and background. As a result, it is imperative to design and test machine learning frameworks that perform well even with images containing the entire plants and not just a single leaf under varied conditions. To address this issue, we collected samples from publicly available resources (e.g., https://edenlibrary.ai/home) and compile a dataset of healthy and diseased plants. We also provide a strong baseline for our new few-shot dataset.

To test our architecture, we conduct extensive FSL experiments on multiple datasets, namely, the Plant and Pests (PP) dataset [21] and the PlantVillage dataset [13]. The former is a dataset of plants and pests and enables us to experiment with mixed and cross-domain settings. We improve the state-of-the-art accuracy of this dataset significantly, both in mixed-domain and cross-domain settings. On the other hand, the PlantVillage dataset consists of images of healthy and diseased plants. Our model outperforms the current best-performing models on this dataset by a significant margin.

The contributions of this work are three-fold: 1) We propose a new architecture for FSL using Transformers for enhancing the feature representations and class-covariance-based distance metric for calculating the feature distance. 2) We conduct extensive experiments under different settings on multiple datasets to establish the superiority of the model. Various settings include single-domain, mixed-

domain, and cross-domain (including cross-dataset). 3) We collect samples from publicly available resources to create a new FSL dataset containing images of healthy and diseased plants which represent the real-world applications of computer vision in agriculture. We also provide strong baselines on the proposed dataset.

2. Related Work

In the field of agriculture, many recent works on vision have tried to solve the classification task with limited training samples. In [12], the authors used DC-GAN [25] to generate augmented images for training. Another approach proposed in [20] is based on prototypical networks [27] and trains a CNN feature extractor followed by Euclidean distance calculation. The framework is trained using a triplet loss function. Another recent work [21], used FSL to train a model with limited samples. In [1], the authors trained a CNN to extract general plant leaf characteristics and used Siamese networks combined with triplet loss for classification. In [1], the authors tested their work using the PlantVillage [13] dataset. In [21], the authors used a part of the PlantVillage [13] dataset.

The existing FSL datasets [13, 21, 1] of plants contain images of plants in an ideal setting with contrasting backgrounds, single leaves, no occlusions, and constant lighting conditions. To facilitate development of robust computer vision approaches in agriculture, we introduce *Plants in Wild*, a new FSL dataset with images of diseased and healthy leaves in real-life setting.

