EAI Endorsed Transactions
on Industrial Networks and Intelligent Systems
EAI.EU Research Article

Few-Shot Classification Of Brain Cancer Images Using Meta-Learning Algorithms

Tuyet-Nhi T. Nguyen[1], Muhammad Fahim[2], Bradley D. E. McNiven[1], Quang Nhat Le[1], Nhan Duc Le[3,*]

[1]Memorial University, St. John's, NL A1B3X5, Canada
[2]Queen's University Belfast, Belfast BT7, 1NN, UK
[3]Danang Hospital, Danang 50000, Viet Nam

Abstract

The primary objective of deep learning is to have good performance on a large dataset. However, when the model lacks sufficient data, it becomes a challenge to achieve high accuracy in predicting these unfamiliar classes. In fact, the real-world dataset often introduces new classes, and some types of data are difficult to collect or simulate, such as medical images. A subset of machine learning is meta learning, or "learning-to-learn", which can tackle these problems. In this paper, a few-shot classification model is proposed to classify three types of brain cancer: glioma brain cancer, meningioma brain cancer, and brain tumor cancer. To achieve this, we employ an episodic meta-training paradigm that integrates the model-agnostic meta-learning (MAML) framework with a prototypical network (ProtoNet) to train the model. ProtoNet focuses on learning a metric space by computing distances to class prototypes of each class, while MAML concentrates on finding the optimal initialization parameters for the model to enable the model to learn quickly on a few labeled samples. In addition, we compute and report the average accuracy for the baseline and our methods to assess the quality of the prediction confidence. Simulation results indicate that our proposed approach substantially surpasses the performance of the baseline ResNet18 model, achieving an average accuracy improvement from 46.33% to 92.08% across different few-shot settings. These findings highlight the potential of combining metric-based and optimization-based meta-learning techniques to improve diagnostic support in healthcare applications.

Received on 26 September 2025; accepted on 04 November 2025; published on 06 November 2025

Keywords: Brain cancer, cancer classification, few-shot learning, meta-learning

doi:10.4108/eetinis.124.10405

1. Introduction

Machine learning has recently achieved notable success in different areas such as autonomous driving cars [1], face recognition [2], and handwritten digit recognition [3]. However, traditional deep learning algorithms generally rely on massive datasets and significant computational resources to achieve good performance. This dependency restricts their use in many practical situations where data are limited or difficult to obtain. In particular, medical imaging poses a considerable challenge because the acquisition of annotated datasets often requires expert knowledge, is time-consuming, and raises ethical concerns about patient privacy [4][5].

In traditional deep learning, models are typically trained for a single and specific task. For instance, when we train a model to perform task A. If we have a new task B that is relevant to task A, we have to write a new model from scratch to perform task B. As a result, for each new task, we have to train a new model from scratch, although those tasks might be related. This leads to inefficiency, especially when the tasks share underlying similarities. To address this issue, researchers have proposed meta-learning [6], or "learning to learn," which focuses on enabling algorithms to adapt rapidly to new tasks using only

*Corresponding author. E-mail: drnhandanang@gmail.com

a few training examples. Recent research on meta-learning applications has shown good performance in various domains, such as computer vision [7], natural language processing [8], and reinforcement learning [9]. In meta-learning, one of the most promising algorithms is few-shot learning (FSL) [10]. This technique aims to generalize well to new tasks by improving the previous experience across multiple tasks during the training process. Indeed, a FSL algorithm can provide a good prediction with only a few samples for each class.

The application of FSL is particularly relevant in healthcare and medical imaging, where data is usually limited. For example, rare cancers often lack sufficient imaging data. As a result, it is hard to train conventional deep learning systems [11]. Studies have shown that FSL techniques can be successfully applied to dermatological diseases [12], histopathological image classification [13], and radiological image analysis [14]. Furthermore, compared to large benchmark datasets such as MiniImageNet [15], CIFAR-100 [16], and CUB-200 [17], medical datasets typically contain fewer samples per class, so there are more generalization challenges due to the limited data. This highlights the urgent need for algorithms that can adapt quickly to small-scale datasets while maintaining high predictive accuracy.

Brain cancer is a critical area because of their limited data. Early and accurate classification of tumor types is important for treatment planning and prediction. However, manual diagnosis through imaging modalities such as magnetic resonance imaging (MRI) or histopathology can be a time-consuming process and is susceptible to mistakes made by humans [18]. Automated classification methods could help radiologists by providing rapid, consistent second opinions, but they must have a good performance even when training data is limited. Thus, integrating FSL into brain cancer classification can improve diagnostic support systems in the real-world healthcare settings.

A further challenge in brain cancer research is the availability and quality of datasets. Many publicly available datasets such as the multiple cancer dataset [19] and the brain tumor segmentation (BraTS) benchmark [20] provide valuable resources, but they still have issues with class imbalance. For example, BraTS collects multi-institutional MRI scans to evaluate algorithms for glioma segmentation, but many subtypes of brain tumors remain underrepresented [21]. These limitations highlight the importance of designing algorithms that not only learn from limited examples but also generalize across complex data sources. Few-shot learning models are suitable for this situation because they can extract important features from small numbers of annotated samples while maintaining the accuracy during training [22].

