Journal of Computing and Electronic Information Management
ISSN: 2413-1660 | Vol. 11, No. 3, 2023

Enhanced few-shot learning for plant leaf diseases recognition

Wei Wu

Taishan University, Tai'an 271000, China

Abstract: With the breakthrough progress of deep learning technology in multiple fields, its application in specialized areas such as plant leaf disease recognition is constrained by the cost of data annotation and the lack of sample diversity. This study proposes an enhanced few-shot learning method that integrates self-supervised learning and semi-supervised learning to improve the model's generalization ability in plant leaf disease recognition tasks. Through self-supervised pre-training and semi-supervised fine-tuning, the model can effectively utilize limited annotated data and expand the training set by generating high-quality pseudo-labels. Experimental results show that this method significantly improves the model's recognition performance on unseen categories. Future research will explore more self-supervised tasks and complex pseudo-label generation algorithms to further enhance the model's accuracy and robustness, promoting the application of few-shot learning technology in the field of agriculture.

Keywords: Plant leaf diseases recognition; Few-shot learning; Self-supervised learning; Semi-supervised learning.

1. Introduction

With the rapid development of artificial intelligence technology, deep learning, as one of its outstanding achievements, has realized revolutionary breakthroughs in fields such as image recognition, speech processing, and natural language understanding. The history of deep learning is long-standing, with its precursor being the artificial neural networks of the last century. However, it was not until the recent decade, thanks to the significant improvement in computing power and the explosive growth of data, that deep learning truly demonstrated its great potential. By simulating the mechanism of human brain information processing, deep learning can automatically learn complex features from data, especially in handling high-dimensional image data, where its multi-level nonlinear transformation advantages have surpassed human performance in many tasks.

Nevertheless, the success of deep learning models largely depends on the support of a large amount of annotated data, which often becomes a limiting factor in real-world applications. Particularly in specialized fields like plant leaf disease recognition, the acquisition of high-quality annotated data is not only costly but also limited by the threshold of professional knowledge, which restricts the speed and scale of data annotation. Moreover, the diversity of diseases and the complexity of environmental factors in plant disease recognition often lead to poor performance of deep learning models on novel data, due to overfitting to the training data.

In response to this issue, few-shot learning has emerged, aiming to enable models to learn quickly from limited data and generalize to new tasks. This learning paradigm mimics human rapid learning ability, i.e., grasping new concepts swiftly through a few examples. In the context of plant leaf disease recognition, few-shot learning can use a small number of annotated samples to quickly adapt to new disease types, which is of significant practical importance for rapid response to emerging diseases. However, few-shot learning models still face the challenge of insufficient generalization ability, as a small number of samples can hardly represent the entire data distribution, limiting the model's predictive power.

To address this challenge, this study proposes a new method that combines self-supervised learning [1] and semi-supervised learning [2]. Self-supervised learning learns data representations by predicting the unannotated parts of the data without external annotation information, allowing the model to use a large amount of unannotated data for pre-training and learning more generalized feature representations. Semi-supervised learning further builds on this by generating pseudo-labels through predictions on unannotated data, which are used for model fine-tuning, increasing the sample diversity of model training, thereby improving the model's generalization ability. This study integrates these two methods with the aim of capturing a wide range of features through self-supervised pre-training and refining the model's adaptability to specific tasks through semi-supervised fine-tuning, significantly enhancing the model's generalization ability in plant leaf disease recognition tasks.

The contribution of this study lies in proposing an enhanced few-shot learning framework that combines the advantages of self-supervised learning and semi-supervised learning, providing a new solution for plant leaf disease recognition. We expect this study to promote the development of precision agriculture and provide strong technical support for the early recognition and prevention of plant diseases.

2. Research Method

(1) Overview of research method

In the research of plant leaf disease recognition, we face the dual challenges of scarce sample quantity and insufficient sample diversity. These challenges limit the performance of traditional deep learning methods in the absence of extensive annotated data support. Meta-learning, or the concept of "learning to learn", thus comes into play. It enables the model to adapt quickly to new tasks through training with a small number of samples. In the scenario of few-shot learning, this concept is particularly crucial as it allows the model to rapidly adjust its parameters after encountering a few new category samples to recognize new plant leaf disease categories.

The method proposed in this study innovatively combines self-supervised learning with semi-supervised learning within the framework of meta-learning. Self-supervised learning, as the core of the pre-training phase, guides the model to capture the intrinsic features of data through auxiliary tasks (such as predicting the rotation angle of an image) without external annotation information. This phase of learning provides the model with rich feature representations, laying the foundation for subsequent tasks.

After self-supervised learning, semi-supervised learning, as the key to the fine-tuning phase, introduces a small amount of annotated data and a large amount of unannotated data, fine-tuning the model using the feature representations learned during the pre-training phase. In this phase, the model generates high-confidence pseudo-labels based on its current parameter state for unannotated data, which, along with real labels, are used for further fine-tuning of the model.

The method that combines self-supervised learning and semi-supervised learning is expected to bring multiple advantages. First, self-supervised learning provides the model with strong feature extraction capabilities during the pre-training phase, which can be effectively utilized in subsequent learning. Second, semi-supervised learning enhances the model's generalization ability for unknown categories by introducing pseudo-labels during the fine-tuning phase. This combination not only improves the model's adaptability to new categories but also enhances its generalization ability in practical applications, bringing a new and effective solution to the field of plant leaf disease recognition.