In general, there are two main components of FSL: 1) learning generalizable instance embeddings and 2) distance computation between support and query instances to classify the input images. In many works, generalizable instance embeddings [11, 23, 32, 29, 3] are learned and they are used for further classification using simple classifiers like nearest-neighbor methods and linear classifiers. In [32], the authors use a nearest neighbor approach. MetaOptNet [18] uses a linear classifier. Siamese networks [15] use a shared feature extractor and classification is done using the smallest L1 distance between the query sample and the support samples. Our work improves on the current models by altering both the components of FSL stated above. Following [35], we change the way instance embeddings are learnt by enhancing the embeddings using a transformer [31] based set-to-set function approximator. Transformers have been effective to contextualize the representations of inputs and have found applications in diverse fields [14, 30]. Secondly, to compute the similarity between query and support examples, we use class-covariance based deterministic metric - Mahalanobis distance [8, 2], which takes into account the distribution in feature space of each class, resulting in improved non-linear classifier decision boundaries.
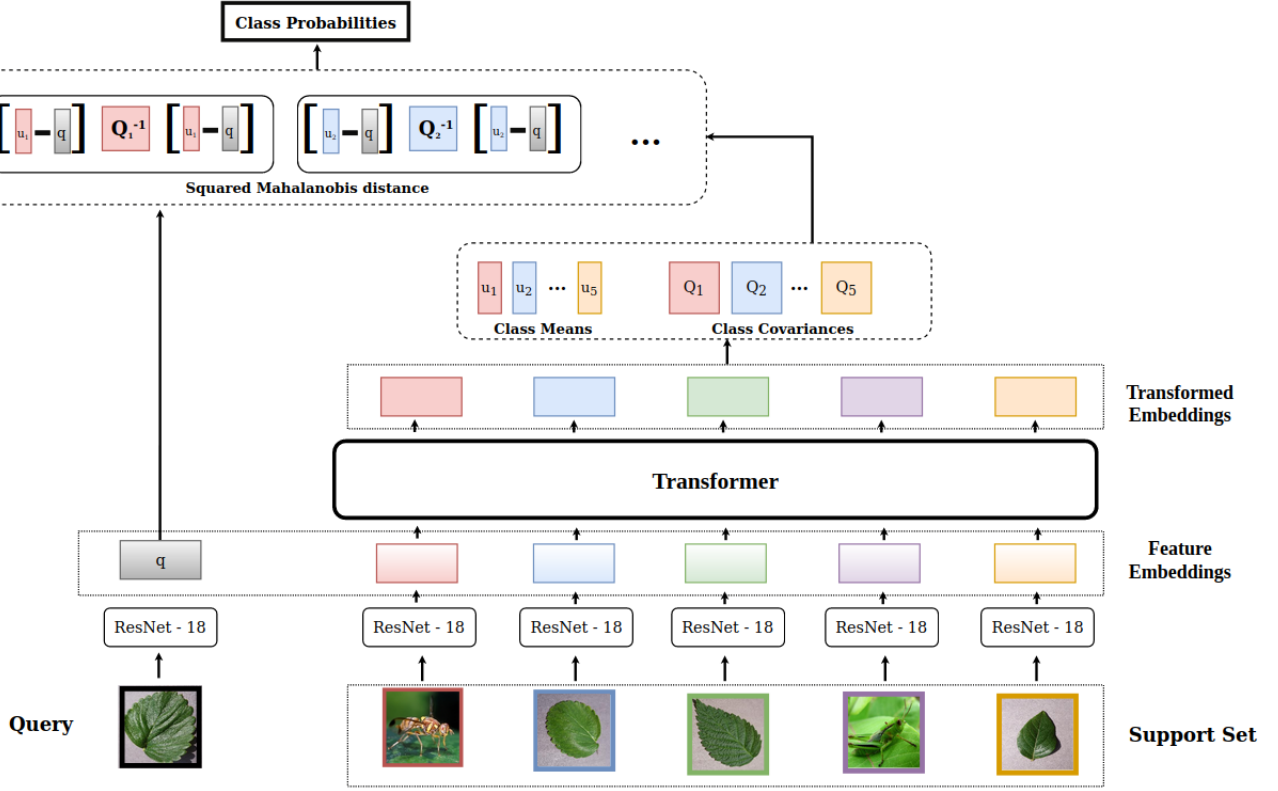
Figure 1. Illustration of the proposed architecture. We adapt the extracted features of the support set for each task using a Transformer before calculating the class-covariance based distance between support set features and query features.

3. Problem formulation

Assume that we are given two labeled sets $\mathbb{D}_{meta\_train}$ and $\mathbb{D}_{meta\_test}$, such that the labels of the two sets are mutually exclusive i.e. $y_{meta\_train} \cap y_{meta\_test} = \phi$. In both the datasets y_i is a one-hot encoded vector denoting the label of the corresponding sample.

In the FSL paradigm, we are given a set of episodes $\{T^i\}$. Each episode T^i is made up of a support set and a query set. We denote support set by $\mathbb{D}_{support}$ and the examples in the support set by (x_s, y_s). Similarly, we denote the examples in a query set ($\mathbb{D}_{query}$) by (x_q, y_q). We also note that some existing works refer to support set, query set, and episode as training set, test set, and task respectively. In this work, we use these terms interchangeably. Each episode is represented as an M-shot N-way classification problem, where N classes are randomly selected from a set of classes y_{meta_test} with M support examples for each of the N classes. The goal is to learn a classifier f such that given a query sample x_q and the support set $\mathbb{D}_{support}$, the classifier is able to predict y_q. In a few-shot scenario the value of M is very small (1, 5, 10).