Moreover, it is important to improve brain cancer image classification. An accurate classification can inform treatment strategies, such as distinguishing between surgical resection for meningiomas and combined chemoradiotherapy for gliomas [23]. Traditional machine learning methods may fail to adapt when the model receives new tumor types or rare subcategories, which limits their use in fast-changing clinical environments. Hence, if we combine meta-learning strategies such as model-agnostic meta-learning (MAML) and prototypical network (ProtoNet), doctors could benefit from decision-support systems that adapt quickly to new cases with minimal additional labeling effort. Furthermore, this approach has the potential to reduce heavy workload, improve early detection, and facilitate more personalized care for patients with brain cancer [24][25].

To overcome the above-mentioned obstacles, our study proposes a hybrid framework that leverages the strengths of both metric-based learning and optimization-based meta-learning. In detail, we integrate ProtoNet [26], which classifies samples by creating a metric space from calculating distances to class prototypes, with the MAML framework [9], which finds optimal initialization parameters for rapid adaptation to new tasks. By combining both approaches, our model achieves high generalization even with limited labeled samples. Especially, it significantly outperforms traditional deep learning baselines such as ResNet18 [27]. This method can demonstrate technical effectiveness and also contribute to more practical and impactful applications in real life.

1.1. Motivation and Main Contributions

In this paper, we focus on implementing a combination algorithms of ProtoNet [26] and MAML [9]. Medical image analysis remains a highly difficult field of study in artificial intelligence because of the limited number of annotated datasets. In brain cancer classification, this challenge is even more critical, as accurate diagnosis is essential for treatment planning and prognosis of disease. Moreover, it is difficult and costly to collect sufficient labeled images for each tumor type. Conventional deep learning approaches, including convolutional neural networks (CNNs) typically require large datasets to generalize well and often fail to adapt when only a few training samples are available. Consequently, there is a strong need for methods that can achieve good performance under data-scarce conditions.

This paper offers the following key contributions:

- **Hybrid meta-learning framework:** We propose a novel FSL framework that integrates ProtoNet with MAML. This technique enhances their

Table 1. Notations

Symbol	Definition
n	Number of classes (ways) in each episode
k	Number of support samples per class (shots)
S_i	Support set containing k labeled samples for each class
Q_i	Query set containing unlabeled samples for evaluation
x_s, y_s	Support sample and its corresponding label
x_q, y_q	Query sample and its ground-truth label
$f_\phi(x)$	Embedding function based on ResNet18 backbone
C_S	Prototype vector of a class, computed as the mean embedding of its support samples
d	Euclidean distance between a query embedding and class prototypes
$p(y = k\|x)$	Predicted probability that query sample x belongs to class k

complementary strengths to improve brain cancer image classification.

- **Application to brain cancer dataset:** We evaluate the proposed approach on the brain cancer subset of the multiple cancer dataset. It shows that our method can generalize effectively with only a few labeled examples per class.
- **Performance gains over baselines:** Our results show that the proposed approach achieves considerably better performance than a strong baseline model, i.e., ResNet18. Our technique improves average accuracy from 46.33% (5-shot) to 92.08% (30-shot).
- **Comprehensive analysis:** We provide detailed evaluations, including confusion matrices and error analysis. These results highlight both the strengths and limitations of the hybrid framework in classifying glioma, meningioma, and tumor cases.

1.2. Paper Organization and Notations

The remainder of the paper is structured as follows: Section 2 provides a literature review of meta-learning, few-shot learning, and related network architectures. In section 3, the process of how we train the model is shown, and its performance is then evaluated. Section 4 explains the results of our method. Finally, a summary of the main points and contributions is provided in section 5. Table 1 summarizes the symbols and their corresponding definitions used in the paper.

2. Related Works

2.1. Meta-Learning

As a subset of machine learning, meta-learning has been widely used in recent years due to its ability to work well with limited data. Traditional deep learning has some limitations, e.g., the model only performs well with a large dataset and sufficient computing resources. Meta-learning, also known as "learning-to-learn", has become a powerful approach for tackling the shortcomings of conventional deep learning when training data are scarce. Instead of relying on large labeled datasets, meta-learning algorithms aim to extract transferable knowledge from a distribution of tasks, so the model can have a rapid adaptation to new problems with few samples. Meta-learning offers a different approach in which a machine learning model learns by accumulating experience across multiple episodes. Recently, meta-learning has been effectively utilized in diverse fields, including few-shot image recognition [26], segment anything model [28], zero-shot learning [29].

One of the most widely adopted optimization-based approaches is MAML [9]. MAML focuses on learning initialization parameters that can be quickly fine-tuned for new tasks by using only a few gradient updates. Its simplicity and broad applicability have inspired many variants, such as first-order MAML [9], which reduce computational overhead while maintaining adaptability. In medical imaging, meta-learning has shown strong promise. For example, Gu et al. [30] demonstrated that MAML could effectively adapt to new radiological image classification tasks with minimal supervision, thus outperforming standard transfer learning. These results suggest that optimization-based meta-learning is well-suited for healthcare, where labeled data are typically difficult to acquire and scarce.

2.2. Few-Shot Learning

In FSL [10], the model aims to learn new classes with limited labeled data. While real-world data scenarios usually arrive continuously, some data such as medical images are hard to collect or simulate. Therefore, it is necessary for a model to achieve high performance on a few samples. There are two main categories of

neural architectures: metric-based and optimization-based networks. Some metric-based methods, such as matching networks [7], ProtoNet [26], and relation networks [31] aim to classify query samples by comparing them with support samples in a learned embedding space. In more detail, matching networks use an attention-based mechanism for the support images, while ProtoNet computes class prototypes and applies Euclidean distance for classification. Moreover, the relation networks focus on learning a deep similarity function. These approaches are simple, efficient, and have shown competitive performance in real-world datasets.

