

Refining Prototypes with Selective Feature Maps: A Lightweight Approach to Few-Shot Medical Image Classification

Ranjana Roy Chowdhury

Department of Computer Science & Engineering
 Indian Institute of Technology, Ropar

Deepti R. Bathula

Department of Computer Science & Engineering
 Indian Institute of Technology, Ropar

ABSTRACT

Prototypical Networks operate by embedding both support and query samples into a common feature space and then representing each class with the mean vector of its support embeddings. Yet, the inherent complexity of medical imagery poses significant challenges for isolating features that are both precise and dependable. Consequently, constructing effective prototypes in this domain demands not only sophisticated preprocessing and more powerful embedding architectures, but also deliberate refinement of feature representations. In this context, most important and representative feature map selection is critical. We introduce Selective Feature Representation in Prototypical Networks, a lightweight yet effective enhancement to prototype-based few-shot learning. Proposed approach explicitly refines support embeddings by ranking and selecting the top feature maps for each class, leveraging an ensemble of channel-wise statistics—Global Average Pooling, Max Pooling, and Variance. Built on a compact CONV4 backbone, proposed method outperforms much larger state-of-the-art models on two medical benchmarks: achieving 67.18% (1-shot) and 78.20% (5-shot) on Derm7pt skin-lesion classification, and 63.39% (1-shot), 77.17% (5-shot), and 83.06% (10-shot) on BloodMNIST pathology classification. These gains demonstrate that targeted feature-map selection significantly improves prototype quality and generalization with minimal complexity, offering a practical solution for resource-constrained clinical applications.

General Terms

Meta Learning, Few Shot Learning, Prototypical Networks

Keywords

Meta Learning, Few Shot Learning, Prototypical Networks, Feature Map Selection

1. INTRODUCTION

Since the past few years, deep learning has brought a revolution in diverse fields of mankind through its astonishing features and extremely powerful prediction capabilities throughout various disciplines [12, 13]. It has brought great prosperity in many industries like finance, automotive, manufacturing including healthcare [15] e.g., computer-aided diagnosis [5], medical image retrieval [17], medical image mining [2], medical image sorting [23], etc. Deep learning and its applications have had a huge impact on the healthcare industry and have uplifted diagnostic capability by their outstanding features, flexibility, and robustness. But the main drawback with these models lies in their hunger for huge, annotated samples to achieve efficient performance. Researchers have explored various techniques to address this challenge, including transfer learning [26] and domain adaptation [8], which leverage large source domains to provide models with rich feature representations and favourable initialization for generalization. However, these approaches often

struggle in extreme scenarios, such as severely tail distributed data or significant domain shift.

The necessity for large numbers of labelled samples poses a huge challenge in data-scarce domains such as medicine. Even though many approaches have been proposed, learning from very few annotated samples remains an open challenge. This contrasts with human intelligence, which adapts quickly in a new environment using prior knowledge. This is where meta-learning comes in—it leverages exactly this human-like ability to learn and react from very few examples [10]. This is the crux of meta-learning, which replicates learning from very few samples. Recently, meta-learning has gained popularity as an alternative in low-data regimes, especially in medicine. For instance, in predicting extremely rare disease classes with almost no samples, meta-learning can be highly significant. This few-sample prediction scenario is termed Few-Shot Learning [25]. Various meta-learning techniques have been applied to Few-Shot problems: metric-based approaches like Matching Networks [24] and Prototypical Networks [22], gradient-based methods such as MAML [7] and Optimization-Based Few-Shot [20].

Prototypical Networks [22] have become popular owing to their simple and intuitive design where they form class prototypes by averaging support-set embeddings and then classify queries according to the nearest-prototype distance. However, this vanilla formulation treats every support embedding equally, which may not reflect their true relevance or importance in class prototype formation. Several studies [9, 21] have introduced support-sample reweighting schemes to improve prototype construction, including the previously proposed work IPNET [18] that employed

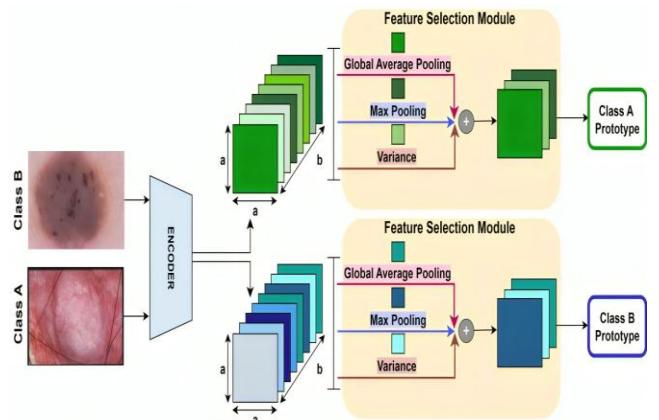


Fig. 1: Overview of the proposed refinement of prototypes in Prototypical Networks, where feature maps from each support embedding are selected using an ensemble strategy that combines global average pooling, max pooling, and variance-based selection

Maximum Mean Discrepancy to identify and emphasize the most influential examples during prototype formation. In parallel, methods such as COMET [3] have explored in extracting semantically meaningful “Concepts” from each class sample to build” Concept-Based Class Prototypes”.