In order to learn the parameters of the classifier f we use a larger dataset called a meta-training dataset denoted by $\mathbb{D}_{meta\_train}$ to sample multiple M-shot N-way episodes [32, 35, 7]. In each of the episodes, the classifier labels the input $x_q^{meta\_train}$ as one of the $N \in y_{meta\_train}$ classes. During the learning process, we minimize the average loss value of all the sampled tasks. To evaluate the performance of the classifier, we followed similar steps on $\mathbb{D}_{meta\_test}$.

4. Architecture

The classifier f consists of two main components: an *embedding function* and a *distance calculation*. Next, we discuss both of these components in detail.

4.1. Embedding function

The *embedding function* is trainable and consists of two parts: feature extraction ϕ_x and transformation ψ_x. The feature extraction step ϕ_x extracts the features and projects them to a d-dimensional space. The feature extractor architecture consists of a ResNet-18 [10] network pre-trained on Imagenet [26]. The extracted features, however, are not ideal and do not capture important discriminative visual features for a specific episode[35]. In other words, simply using an embedding generated by ϕ_x does not incorporate any information about other support samples in the current
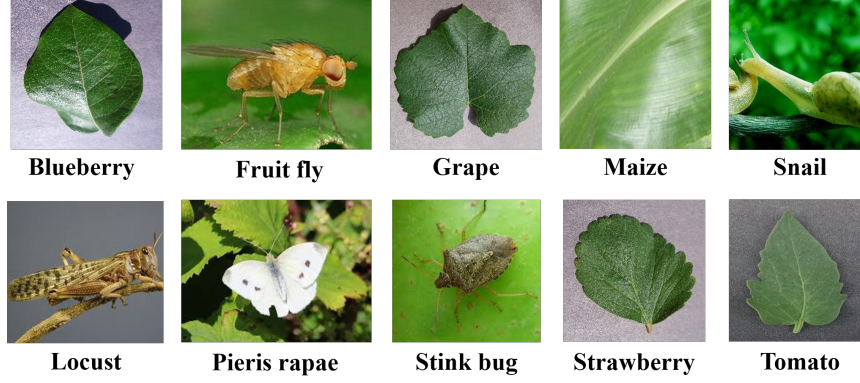
Figure 2. An example of meta-train set (top row) and meta-test set (bottom row) for mixed domain classification on PP dataset [21].


Figure 3. An example of set containing samples only from plant classes from PP dataset [21].

episode. As a result, irrespective of the cohort, the embedding of a given image is always the same. Having such deterministic features is detrimental for discriminating the elements of the set. To overcome this, the embeddings need to be contextualized. Inspired by FEAT algorithm proposed in [35], the outputs of ϕ_x are transformed using a set-to-set function (ψ_x), which enriches the embedding of each image by considering those of all other images in the support set. By enabling interaction between embeddings of various images in the set it leads to richer representations and discriminative features. Another desired feature of the set-to-set function is that it should be permutation-invariant. Following [35], we implement ψ_x using *Transformer* architecture [31], which fit the criteria of the required set-to-set function. Transformers rely on the self-attention mechanism which updates the representation of every input by weighing the relevance of all the other inputs. They map a query and a set of key-value pairs to an output.

$$\psi_{x_i} = \phi_{x_i} + \left(\sum_{\forall j} V(s_j) f(\phi_{x_i}, \phi_{x_j}) \right), \quad (1)$$

where $f(\phi_{x_i}, \phi_{x_j})$ is used to measure the similarity between ϕ_{x_i}, ϕ_{x_j}.

$$f(\phi_{x_i}, \phi_{x_j}) = softmax_j \left(\frac{Q(\phi_{x_i})^T K(\phi_{x_j})}{\sqrt{d}} \right), \quad (2)$$

where the functions Q and K are learned linear projec-

tions. Combined with V, which is also a learned linear projection, the functions Q and K project the inputs to a common representation space before applying the similarity measure.

The resultant transformed embeddings (ψ_x) act as inputs to the distance computation module.