ProtoNet [26] is frequently utilized because it is both simple and effective. By computing class prototypes in an embedding space and classifying queries based on Euclidean distance, ProtoNet has achieved competitive results on benchmark datasets like MiniImageNet [15]. In the medical domain, metric-based methods have also shown success. For example, when applying ProtoNet to skin disease image datasets [12], the results show that few-shot classification could achieve performance comparable to fully supervised models. On the other hand, optimization-based approaches such as MAML [9] emphasize rapid adaptation by learning task-agnostic initialization parameters. Both categories complement each other: metric-based methods are good at learning robust representations, while optimization-based methods enable efficient task adaptation. This motivates hybrid strategies, such as the approach adopted in this study, which combine both paradigms to leverage their respective strengths.

2.3. Related Network Architectures

The choice of backbone network is important in few-shot learning because it determines the quality of feature embeddings. ResNet [27] has been widely used as the backbone of both ProtoNet and MAML due to its balance between efficiency and representational power. However, recent advances suggest that alternative architectures may offer improvements. DenseNet [32] has been shown to capture richer feature hierarchies with fewer parameters, while EfficientNet [33] introduces a compound scaling strategy that attains leading accuracy while lowering computational cost. In addition, vision transformers [34] are popular for their ability to model long-range dependencies, and recent works have adapted them for FSL with promising results in medical imaging tasks.

In healthcare, the application of meta-learning with advanced backbones has notable improvements. For instance, a transformer-based ProtoNet framework for histopathological image classification [13] achieves higher performance against inter-class similarity. Similarly, hybrid models that combine CNN-based embeddings with meta-learning strategies, such as MAML [9] with siamese networks, have been applied to MRI brain tumor classification with encouraging results [35]. These findings suggest that combining metric-based and optimization-based meta-learning with carefully chosen backbones can enhance diagnostic accuracy, especially in limited medical imaging scenarios. In this paper, we adopt ResNet18 as the backbone for its proven reliability in medical imaging and combine it with ProtoNet and MAML to balance between efficiency and accuracy for few-shot brain cancer classification.

3. System Model and Problem Formulation

In this section, we propose our method for classifying brain cancer images with a few samples per class. ProtoNet [26] is used to calculate the prototype distances between a query image and support images. Then, MAML [9] is applied to find optimal initialization parameters.

3.1. Dataset

Deep learning models typically achieve strong performance on the tasks they have been trained on before. However, they have problems with generalizing on the tasks that they have been trained on before with a few samples. FSL aims to tackle this problem by focusing on achieving good performance on the tasks that using only a few examples.

In our study, we train the model to classify three types of brain cancer. For this, we use the brain cancer dataset taken from the multiple cancer dataset [19]. The multiple cancer dataset includes 130,000 histopathological images from eight cancer types: *acute lymphoblastic leukemia (ALL)*, *brain cancer*, *breast cancer*, *cervical cancer*, *kidney cancer*, *lung and colon cancer*, *lymphoma*, and *oral cancer*. For the brain cancer dataset, each class contains around 5,000 X-ray images. These images are divided into training, testing, and validation sets with a ratio of 70:15:15, which means 3,500 images for training, 750 images for testing, and the remaining for validation.

The primary aim of our study is to construct a strong generalization to those three types of brain cancer with a few samples in each class. To achieve this, we employ an n-way k-shot episodic training paradigm. In each episode, the training data $\mathcal{D}_{train}$ consists of a support set $\mathcal{S}_i$ and a query set $\mathcal{Q}_i$: $\mathcal{D}_{train} = \{\mathcal{S}_i, \mathcal{Q}_i\}$. We train our n-way k-shot model on these classes to classify the query images. n-way = 3 denotes the number of brain cancer classes, and k-shot denotes used number of images of each support class, e.g., $k = 5, 10$. The support set $\mathcal{S}_i = \{x_s, y_s\}$ includes x_s images with their respective labels y_s. Similarly, the query set $\mathcal{Q}_i = \{x_q, y_q\}$.

