Few-Shot Image Classification using Meta-Learning

CS6350 : Computer Vision

TPA No: 18

Introduction

Conventional deep learning models have achieved remarkable success in image classification tasks, primarily due to the availability of large-scale labeled datasets. However, collecting and annotating such datasets is labor-intensive and often impractical, especially for rare or specialized categories. This limitation motivates the field of few-shot learning (FSL), which aims to recognize new classes using only a small number of labeled examples. Inspired by human learning, few-shot learning focuses on rapidly adapting to new tasks with minimal supervision. Recent advances in meta-learning or "learning to learn" have shown promise in addressing this challenge by training models on a distribution of tasks rather than a single task. This project explores few-shot image classification using meta-learning methods (e.g. prototype-based), aiming to evaluate and benchmark their performance on standard few-shot datasets.

Problem Statement

Deep learning models typically require large amounts of labeled data to generalize well. However, in many real-world scenarios, acquiring abundant labeled examples per class is infeasible. This project aims to explore few-shot learning approaches for image classification, where the goal is to correctly classify query images given only a few labeled examples (support set) per class. We will implement and evaluate prototype-based methods using popular few-shot learning datasets.

Inputs

- A few labeled examples per class (support set), e.g., 5-way 1-shot: 5 classes, 1 image per class.
- Unlabeled query images from the same set of classes.

Figure 1 provides a sample for a single training episode for 3-way 3-shot setting.

Figure 1: Example of a Single Episode

Expected Output

The expected outputs are

- The predicted class label for each query image.
- Accuracy of the model across test episodes.

Dataset

The datasets are :

- **Mini-ImageNet**
  - **https://www.kaggle.com/datasets/arjunashok33/miniimagenet**
- **Tiered-ImageNet**
  - **https://github.com/yaoyao-liu/tiered-imagenet-tools**
  - **https://www.kaggle.com/datasets/arjun2000ashok/tieredimagenet**
- **CIFAR-100**
  - **https://www.cs.toronto.edu/ kriz/cifar.html**

- **CIFAR-FS**
  - **https://www.kaggle.com/datasets/keywhere/cifarfs**

References

1. Snell, J., Swersky, K., & Zemel, R. (2017). Prototypical networks for few-shot learning. Advances in neural information processing systems, 30.

2. Xu, J., & Le, H. (2022). Generating representative samples for few-shot classification. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 9003-9013).

3. Ye, H. J., Hu, H., Zhan, D. C., & Sha, F. (2020). Few-shot learning via embedding adaptation with set-to-set functions. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 8808-8817).

4. Cheng, H., Yang, S., Zhou, J. T., Guo, L., & Wen, B. (2023). Frequency guidance matters in few-shot learning. In Proceedings of the IEEE/CVF international conference on computer vision (pp. 11814-11824).

5. Sun, S., & Gao, H. (2023). Meta-AdaM: An meta-learned adaptive optimizer with momentum for few-shot learning. Advances in Neural Information Processing Systems, 36, 65441-65455.

6. Zhang, H., Xu, J., Jiang, S., & He, Z. (2024). Simple semantic-aided few-shot learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 28588-28597).