4.2. Distance computation

In this step, we calculate the distance between the embeddings of the support samples with the query sample to finally classify the query. Following [2], we use Mahalanobis distance [8], a class-covariance based distance metric estimated at test-time, which has been shown to be better than using other metrics such as Euclidean or cosine. The final prediction is calculated as follows:

$$p(y_q = n | \psi_{x_q}, S^t) = \text{softmax}(-d_n(\psi_{x_q}, \mu_n)), \quad (3)$$

where $n \in N$ is one of the support classes in the current episode/task $t \in T$, S^t is the support set for the current episode/task, μ_n is the mean transformed embedding (ψ_x) of all the M samples corresponding to the class n and d_n is the squared Mahalanobis distance defined as below:

$$d_n(x, y) = \frac{1}{2}(x - y)^T (Q_n^t)^{-1}(x - y), \quad (4)$$

where Q_n^t is the covariance matrix of class $n \in N$ for episode/task $t \in T$. We calculate the covariance matrix Q_n^t using a regularized estimator

| Cydia pomonella | Gryllotalpa | Leaf hopper | Spodoptera litura | Weevil |

Figure 4. An example of set containing samples only from pest classes from PP dataset [21].

$$Q_n^t = \lambda_k^t \Sigma_n^t + (1 - \lambda_k^t)\Sigma^t + \beta I, \tag{5}$$

where Σ^t is the covariance matrix of all the classes in the task t, Σ_n^t is the covariance matrix of the class n in task t and λ is a weigthing factor defined as follows:

$$\lambda_k^t = \frac{|S_k^t|}{|S_k^t| + 1}, \tag{6}$$

where $|S_k^t|$ is the number of elements in the support set of the task t corresponding to class k. For a theoretical explanation of the superiority of Mahalanobis distance to other metrics, we refer the reader to [2].

4.3. Loss function

To ensure that the transformation function ψ_x pulls the embeddings of same-class instances closer and drives different-class instances farther, we use a contrastive loss function. To achieve this, the transformation function is applied to instances of each of the N classes present in the given episode/task, giving a transformed embedding ψ_x' and the mean class centers $\{c_n\}_{n=1}^N$. We then calculate the similarity of individual embedding to the class center corresponding to the individual. Apart from this we also use a standard cross-entropy loss function to calcualte the loss using the final prediction. Hence, the complete loss function is as follows:

$$\begin{aligned}\mathcal{L}(\hat{y}_q, y_q) &= l(\hat{y}_q, y_q) \\ &+ \lambda \times l(\text{softmax}(\text{sim}(\phi_x, c_n)), y_q),\end{aligned} \tag{7}$$

where sim is the similarity function calculated using the class-covariance-based distance metric, l is the standard cross entropy loss function and λ is the weighting factor.

4.4. Implementation Details

The proposed model has been implemented using PyTorch [24]. We use stochastic gradient descent with Nesterov acceleration with an initial learning rate of 0.0002, weight decay of 5e-4, momentum of 0.9, and a learning rate scheduler with step size of 40 and gamma of 0.5 to optimize the model. We resize all the images to 84×84×3 before using the feature extractor. The value of λ is set to 0.1.

Table 1. Various classes in our newly proposed *Plants in Wild* dataset

Healthy	Diseased
Celery	Corn Leaf Blight
Chinese Cabbage	Chinese Cabbage Fusarium
Cotton	Corn Rust Leaf
Grapevine	Grapevine Esca
Potato	Potato Early Blight
Red Cabbage	Cucumber Tetranychus
Tomato	Tomato Fruit Virus
Watermelon	Cucumber Thrips
Zucchini	Grapevine Powdery Mildew
Bell Pepper	Bell Pepper Leaf Spot

5. Experiments

In this section, we provide an overview of the datasets used, describe the experimental setup and provide quantitative results.

5.1. Datasets

We test our proposed architecture on two different challenging datasets and also provide baselines for the new dataset that we propose. We next describe the two external datasets that we use.