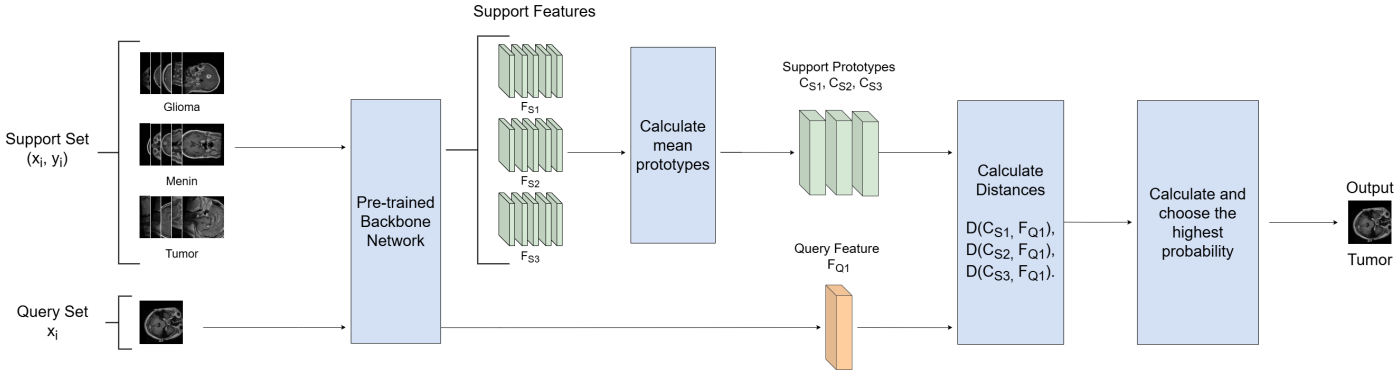

Figure 1. Each episode of the training process consists of a support set and a query set, and they are taken without replacement. These sets are passed through a pre-trained backbone network, i.e., ResNet18, to extract support features and query features. The model then calculates support prototypes by using support features. Next, Euclidean distances are computed between each query feature and all support prototypes: $d_{\text{distances}} = \left\{d(C_{S_1}, F_{Q_1}), d(C_{S_2}, F_{Q_1}), \ldots, d(C_{S_k}, F_{Q_1})\right\}$. After that, these distances are passed through a softmax function to obtain a probability distribution over the support classes. Finally, the model outputs the predicted class label $\hat{y}_i$ for each query sample based on the highest probability.

For evaluation, we sample episodes in the same format $\mathcal{D}_{test} = \{\mathcal{S}_i, \mathcal{Q}_i\}$ with n-way k-shot. In this process, we use the same k-shot values as in training. Since the testing samples are disjoint from the training samples, this setup provides a way to examine how well the model can generalize to unknown cancer samples with a few images.

3.2. Prototypical Network

ProtoNet [26] is a metric-based FSL approach that classifies samples based on their distance to class prototypes in an embedding space. Figure 1 shows that the training data $\mathcal{D}_{train}$ consists of the support set $\mathcal{S}_i$ and the query set $\mathcal{Q}_i$: $\mathcal{D}_{train} = \{\mathcal{S}_i, \mathcal{Q}_i\}$ in each episode. These support set and query set are inputs of the model and are used for training at the same time. The support set includes random k images of each class with their respective labels: $\mathcal{S}_i = \{x_s, y_s\}$, where x_s is the support image and y_s is the corresponding class label. The query set only contains some images $\mathcal{Q}_i = \{x_q\}$, where x_q is the query image. The query label y_q is not used during training to ensure a fair and unbiased distance-based classification. Instead, it is used to compute the loss for backpropagation. This method encourages the model to learn a transferable metric that generalizes well to unseen samples, rather than memorizing specific class-label associations.

We pass these input images to the pretrained backbone network. The goal of this step is to use a ResNet18 [27] as a pretrained backbone network to extract features of the support set and the query set. The ResNet18 is trained with the ImageNet1K dataset [36], which contains 1,000 object categories. The ResNet18 network extracts support features F_s and query features F_q. For the support features, the model computes a prototype vector $\mathcal{C}_S$ for each class, i.e., each type of brain cancer, by averaging the feature embeddings of the support samples belonging to that class. The formula for calculating the mean embeddings of the data points under each class is shown as follows:

$$\text{Class Prototype } (\mathcal{C}_S) = \frac{1}{S} \sum_{(x_s, y_s) \in S} f_\phi (x_s), \qquad (1)$$

where $\mathcal{C}_S$ represents the class prototype of the corresponding embedding support features F_S taken from the ResNet18, $\mathcal{S}$ denotes the support set with support images and their respective labels $\{x_s, y_s\}$, and f_ϕ is the embedding function pretrained ResNet18.

Once we have the support class prototypes $\mathcal{C}_S$ and query set embeddings F_Q, the Euclidean distance d between query set embeddings F_Q and the class prototypes $\mathcal{C}_S$ is calculated. Figure 2 illustrates the distance d between a query point and the three class prototypes of three types of brain cancer. d is denoted as

$$\text{Distance } (d) = \text{ Euclidean\_Distance } \left(\mathcal{C}_S, F_Q\right). \qquad (2)$$

Once we get the three distances d between the query point and the three class prototypes, the class probability p of a query set is then computed by applying softmax over the distance d, which is given by

$$p_(y = k \mid x) = \frac{\exp\left(-d\left(f_\phi(x_q), \mathcal{C}_S\right)\right)}{\sum_k \exp\left(-d\left(f_\phi(x_q), \mathcal{C}_S\right)\right)}, \qquad (3)$$

where $d(f_\phi(x_q), \mathcal{C}_S)$ denotes the Euclidean distance between the query embeddings and the class prototype.

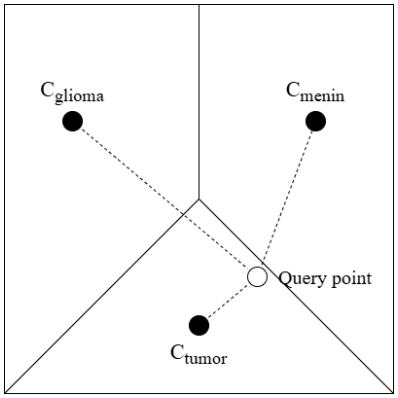

Figure 2. The use of Euclidean distance for classification in the feature space is illustrated in the above picture. The open circle represents a query sample that needs to be classified, and the filled circles denote the class centroids for *glioma* (C_{glioma}), *meningioma* (C_{menin}), and *tumor* (C_{tumor}). The decision is made by calculating the Euclidean distance between the query sample and each class centroid. The sample is then assigned to the class with the nearest centroid, thus effectively demonstrating the principle of distance-based classification in FSL.