Feature selection plays a critical role in enhancing the discriminative power and robustness of deep representations, particularly in high-dimensional settings like few-shot learning. Recent studies have highlighted that not all features contribute equally to the final prediction, and the presence of redundant or noisy dimensions can adversely affect the learning process [14, 27]. Works such as [4] and [11] have shown that adaptive filtering or attention-based weighting of features can significantly improve generalization and reduce overfitting, especially in low-data regimes. In the context of prototype-based few-shot methods, this becomes even more pertinent since prototype computation depends heavily on the aggregated feature embeddings.

Despite these advances, none of the existing approaches explicitly emphasize filtering out noisy or less informative feature maps within each embedding; they still aggregate all feature maps when computing prototypes, potentially diluting the representation with irrelevant information. In this work, we propose Selective Feature Representation in Prototypical Networks which is a novel feature selection strategy that retains only the most salient feature maps in support sample embeddings, thereby reducing noise and enhancing class prototype quality. Proposed method determines important feature maps through an ensemble of three techniques—Global Average Pooling, Max Pooling, and Variance ensuring both diversity and robustness in the resulting class prototypes. The major contributions of this paper are as follows:

- This work proposes refining prototypes in prototypical networks through selective feature maps, where only the most influential support embedding feature maps are retained for class prototype formation.
- The feature maps are selected using an ensemble of three techniques—Global Average Pooling, Max Pooling, and Variance—ensuring a diverse and informative representation by capturing complementary aspects of the feature space.
- The effectiveness of the proposed method is demonstrated on two medical imaging domains: dermatoscopic imaging using the Derm7pt [6] dataset and pathology using the BloodMNIST [1] dataset.
- t-SNE visualizations on the BloodMNIST dataset show that Selective Feature Representation in Prototypical Networks yields more compact intra-class clusters and clearer inter-class separations than vanilla Prototypical Networks.

The paper begins by introducing the preliminaries in Section 2, then describes the proposed methodology in Section 3. Section 4 details the experimental setup and analysis, followed by the presentation of results in Section 5. Finally, Section 6 concludes the paper and outlines directions for future work.

2. PRELIMINARIES

2.1 Meta Learning

Meta-learning, or “learning to learn” [7], aims to train models that can swiftly master new tasks from only a few examples by optimizing for adaptability rather than a single fixed objective. Given a labelled dataset D with C classes, we first partition it into

two non-overlapping subsets, $D_{\text{Meta-Train}}$ and $D_{\text{Meta-Test}}$, ensuring that classes seen during meta-training do not appear in meta testing. During the meta-training phase, over E_{train} episodes we sample a task T_i from $D_{\text{Meta-Train}}$ and split it into a small support set S_i (used to update the meta-learner f_θ via gradient descent on the loss $L(f_\theta(S_i), y_{S_i})$) and a query set Q_i (used to evaluate $L(f_\theta(Q_i), y_{Q_i})$), thereby refining the model’s few-shot learning strategy through episodic training. In the meta-testing phase, we assess generalization over E_{test} episodes by sampling tasks T_j from $D_{\text{Meta-Test}}$, fine-tuning f_θ on each support set, and measuring performance on the corresponding query set; the aggregate query loss across these episodes quantifies the model’s ability to adapt to entirely new tasks with minimal data.

2.2 Few Shot Learning

Conventional deep learning approaches typically depend on large volumes of labelled data, which becomes a significant limitation in domains such as medical imaging, where data imbalance and scarcity especially for rare classes are common. Few-Shot Learning (FSL) [20] offers an effective alternative by enabling models to perform classification tasks using only a small number of labelled instances.

In Few-Shot Learning, tasks are framed as N -way K -shot classification problems, where N indicates the number of distinct classes involved in each task and K specifies the number of available support samples per class. This episodic setting is designed to simulate the conditions under which the model will be evaluated. As described in Algorithm 1, the meta-training phase involves generating E_{Train} such tasks T_i , each structured according to the N -way K -shot format. An analogous process is followed during meta-testing, where the model is assessed on E_{Test} previously unseen tasks drawn from a separate test distribution.