Plant and Pest (PP) [21] dataset contains 6000 images of both pests and plants. We use this dataset to test the mixed-domain and cross-domain performance of our proposed architecture. The dataset contains 20 classes, 10 each for plants and pests. During experimentation, the dataset is split into two: meta-train set and meta-test set. The labels of both sets do not overlap and the distribution of the sets may also differ. In mixed domain settings, both meta-train and meta-test sets contain instances of plants and pests. We show an example of this in Fig 2. On the other hand, there are two cross-domain settings. In the cross-domain 1 setting, the meta-train set contains examples of pests and the meta-test set is made up of plant instances. For example, Fig 4 can make up the meta-train set and the meta-test set contains images from Fig 3. Finally, the cross-domain 2 setting, refers to the meta-train set containing plants and the meta-test set containing pest instances. In this setting, the

Figure 5. Examples of healthy and diseased tomato leaves from PlantVillage dataset [13].

Figure 6. A few examples of real world images from our Plants in Wild (PiW) dataset containing both healthy and diseased plants.

images in Fig 3 make up the meta-train set, whereas instances in Fig 4 constitute the meta-test set. While being challenging, the cross-domain setting is vital owing to its practical implications.

PlantVillage [13] is a public dataset containing 38 classes of healthy and diseased leaves of different crops. The dataset consists of 61,486 images. We split the dataset into 3 different and mutually exclusive meta-training and meta-testing sets, in the same manner followed in [1]. In split 1, the meta-test set consists of 10 different classes of Tomato (1 healthy, and 9 diseased) as shown in Fig 5, and the meta-training set consists of the rest of the 28 classes. In split 2, the meta-test set consists of 4 classes of Apple (3 diseased, 1 healthy): apple scab, black rot, cedar apple rust, and healthy; 4 classes of Grape (3 diseased, 1 healthy): black rot, esca, leaf blight, and healthy; 2 classes of Cherry

(1 diseased, 1 healthy): powdery mildew, and healthy; and the meta-train consists of rest of the 28 classes. Similarly, in split 3, the meta-test set consists of 4 classes of Corn (1 healthy, 3 diseased), 4 classes of Grape (1 healthy, 3 diseased), and 2 classes of peach (1 healthy, 1 diseased).

Plants in Wild (PiW): In addition to these two datasets, we curated another dataset to validate our model's performance across domains and under real-life settings. The images of leaves in the Plantvillage and PP dataset are in a laboratory setting such that each image contains only one leaf in a contrasting background. However, in a real-life setting, the images are less likely to be taken in a controlled setting and will be of varying backgrounds, lighting conditions, and angles. We used the publicly available images available on the internet (for instance, the Eden Li-

Table 2. Average accuracy achieved on different setting of N and M on mixed domain on PP dataset [21].

	K = 1		K = 5		K = 10	
N_{train}, N_{test}	Li et al. [21]	Ours	Li et al. [21]	Ours	Li et al. [21]	Ours
3, 3	0.811	**0.841**	0.870	**0.900**	0.904	**0.929**
3, 5	0.677	**0.723**	0.794	**0.839**	0.821	**0.873**
5, 3	0.790	**0.828**	0.864	**0.912**	0.889	**0.924**
5, 5	0.676	**0.739**	0.780	**0.837**	0.825	**0.872**

Table 3. Average accuracy achieved on different settings of N and M on cross-domain 1 on PP dataset [21].

	M = 1		M = 5		M = 10	
N_{train}, N_{test}	Li et al. [21]	Ours	Li et al. [21]	Ours	Li et al. [21]	Ours
3, 3	0.697	**0.792**	0.849	**0.937**	0.865	**0.955**
3, 5	0.614	**0.705**	0.774	**0.894**	0.812	**0.930**
5, 3	0.721	**0.759**	0.814	**0.923**	0.871	**0.936**
5, 5	0.531	**0.677**	0.760	**0.884**	0.810	**0.902**

Table 4. Average accuracy achieved on different setting of N and M on cross-domain 2 on PP dataset [21]

	M = 1		M = 5		M = 10	
N_{train}, N_{test}	Li et al. [21]	Ours	Li et al. [21]	Ours	Li et al. [21]	Ours
3, 3	0.441	**0.518**	0.526	**0.660**	0.548	**0.716**
3, 5	0.290	**0.381**	0.380	**0.547**	0.427	**0.594**
5, 3	0.425	**0.520**	0.531	**0.665**	0.558	**0.708**
5, 5	0.292	**0.385**	0.374	**0.545**	0.439	**0.603**

brary website[1]) to collect images of diseased and healthy plants. All the images were taken using smartphone cameras in different fields and farms. We preprocessed all the images by resizing and center-cropping the images. Fig 6 shows some of the classes from our dataset. Comparing our dataset with the existing datasets, shown in Figures 3 and 5, highlights the differences between the datasets - lab controlled and real-life images.