Table 2. Configuration of MAML in our proposed method

Component	Description
Backbone	ResNet18 pretrained on the ImageNet1K dataset, removing the final layer
Meta-learner	ProtoNet
Episode sampling	Randomly sampling N classes, each with K support and Q query images
Inner loop	Stochastic gradient descent (SGD) [37] on support set
Outer loop	Meta-update using query loss and Adam optimizer [38]
Loss function	Cross-entropy [39] on query predictions
Learning rate	Outer: 0.0001 (for Adam), Inner: 0.01 (for SGD)

Here, we take the negative value because when the loss function is calculated later, the *cross entropy loss* takes a negative log probability. Thus, it is necessary to use the negative value of the probability.

3.3. Model-Agnostic Meta-Learning

After each episodic training, the model updates their parameters to minimize the loss value. To enable the model to quickly learn from a few examples, we apply

Algorithm 1 The proposed MAML with ProtoNet procedure.

Require: $p(\mathcal{T})$: Distribution over tasks
Require: α, β: Inner and outer learning rates
1: Randomly initialize θ
2: **while** not done **do**:
3: Sample batch of tasks $\mathcal{T}_i \sim p(\mathcal{T})$
4: **for all** $\mathcal{T}_i$ **do**
5: Construct support set $\mathcal{S}_i$ and query set $\mathcal{Q}_i$
6: Compute class prototypes
7: Compute inner loss $\mathcal{L}_{\mathcal{T}_i}^{\text{support}}$ using distances to prototypes
8: Compute adapted parameters: $\theta_i' = \theta - \alpha \nabla_\theta \mathcal{L}_{\mathcal{T}_i}^{\text{support}}$
9: Re-compute prototypes c_k' using $f_{\theta_i'}$
10: Compute meta-loss $\mathcal{L}_{\mathcal{T}_i}^{\text{query}}$ on $\mathcal{Q}_i$
11: **end for**
12: Update $\theta \leftarrow \theta - \beta \nabla_\theta \sum_{\mathcal{T}_i} \mathcal{L}_{\mathcal{T}_i}^{\text{query}}$
13: **end while**

MAML algorithm [9]. MAML is an optimization-driven meta-learning framework designed to identify optimal model initialization parameters, allowing the model to quickly adjust to new tasks with just a few gradient steps.

To implement the MAML algorithm, some tasks $\mathcal{T}$ are created to quickly find the initialization parameters. Each task $\mathcal{T}_i$ is a few-shot classification problem and consists of a support set $\mathcal{S}_i = (x_j, y_j)_{j=1}^{N \cdot K}$ and a query set $\mathcal{Q}_i = (x_j, y_j)_{j=1}^{Q}$, where N denotes the number of classes and K is the number of support samples per class. We use the support set for inner loop adaptation and the query set for evaluating the generalization performance of the updated model.

In the inner loop, a temporary copy of the ProtoNet model is created and then updated by using gradient descent on the support set. After a small number of gradient steps, the model evaluates their updated parameters on the query set. *Cross-entropy loss* [39] is used to evaluate the temporary model, and the function for calculating the loss value is defined as

$$\mathcal{L}_{CE} = -\sum_{c=1}^{N} y_q^{(c)} \log \hat{y}_q^{(c)}, \tag{4}$$

where $y_q^{(c)}$ and $\hat{y}_q^{(c)}$ denote the ground-truth one-hot encoded label and the predicted probability for class c over N classes, respectively. This loss enhances the model to produce higher probabilities for the correct classes. The gradients from this loss are then used to update the meta-model parameters across multiple tasks in the outer loop.

This technique helps the model quickly learn from a few data points. For the outer loop, we compute the loss of the updated model on the query set. This loss serves as a meta-objective, and its gradient is used to update the parameters of the original (meta) model using the Adam optimizer [38]. This process is repeated over many episodes. As a result, this algorithm enhances the model in learning parameters that can be quickly adjusted to new tasks with only a few labeled samples. The MAML setting is shown in table 2. The procedure of our proposed method is summarized in algorithm 1.

3.4. Evaluation Metrics

We evaluate the effectiveness of our proposed FSL model by comparing it against a baseline method using the same dataset and episode configuration. For the baseline scheme, we implement a standard supervised learning approach by training a pretrained ResNet18 backbone on a limited number of labeled samples. In particular, for each few-shot setting (N-way, K-shot), we randomly choose N classes, K labeled support samples, and Q unlabeled query samples per class. The model is trained in a conventional supervised learning model using only these samples, without any meta-learning strategy.

In contrast, our method uses an episodic meta-training paradigm based on the MAML framework combined with a ProtoNet classifier. This technique improves the model to rapidly adapt to new tasks by using only a few gradient steps. For both the baseline and our methods, we follow the standard few-shot evaluation protocol. Both methods are evaluated under 1,000 randomly sampled episodes for each N-way, K-shot setting.

The primary evaluation metric is the top-1 classification accuracy, which measures the percentage of correctly classified query samples. We report the average top-1 accuracy across all episodes, along with the standard deviation to reflect result stability. In addition, the average cross-entropy loss for both methods are calculated and reported to assess the quality of the prediction confidence. To provide further insights into classification performance, confusion matrices are also presented in figure 3.