2.3 Prototypical Networks

Meta-learning approaches have become key to addressing Few-Shot Learning (FSL) challenges especially in data scarce fields like medical imaging by enabling rapid adaptation from limited examples. Among these, metric-based strategies [16] excel in both simplicity and accuracy; notably, Prototypical Networks [22] perform N -way K -shot classification by embedding support set examples, computing each class prototype as the mean of its feature vectors, and assigning queries to the nearest prototype in embedding space.

Formally, given a labelled dataset $\mathcal{D} = \{(x_k, y_k)\}_{k=1}^n$ drawn from class set C , we split C into disjoint training and testing subsets C_{tr} and C_{ts} (so that $C_{\text{tr}} \cap C_{\text{ts}} = \emptyset$), thereby ensuring that during evaluation the model must generalize to entirely unseen classes using only the few support examples provided. The training procedure follows a meta-learning framework consisting of repeated episodes. During each meta-training episode, a task T_i is created by randomly sampling N classes from C_{tr} . For each class c within the task, K samples from the support set S_i^c , and Q samples from the query set Q_{ci} , both drawn from $D_{\text{Meta-Train}}$. A class prototype p_i^c for each class c is computed by averaging the feature representations of its support set:

$$p_i^c = \frac{1}{K} \sum_{x_k \in S_i^c} f_\theta(x_k) \quad (1)$$

For each query sample $(x_q, y_q) \in Q_{ci}$, the distance to each prototype is calculated. The probability that the query belongs to class c is obtained by applying a SoftMax over the negative distances:

$$p_\theta(y = c | x_q) = \frac{\exp(-d(f_\theta(x_q), p_i^c))}{\sum_{c'} \exp(-d(f_\theta(x_q), p_i^{c'}))} \quad (2)$$

Here, $d(f_\theta(x_q), p_i^c)$ denotes the distance between the query embedding $f_\theta(x_q)$ and the prototype p_i^c . The model is optimized by minimizing the negative log-likelihood of the true class label over all query samples, thereby updating the parameters θ of the embedding function f_θ .

In the meta-testing phase, the trained model is evaluated on new tasks sampled from C_{ts} , which contain previously unseen classes. This phase serves to assess how well the model generalizes beyond the training distribution.

3. PROPOSED METHODOLOGY

The core objective of this study is to eliminate redundant or noisy feature maps and preserve only those that are most informative, thereby enabling the construction of more robust class prototypes. We posit that not all extracted feature maps contribute positively to prototype formation some may introduce extraneous information that could impair predictive performance. Consider an input image $\mathbf{x} \in \mathbb{R}^{4 \times A \times B}$, where $A \times A$ denotes the spatial resolution and B the number of input channels. When processed by the encoder $f_\theta : \mathbb{R}^{4 \times A \times B} \rightarrow \mathbb{R}^{a \times a \times b}$, a pre-final feature tensor $\mathbf{z} = f_\theta(\mathbf{x}) \in \mathbb{R}^{a \times a \times b}$ is extracted where $a \times a$ is the reduced spatial resolution and b is the number of feature channels. From this tensor, we select the top M feature maps by integrating three complementary selection strategies: Global Average Pooling, Max Pooling, and Variance analysis. Global Average Pooling, Max Pooling, and Variance each capture complementary aspects of feature maps, and their integration leads to richer prototype representations. Global Average Pooling summarizes the overall activation strength across spatial locations, offering a sense of the dominant trends or global context within each channel. Max Pooling emphasizes the most strongly activated regions, often corresponding to highly discriminative or distinctive visual cues that are critical for class separation. Variance, in contrast, quantifies the degree of fluctuation within a channel, highlighting feature maps where activations exhibit strong diversity and thus capturing localized complexity or structural variation. When combined, these three statistical perspectives provide a holistic characterization of the support samples: the average conveys representativeness, the maximum emphasizes saliency, and the variance reflects diversity and richness of visual patterns. This multi-faceted view ensures that prototypes are not only broadly representative but also sensitive to fine-grained details and discriminative features qualities that are particularly essential for image-based tasks where subtle local variations often determine class identity. Let suppose, each support sample embedding from a class c is

represented as a feature map of dimension $a \times a \times b$, where b denotes the number of channels and $a \times a$ is the spatial resolution.