We collected 20 categories of plants: 10 of which are healthy and 10 classes correspond to diseased plants. The dataset consists of 1980 images, which were taken by users using different smartphone cameras. Table 1 shows all the available classes in the dataset. We evaluate the model using three different splits of the dataset. In split 1, the meta-train set consists of Bell pepper healthy, Bell pepper leaf spot, Celery healthy, Chinese cabbage healthy, Chinese cabbage fusarium, Corn leaf blight, Corn rust leaf, Cucumber tetranychus, Cucumber thrips, and Red cabbage healthy. The meta-test set consists of the rest of the 10 classes. In split 2, the meta-test and meta-train sets are reversed. In split 3, the meta-train set consists of all the healthy plant leaf images, while the meta-test set consists of all the diseased plant leaf images.

5.2. Results

We evaluate the performance of our proposed model on the PP dataset [21] and the PlantVillage dataset [13] to establish the superiority of our model compared to existing best-performing methods on these datasets. We then use our model and provide baselines for our newly proposed dataset. To evaluate the model, we generate 600 support-and-query sets from the meta-test set. To report the final accuracy, we use the average accuracy on all the 600 sets.

Plant and Pest (PP): On the *PP dataset* [13], we perform 36 sets of experiments, under three different settings: 2 cross-domain and mixed-domain following [21]. Please refer to the 5.1 for the details on different settings. For each of the domains we evaluate with different values of *M* and *N*. We also vary the value of *N* during the meta-train and meta-test phase, which we represent using N_{meta_train} and N_{meta_test}. For example, a value of 3, 5 corresponding to N_{meta_train} and N_{meta_test} respectively means that the number of support classes in each episode is 3 during meta-training and 5 during the model evaluation (or the meta-testing) phase. We also experiment with 3 different values of M (1, 3, and 5). We observe significant performance gains across all the settings. Specifically, we observe around 3 to 6% gains in performance on the mixed-domain setting, approximately 7 to 11% improvement on cross-domain 1,

Table 5. Average accuracy achieved on different settings on PlantVillage dataset [13]

	Split 1		Split 2		Split 3	
Method used	5 way 1 shot	5 way 5 shot	5 way 1 shot	5 way 5 shot	5 way 1 shot	5 way 5 shot
Argueso et al. [1]	0.34	0.531	0.464	0.769	0.552	0.693
Ours - Cross Dataset	**0.38**	**0.54**	**0.573**	**0.772**	**0.716**	**0.861**
Ours	**0.466**	**0.635**	**0.709**	**0.87**	**0.754**	**0.885**

and about 12 to 14 % improvement on cross-domain 2 setting. We observe that gains on cross-domain settings are much higher compared to the mixed domain setting. This shows that our architecture is much better at extracting discriminative features even on unseen domains. We also observe that the accuracy improvements when using a higher number of support classes during meta-testing are greater i.e. when using a value of 5 for N_{meta_test} compared to 3. Specifically, the average improvement across all the settings with a N_{meta_test} value of 5 is 9% while it is equal to 5% when N_{meta_test} is set to 3. This is due to the fact that, with lower values of N_{meta_test} even a random classifier can correctly predict the outcome with high probability (given just 3 three classes, a random classifier is correct 33% of the time). Increasing the value demands robust architectures for accurate predictions. In summary, we see the trend of the increasing performance gap with increasing difficulty of the task throughout our experiments: performance improvements on cross-domain settings are more pronounced compared to single-domain and the accuracy gap widens when we increase the number of support classes in meta-test set.