It is important to note that no fine-tuning is applied during the evaluation process. The model is evaluated in a purely few-shot setting to show its ability to generalize with a limited dataset.

4. Simulation Results

Table 3 presents the top-1 classification accuracy of our method compared to the baseline scheme across different values of k support samples per class. Based on supervised fine-tuning of a ResNet18 backbone, the baseline shows modest gains with the increase in k, from 46.33% at 5-shot to 52.09% at 30-shot. In contrast, our method, which integrates ProtoNet with the MAML optimization framework, consistently outperforms the baseline by a large margin. Notably, our model achieves 77.75% accuracy at 5-shot and improves to 92.08% at 30-shot, which shows a strong generalization from limited examples. These findings align with prior research showing that meta-learning significantly enhances performance in data-scarce environments. For example, the authors in [26] reported that ProtoNets outperform conventional supervised models on MiniImageNet [15], and the authors in [9] demonstrated that MAML can rapidly adapt to unseen tasks. Our work extends these insights into the medical imaging domain, specifically brain cancer classification, by combining both metric-based and optimization-based strategies.

In Figure 3, the normalized confusion matrices of our proposed method are presented under different k-shot settings. It can be observed that the model obtains high classification accuracy across all cancer types. Especially with $k = 30$, the model achieves correct classification rates of approximately 90% for glioma, meningioma, and tumor brain cancer. While the proposed method achieves strong performance across all few-shot settings, some misclassifications persist, particularly between glioma and meningioma cases. This may be attributed to the high morphological similarity of tumor subtypes in medical images. Such errors highlight an important limitation of purely image-based approaches. In clinical practice, combining imaging data with complementary modalities, e.g., patient history, genomic data, may further improve reliability.

5. Conclusions and Future Directions

In this paper, we introduced a FSL model that combines ProtoNet with MAML for brain cancer classification. The proposed method consistently outperformed a supervised ResNet18 baseline scheme. Especially, our technique achieves up to 92.08% accuracy in the 30-shot setting. These results demonstrate the effectiveness of integrating metric-driven and optimization-driven meta-learning strategies in handling data-scarce medical imaging tasks.

In future research, we intend to expand our framework into an end-to-end network and evaluate it on larger, multi-center medical datasets to ensure a strong study across different imaging conditions. Furthermore, our objective is to build a multimodal system incorporating multimodal data, such as genomic profiles or clinical records, to provide complementary information beyond visual patterns, as suggested by [13] in their transformer-based ProtoNet for histopathology. Moreover, exploring alternative backbone networks like DenseNet [32], EfficientNet [33], or vision transformers