For each class $c \in \{1, \dots, N\}$ in an N -way classification setting, we extract feature maps for K support samples using the encoder f_θ . Each support sample of c^{th} class is represented as $\mathbf{z}_{c,j} \in \mathbb{R}^{a \times a \times b}$, where $j = 1, \dots, K$. These samples collectively form the support set S_i^c , corresponding to class c in task T_i . For each support sample $\mathbf{z}_{c,j}$, we identify the most informative feature maps by using three different channel-wise statistics: the average pooled vector $\mathbf{a}_{c,j} = \text{GlobalAvgPool}(\mathbf{z}_{c,j})$, the max pooled vector $\mathbf{m}_{c,j} = \text{MaxPool}(\mathbf{z}_{c,j})$, and the per-channel variance vector $\mathbf{v}_{c,j} = \text{Var}(\mathbf{z}_{c,j})$, all of which lie in \mathbb{R}^b . These methods are chosen because they each highlight unique aspects of the feature activations. A detailed picture of proposed feature selection strategy is depicted in Figure 1. The calculated vectors from these channels are then ranked based on descending values of $\mathbf{a}_{c,j}$, $\mathbf{m}_{c,j}$, and $\mathbf{v}_{c,j}$. A fixed number of top M unique channels (denoted as $\mathbf{I}_{c,j}$, where $|\mathbf{I}_{c,j}| = M$) are selected by interleaving the sorted indices from the three metrics. After constructing the selected index set $\mathbf{I}_{c,j}$ for a support embedding $\mathbf{z}_{c,j} \in \mathbb{R}^{a \times a \times b}$, we apply a binary channel-wise mask that preserves the chosen channels and suppresses the rest. Formally, let $\mathbf{M}_{c,j} \in \{0, 1\}^b$

be the binary mask such that $M_{c,j}(k) = 1$ if $k \in \mathbf{I}_{c,j}$ and 0 otherwise. The masked tensor retains the original shape $\mathbb{R}^{a \times a \times b}$, which simplifies batching and downstream processing. In our experiments, this zeroing operation (retaining selected channels and suppressing others) proved effective for preserving tensor shapes and enabling efficient batching. The resulting masked embedding is denoted as $\hat{\mathbf{z}}_{c,j}$. The prototype for class c in task T_i is then obtained by averaging the corresponding masked support features:

$$p_i^c = \frac{1}{K} \sum_{j=1}^{S_i^c} \hat{\mathbf{z}}_{c,j}, \quad p_i^c \in \mathbb{R}^b \quad (3)$$

Each class prototype p_i^c captures the most informative channel activations from its support samples based on average, maximum, and variance-based criteria. For a given query sample $(x_k, y_k) \in Q_i^c$, we first extract feature embedding $\hat{\mathbf{z}}_k$ for each query sample using the proposed feature selection strategy, and its distance to each class prototype is obtained by computing the corresponding Euclidean distance. The probability $p_\theta(y = c | \hat{\mathbf{z}}_k)$ that the query sample is classified into class c is then expressed using a SoftMax function applied over the negative distances:

$$p_\theta(y = c | \hat{\mathbf{z}}_k) = \frac{\exp(-d(\hat{\mathbf{z}}_k, p_i^c))}{\sum_{c'} \exp(-d(\hat{\mathbf{z}}_k, p_i^{c'}))}. \quad (4)$$