PlantVillage: On the PlantVillage dataset [13], we perform experiments under 6 different settings. We use three different splits (please refer to the Section 5.1 for details on splits) and use two different values of M (1, 5) while keeping N constant. Please refer to the row corresponding to *Ours* in Table 5 for the results. Experiments using our model outperform the current state-of-the-art architecture [1] on this dataset by 10 to 24%. Moreover, the performance gains are much higher in 1-shot settings with a mean improvement of 18% compared to an average improvement of 11% in 5-shot settings. This is in line with our observations on the PP dataset [21] that the performance gains are much higher as the tasks get tougher.

Plants in Wild (PiW): The PiW dataset, in contrast to the previous two datasets, consists of diseased and healthy images of plants taken in a natural setting. We perform six experiments using three different data splits (Please refer to the 5.1 for the details on different splits) to validate the performance of our model in a natural setting. As shown in Table 6, our model performs the same, if not better when compared to its performance on plant images in a controlled setting. These experiments show the effectiveness of our model even with real-life images taken with smartphone cameras.

Cross Dataset Performance: To further establish the

superiority of our model in adapting to unseen domains, we only use the images of pests from the PP dataset [21] as our meta-train set and use plant images in the PlantVillage dataset [13] as meta-test set. We show that even under such a stringent *cross-domain / cross-dataset* setting, our algorithm still shows notable improvements (up to 17% gains). We record these values in the *Ours - Cross Dataset* row in Table 5.

In summary, results across multiple settings on various datasets indicate that using Transformers for generating episode-specific rich instance embeddings (ψ_x) provide for a useful calculation of episode and class specific covariance matrix (Q_n^t) in few-shot settings.

Table 6. Average accuracy on different settings on our PiW dataset.

Splits	N,M	**Ours**
Split 1	5 way 1 shot	0.768
	5 way 5 shot	0.905
Split 2	5 way 1 shot	0.714
	5 way 5 shot	0.846
Split 3	5 way 1 shot	0.767
	5 way 5 shot	0.933

6. Conclusion

In this paper, we address the problem of few-shot learning in agriculture. Most of the current works in the field use instance embeddings that are not tailored to the task at hand. Using a Transformer based architecture [31, 35], we enrich the embeddings of the images by considering all the images in the support set of the given episode. We further incorporate a class-covariance-based deterministic metric - Mahalanobis distance [8, 2] to calculate the similarity of the query vector with the candidates of the support set. We show that our architecture significantly improves the performances on two different datasets (alongside one cross-dataset experiment) under multiple settings.

We also provide a new dataset (*Plants in Wild*) mimicking the real-world scenarios. Our newly compiled dataset consists of plants that are both healthy and diseased. Using our model which has been shown to outperform several existing FSL architectures on agriculture datasets, we provide strong baselines on our new dataset. We hope that our new dataset along with our new models opens up new research directions in the field of agriculture.

References

[1] David Argüeso, Artzai Picon, Unai Irusta, Alfonso Medela, Miguel G San-Emeterio, Arantza Bereciartua, and Aitor Alvarez-Gila. Few-shot learning approach for plant disease classification using images taken in the field. *Computers and Electronics in Agriculture*, 175:105542, 2020. 2, 6, 8

[2] Peyman Bateni, Raghav Goyal, Vaden Masrani, Frank Wood, and Leonid Sigal. Improved few-shot visual classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14493–14502, 2020. 2, 4, 5, 8

[3] Soravit Changpinyo, Wei-Lun Chao, and Fei Sha. Predicting visual exemplars of unseen classes for zero-shot learning. In *Proceedings of the IEEE international conference on computer vision*, pages 3476–3485, 2017. 2

[4] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. *arXiv preprint arXiv:1904.04232*, 2019. 1

[5] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *International Conference on Learning Representations*, 2019. 1

[6] Konstantinos P Ferentinos. Deep learning models for plant disease detection and diagnosis. *Computers and Electronics in Agriculture*, 145:311–318, 2018. 1

[7] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017. 3

[8] Pedro Galeano, Esdras Joseph, and Rosa E Lillo. The mahalanobis distance for functional data with applications to classification. *Technometrics*, 57(2):281–291, 2015. 2, 4, 8