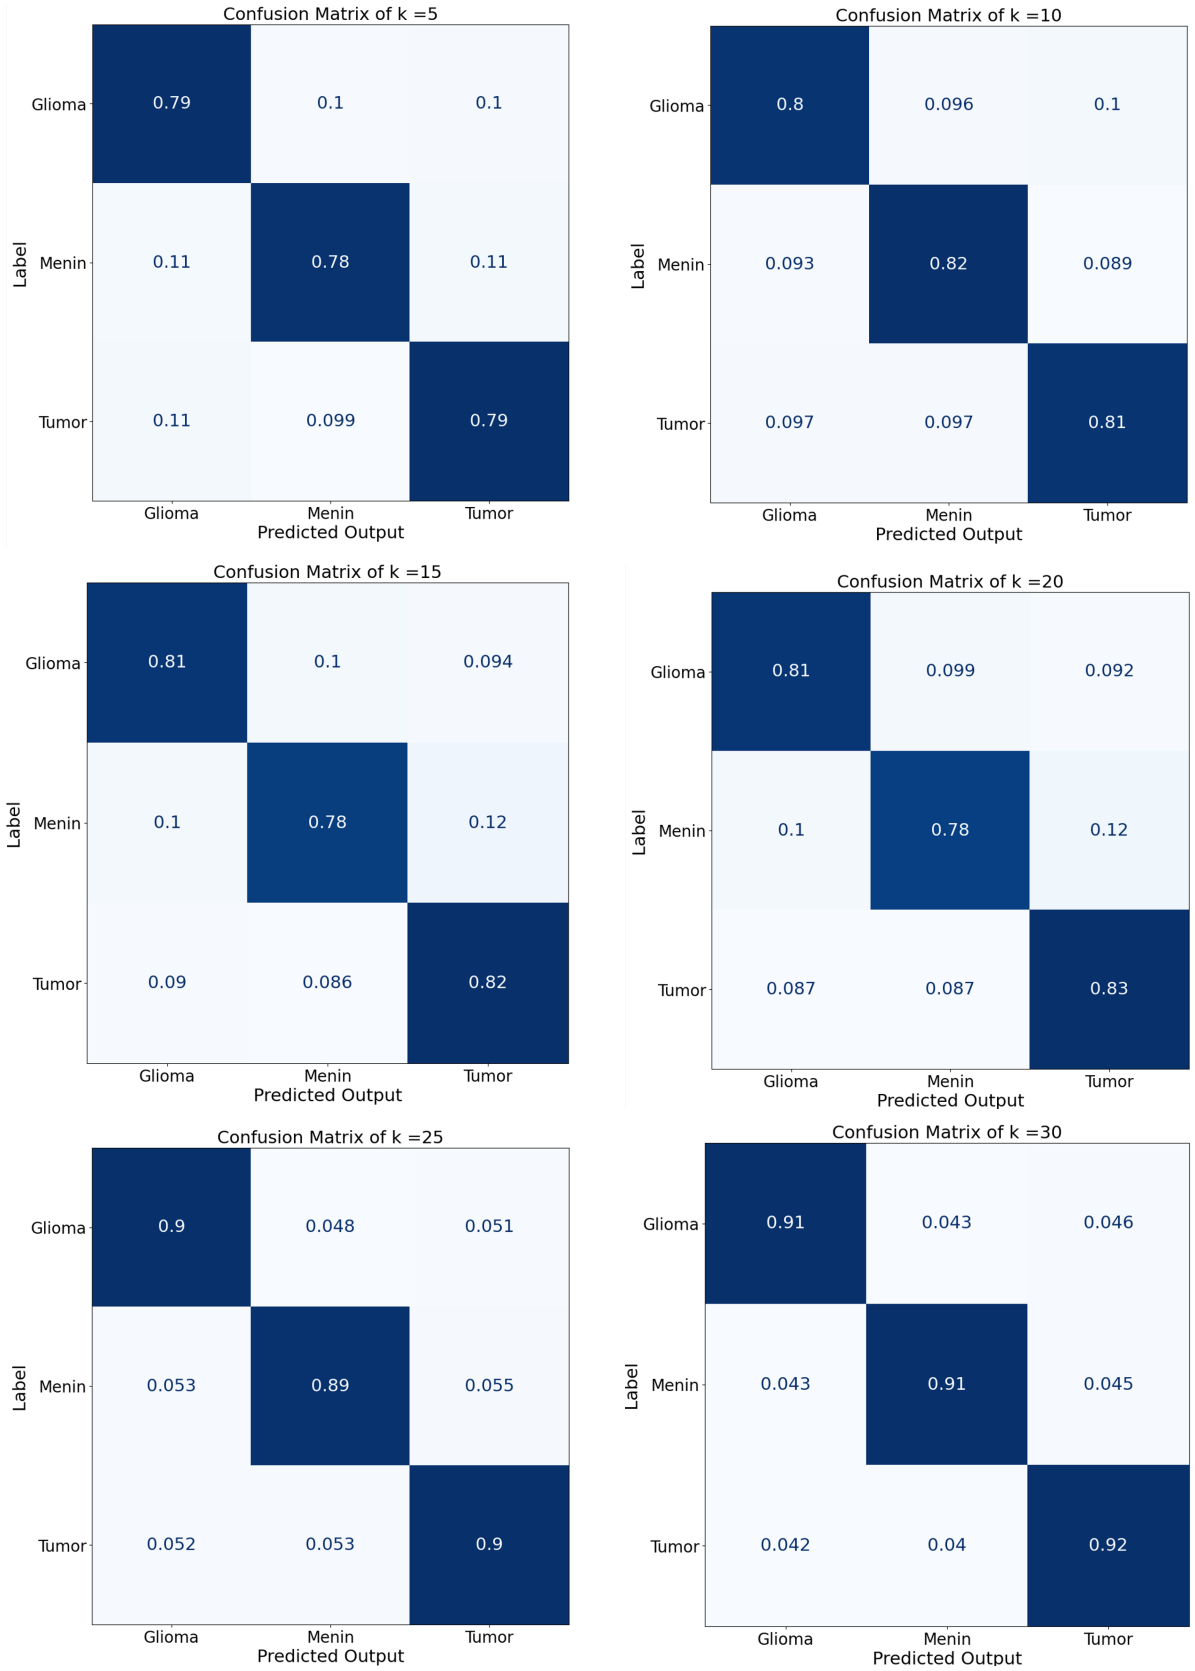

Figure 3. Confusion matrices of the k-NN classifier under six different values of k, i.e., $k = 5, 10, 15, 20, 25, 30$. Each matrix compares predicted labels with ground truth classes: *glioma*, *meningioma*, and *tumor*.

Table 3. Comparison of average accuracy (%) for 3-way k-shot classification using ResNet18 as baseline scheme and our proposed method (ProtoNet + MAML).

k-shot	5	10	15	20	25	30
Baseline: ResNet18 [27]	46.33	49.83	49.49	51.16	51.88	52.09
Our Method: ProtoNet [26] + MAML [9]	**77.75**	**80.02**	**82.29**	**85.47**	**89.24**	**92.08**

[34] may further enhance generalization across diverse imaging conditions.

References

[1] Z. Chen and X. Huang, "End-to-end learning for lane keeping of self-driving cars," in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2017, pp. 1856–1860.

[2] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "DeepFace: Closing the gap to human-level performance in face verification," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2014, pp. 1701–1708.

[3] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.

[4] G. Litjens *et al.*, "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, Dec. 2017.

[5] H. Greenspan, B. van Ginneken, and R. M. Summers, "Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique," *IEEE Transactions on Medical Imaging*, vol. 35, no. 5, pp. 1153–1159, May 2016.

[6] R. Vilalta and Y. Drissi, "A perspective view and survey of meta-learning," *Artificial Intelligence Review*, vol. 18, no. 2, pp. 77–95, Jun. 2002.

[7] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra, "Matching networks for one shot learning," in *Proceedings of the 30th International Conference on Neural Information Processing Systems (NeurIPS)*, Dec. 2016, pp. 3637–3645.