(2) Self-supervised learning-based pre-training

In this study, the plant leaf disease recognition method is based on an enhanced few-shot learning strategy, divided into meta-training and meta-testing phases. The meta-training phase uses the source domain dataset for model training, which is task-driven. In each task, samples of N categories are randomly sampled, with each category containing T data points. The first K (K<T) data points of all categories form the support set, used for establishing the plant leaf disease classification pattern, and the remaining T-K data points form the query set, used to validate and fine-tune this classification pattern. This is the N-way K-shot method in few-shot learning. During model training, the Prototype method [3] is used to construct prototypes for each category. The goal of model training is to minimize the classification loss of labeled training samples.

To further enhance the model's generalization ability, the meta-training phase adopts a joint training method of supervision and self-supervision. That is, while conducting supervised training, a large amount of unlabeled data is used to construct an image rotation task as a self-supervised task. Specifically, the original image is rotated by 0 degrees, 90 degrees, 180 degrees, and 270 degrees, and the model needs to predict the rotation angle of the image. The goal of this self-supervised task is to minimize the self-supervised loss. In this way, the model not only learns the classification knowledge from labeled data but also learns more feature representations through the self-supervised task.

(3) Semi-supervised learning-based fine-tuning

The meta-testing phase uses the target domain dataset for model fine-tuning and testing. In this phase, we adopt a semi-supervised learning strategy. Specifically, the model is initially fine-tuned using a small amount of labeled data, followed by the introduction of a large amount of unlabeled data. For the unlabeled data, the model classifies and predicts based on its current parameter state and generates pseudo-labels based on the confidence of the prediction results. These pseudo-labels, together with real labels, are used for further fine-tuning of the model.

Through this semi-supervised learning method, the model can learn from a large amount of unlabeled data with limited annotated data support, thereby improving its recognition ability for new categories. The core of this method is that by introducing pseudo-labels, the model can continuously adjust its parameters during the learning process to adapt to the new data distribution, thereby enhancing its generalization ability.

3. Experiments and analysis

To verify the effectiveness of the proposed method, we chose the widely used Plant Village [4] dataset for experiments in plant disease recognition. This dataset contains 54,303 images of plant leaves, covering 14 plants and 26 types of leaf diseases, forming 38 varieties (including healthy types) based on the combination of plants and leaf diseases. In this study, we randomly selected samples from 28 categories as the meta-training dataset, and samples from the remaining 10 categories were used as the meta-testing dataset. The network structure before the fully connected layer in the Inception V3 model [5] was used for feature extraction.

In the meta-training phase, to adapt to the scenario of few-shot learning, we selected a limited but diverse set of samples from the dataset. Specifically, we constructed different N-way K-shot tasks, that is, randomly selecting 3 or 5 categories from the meta-training set, choosing 5 or 10 annotated samples per category as the support set, and selecting 15 different samples per category as the query set for model evaluation. The meta-testing phase fixedly generated 500 tasks, with the N-way K-shot parameters the same as those in the meta-training set. When predicting pseudo-labels for unlabeled samples, only samples with a prediction confidence greater than 99.3% were sent to the model's secondary fine-tuning phase.

Table 1. Comparison of classification accuracy for different parameters under different few shot method

Method	Category number	K =5	K =10
Standard few shot method	N =3	0.752	0.826
	N =5	0.713	0.805
Our enhanced few shot method	N =3	0.791	0.863
	N =5	0.756	0.837

Table 1 compares the classification accuracy of plant leaf disease under the standard few-shot setting and the enhanced few-shot setting adopted in this paper for different N-way K-shot configurations. The experimental results show that the model combining self-supervised learning and semi-supervised learning has a significant performance improvement in few-shot plant disease recognition tasks compared to the standard few-shot learning method. This verifies the effectiveness of self-supervised pre-training in the meta-training phase combined with semi-supervised learning in the meta-testing phase in improving model performance. When a certain few-shot learning setting is fixed, a larger K value under the same N value obtained a higher classification accuracy. This is because the K value represents the number of samples per category when the network model is trained, and a larger K value means that there is more information in each category, and the prototype features calculated can more accurately represent the corresponding category, thereby obtaining a higher classification accuracy.

4. Conclusion

This study explored a new way to improve the model's generalization ability in few-shot plant leaf disease classification tasks by combining self-supervised learning and semi-supervised learning. Our experimental results show that this method can effectively use limited annotated data and expand the training set by generating high-quality pseudo-labels, thereby significantly improving the model's recognition performance on unseen categories. In future research, we will explore more self-supervised tasks beyond image rotation to further enrich the feature representations learned by the model. In addition, we will study more complex pseudo-label generation algorithms, such as cluster-based methods, to improve the accuracy of pseudo-labels and the robustness of the model. We look forward to future research continuing to delve deeper on the basis of this paper, continuously promoting the development and perfection of few-shot learning technology in the field of agriculture.

References

[1] Gidaris S, Bursuc A, Komodakis N, et al. Boosting few-shot visual learning with self-supervision[C]//Proceedings of the IEEE/CVF international conference on computer vision. 2019: 8059-8068.

[2] Li Y, Chao X. Semi-supervised few-shot learning approach for plant diseases recognition[J]. Plant Methods, 2021, 17: 1-10.

[3] Snell J, Swersky K, Zemel R. Prototypical networks for few-shot learning[J]. Advances in neural information processing systems, 2017, 30.

[4] Hughes D, Salathé M. An open access repository of images on plant health to enable the development of mobile disease diagnostics[J]. arXiv preprint arXiv:1511.08060, 2015.

[5] Szegedy C, Vanhoucke V, Ioffe S, et al. Rethinking the inception architecture for computer vision[C]//Proceedings of the IEEE conference on computer vision and pattern recognition. 2016: 2818-2826.