[9] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4367–4375, 2018. 1

[10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3

[11] Kyle Hsu, Sergey Levine, and Chelsea Finn. Unsupervised learning via meta-learning. *arXiv preprint arXiv:1810.02334*, 2018. 2

[12] Gensheng Hu, Haoyu Wu, Yan Zhang, and Mingzhu Wan. A low shot learning method for tea leaf's disease identification. *Computers and Electronics in Agriculture*, 163:104852, 2019. 2

[13] David Hughes, Marcel Salathé, et al. An open access repository of images on plant health to enable the development of mobile disease diagnostics. *arXiv preprint arXiv:1511.08060*, 2015. 2, 6, 7, 8

[14] Yash Kant, Dhruv Batra, Peter Anderson, Alex Schwing, Devi Parikh, Jiasen Lu, and Harsh Agrawal. Spatially aware multimodal transformers for textvqa. *arXiv preprint arXiv:2007.12146*, 2020. 2

[15] Gregory Koch, Richard Zemel, Ruslan Salakhutdinov, et al. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2. Lille, 2015. 2

[16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. 1

[17] Brenden Lake, Ruslan Salakhutdinov, Jason Gross, and Joshua Tenenbaum. One shot learning of simple visual concepts. In *Proceedings of the annual meeting of the cognitive science society*, volume 33, 2011. 1

[18] Kwonjoon Lee, Subhransu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10657–10665, 2019. 2

[19] Fei-Fei Li, Rob Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE transactions on pattern analysis and machine intelligence*, 28(4):594–611, 2006. 1

[20] Yang Li and Jiachen Yang. Few-shot cotton pest recognition and terminal realization. *Computers and Electronics in Agriculture*, 169:105240, 2020. 1, 2

[21] Yang Li and Jiachen Yang. Meta-learning baselines and database for few-shot classification in agriculture. *Computers and Electronics in Agriculture*, 182:106055, 2021. 1, 2, 4, 5, 7, 8

[22] Yang Lu, Shujuan Yi, Nianyin Zeng, Yurong Liu, and Yong Zhang. Identification of rice diseases using deep convolutional neural networks. *Neurocomputing*, 267:378–384, 2017. 1

[23] Luke Metz, Niru Maheswaranathan, Brian Cheung, and Jascha Sohl-Dickstein. Meta-learning update rules for unsupervised representation learning. *arXiv preprint arXiv:1804.00222*, 2018. 2

[24] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. pages 8026–8037, 2019. 5

[25] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 2

[26] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 3

[27] Jake Snell, Kevin Swersky, and Richard S Zemel. Prototypical networks for few-shot learning. *arXiv preprint arXiv:1703.05175*, 2017. 1, 2

[28] Edna C Too, Yujian Li, Pius Kwao, Sam Njuki, Mugendi E Mosomi, and Julius Kibet. Deep pruned nets for efficient image-based plants disease classification. *Journal of Intelligent & Fuzzy Systems*, 37(3):4003–4019, 2019. 1

[29] Eleni Triantafillou, Richard Zemel, and Raquel Urtasun. Few-shot learning through an information retrieval lens. *arXiv preprint arXiv:1707.02610*, 2017. 2

[30] Anirudh Tunga, Sai Vidyaranya Nuthalapati, and Juan Wachs. Pose-based sign language recognition using gcn and bert. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 31–40, 2020. 2

[31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017. 2, 4, 8

[32] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29:3630–3638, 2016. 2, 3

[33] Chengjun Xie, Rujing Wang, Jie Zhang, Peng Chen, Wei Dong, Rui Li, Tianjiao Chen, and Hongbo Chen. Multi-level learning features for automatic classification of field crop pests. *Computers and Electronics in Agriculture*, 152:233–241, 2018. 1

[34] Leiming Yan, Yuhui Zheng, and Jie Cao. Few-shot learning for short text classification. *Multimedia Tools and Applications*, 77(22):29799–29810, 2018. 1

[35] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8808–8817, 2020. 1, 2, 3, 4, 8