[8] Supriyono, A. P. Wibawa, Suyono, and F. Kurniawan, "Advancements in natural language processing: Implications, challenges, and future directions," *Telematics and Informatics Reports*, vol. 16, pp. 100 173–100 189, Dec. 2024.

[9] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, Aug. 2017, pp. 1126–1135.

[10] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a few examples: A survey on few-shot learning," *ACM Computing Surveys*, vol. 53, no. 3, pp. 1–34, Jun. 2020.

[11] A. Esteva *et al.*, "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, Jan. 2017.

[12] V. Prabhu *et al.*, "Few-shot learning for dermatological disease diagnosis," in *Proceedings of the 4th Machine Learning for Healthcare Conference*, Aug. 2019, pp. 532–552.

[13] J. C. K. R. Hasan, S. Kim and H. S. Han, "Prototypical few-shot learning for histopathology classification: Leveraging foundation models with adapter architectures," *IEEE Access*, vol. 13, pp. 86 356–86 379, May 2025.

[14] H. Oliveira, P. H. Gama, I. Bloch, and R. M. C. Jr., "Meta-learners for few-shot weakly-supervised medical image segmentation," *Pattern Recognition*, vol. 153, pp. 110 471–110 483, Sep. 2024.

[15] S. Ravi and H. Larochelle, "Optimization as a model for few-shot learning," in *Proceedings of the International Conference on Learning Representations (ICLR)*, Apr. 2017, pp. 1–11.

[16] A. Krizhevsky, "Learning multiple layers of features from tiny images," Ph.D. dissertation, University of Toronto, Apr. 2009.

[17] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, "The caltech-ucsd birds-200-2011 dataset," California Institute of Technology, Tech. Rep., Jul. 2011.

[18] S. Bauer, R. Wiest, L.-P. Nolte, and M. Reyes, "A survey of MRI-based medical image analysis for brain tumor studies," *Physics in Medicine & Biology*, vol. 58, no. 13, pp. R97–R129, Jun. 2013.

[19] O. S. Naren, "Multi cancer dataset," https://doi.org/10.34740/KAGGLE/DSV/3415848, 2022.

[20] B. H. Menze *et al.*, "The multimodal brain tumor image segmentation benchmark (BRATS)," *IEEE Transactions on Medical Imaging*, vol. 34, no. 10, pp. 1993–2024, Oct. 2015.

[21] A. Verma and A. K. Yadav, "Brain tumor segmentation with deep learning: Current approaches and future perspectives," *Journal of Neuroscience Methods*, vol. 418, pp. 1–33, Jun. 2025.

[22] Y. Deo *et al.*, "Few-shot learning in diffusion models for generating cerebral aneurysm geometries," in *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI)*, May 2024, pp. 1–5.

[23] M. Weller *et al.*, "European association for neuro-oncology (EANO) guidelines for the diagnosis and treatment of adult astrocytic and oligodendroglial gliomas," *The Lancet Oncology*, vol. 18, no. 6, pp. e315–e329, Jun. 2017.

[24] S. Rathore, M. Habes, M. A. Iftikhar, A. Shacklett, and C. Davatzikos, "A review on neuroimaging-based classification studies and associated feature extraction methods for alzheimer's disease and its prodromal

stages," *NeuroImage*, vol. 155, pp. 530–548, Jul. 2017.
[25] A. R. Dudhe, P. G. Burade, and M. Nikose, "Comprehensive review of brain tumor segmentation and classification using deep learning techniques," in *Proceedings of the IEEE International Students' Conference on Electrical, Electronics and Computer Science (SCEECS)*, Jan. 2025, pp. 1–6.
[26] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS)*, Dec. 2017, pp. 4080–4090.
[27] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2016, pp. 770–778.
[28] A. Kirillov *et al.*, "Segment anything," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct. 2023, pp. 4015–4026.
[29] Y. Xian, B. Schiele, and Z. Akata, "Zero-shot learning – a comprehensive evaluation of the good, the bad and the ugly," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 9, pp. 2251–2265, Sep. 2019.
[30] J. Gu, Y. Wang, Y. Chen, K. Cho, and V. O. Li, "Meta-learning for low-resource neural machine translation," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Nov. 2018, pp. 3622–3631.
[31] F. Sung *et al.*, "Learning to compare: Relation network for few-shot learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2018, pp. 1199–1208.
[32] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jul. 2017, pp. 4700–4708.
[33] M. Tan and Q. V. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *Proceedings of the 36th International Conference on Machine Learning*, Jun. 2019, pp. 6105–6114.
[34] A. Dosovitskiy *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proceedings of the International Conference on Learning Representations (ICLR)*, May 2021, pp. 1–21.
[35] S. Gull and J. Kim, "Metric-based meta-learning approach for few-shot classification of brain tumors using magnetic resonance images," *Electronics*, vol. 14, no. 9, pp. 1–21, May 2025.
[36] O. Russakovsky *et al.*, "ImageNet large scale visual recognition challenge," *International Journal of Computer Vision (IJCV)*, vol. 115, pp. 211–252, Apr. 2015.
[37] L. Bottou, "Large-scale machine learning with stochastic gradient descent," in *Proceedings of COMPSTAT*, Aug. 2010, pp. 177–186.
[38] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proceedings of the International Conference on Learning Representations (ICLR)*, May 2015.
[39] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.