Algorithm 1: Meta-learning with Selective Feature Representation in Prototypical Networks

```

1: Input:
2: * Dataset  $\mathcal{D} = \{(x_k, y_k)\}_{k=1}^n$  with  $y_k \in \{1, \dots, C\}$ 
3: *  $C_{tr}, C_{ts}$ : disjoint train/test class splits ( $|C_{tr}| + |C_{ts}| = C$ )
4: *  $D_{\text{Meta-Train}}, D_{\text{Meta-Test}}$ : subsets of  $\mathcal{D}$  for train/test classes
5: *  $E_{\text{Train}}, E_{\text{Test}}$ : number of meta-training/testing episodes
6: *  $N$ : classes per episode
7: *  $K, N_Q$ : support/query samples per class
8: *  $M$ : number of top feature-map channels per support sample to be chosen by Global Average Pooling, Max Pooling and Variance
9: Randomly initialize parameters  $\theta$  of Meta-Learner  $f$ 
10: Meta-Training Phase: (for  $i = 1, \dots, E_{\text{Train}}$ ):
    11:   Sample  $N$  classes  $C_i \subset C_{tr}$ 
    12:   Construct support set  $S_i$  and query set  $Q_i$ 
    13:   for  $c \in C_i$  do
    14:      $D_{ic} \leftarrow \{(x, y) \in D_{\text{Meta-Train}} \mid y = c\}$ 
    15:      $S_i^c \leftarrow$  random  $K$  samples from  $D_{ic}$  ▷ Form support set pf class  $c$ 
    16:      $Q_i^c \leftarrow$  random  $N_Q$  samples from  $D_{ic} \setminus S_i^c$  ▷ Form query set pf class  $c$ 
    17:      $S_i \leftarrow S_i \cup S_i^c, \quad Q_i \leftarrow Q_i \cup Q_i^c$  ▷ Feature-map selection for each support sample
    18:     for  $\mathbf{z}_{c,j} = f_\theta(x_{c,j})$  in  $S_i^c$  do
    19:        $\mathbf{a}_{c,j} \leftarrow \text{GlobalAvgPool}(\mathbf{z}_{c,j}) \in \mathbb{R}^4$  ▷ Calculate Global Average Pooling across all channels
    20:        $\mathbf{m}_{c,j} \leftarrow \text{MaxPool}(\mathbf{z}_{c,j}) \in \mathbb{R}^4$  ▷ Calculate Max Pooling across all channels
    21:        $\mathbf{v}_{c,j} \leftarrow \text{Var}(\mathbf{z}_{c,j}) \in \mathbb{R}^4$  ▷ Calculate Variance across all channels
    22:       Sort channels by descending  $\mathbf{a}_{c,j}, \mathbf{m}_{c,j}, \mathbf{v}_{c,j}$ 
    23:       Interleave top indices to form  $\mathbf{I}_{c,j}$  with  $|\mathbf{I}_{c,j}| = M$ 
    24:        $\hat{\mathbf{z}}_{c,j} \leftarrow \text{Mask}(\mathbf{z}_{c,j}, \mathbf{I}_{c,j})$ 
    25:     end for ▷ Class Prototype with masked feature map
    26:      $p_i^c = \frac{1}{K} \sum_{j=1}^{S_i^c} \hat{\mathbf{z}}_{c,j}, \quad p_i^c \in \mathbb{R}^{B \times B \times K}$ 
    27:   end for
    28:    $\mathcal{L} \leftarrow 0$ 
    29:   for  $(x_k, y_k) \in Q_i$  do
    30:     Extract  $f_\theta(x_k)$  and compute distances  $d(f_\theta(x_k), p_i^c)$ 
    31:     Accumulate cross entropy loss into  $\mathcal{L}$ 
    32:   end for
    33:   Update  $\theta \leftarrow \theta - \alpha \nabla_\theta \mathcal{L}$  ▷ Update parameter  $\theta$  of model  $f$ 
34: Meta-Testing Phase: repeat from step 10 for  $E_{\text{Test}}$  episodes on  $D_{\text{Meta-Test}}$  to evaluate the performance of  $f_\theta$ 

```

where the SoftMax function is applied to the negative distances between the query embedding $\hat{\mathbf{z}}_k$ and the class prototypes p_i^c . Here, as mentioned in Section 2.3 also, $d(\hat{\mathbf{z}}_k, p_i^c)$ denotes the distance between the feature representation of the query sample $\hat{\mathbf{z}}_k$ and the corresponding class prototype p_i^c . The prototypical network is trained by minimizing the negative log-likelihood of the true class for each query sample, thereby learning the parameters θ of the feature extractor f_θ . A step-by-step overview of proposed approach is outlined in Algorithm 1.

4. EXPERIMENTATION DETAILS

4.1 Datasets

To ensure broad applicability, the proposed method was evaluated on two diverse medical imaging benchmarks, covering dermatology and pathology, thereby offering a comprehensive evaluation of its performance. The details of dataset used and its configuration are provided in Table 1. Some sample images from random classes of each dataset are shown in Figures 2 and 3 as a reference.

Table 1: Configuration details of datasets used in this study

Dataset Summary				
Dataset	#Images	Resolution	#Train Classes	#Test Classes
Derm7pt [6]	2,000	224×224	13	7
BloodMNIST [1]	17092	84×84	4	4

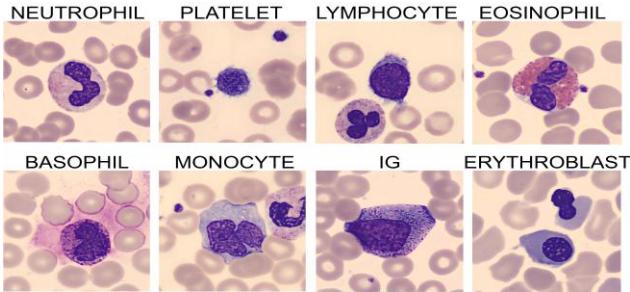


Fig. 2: Sample images from random classes of BloodMNIST dataset

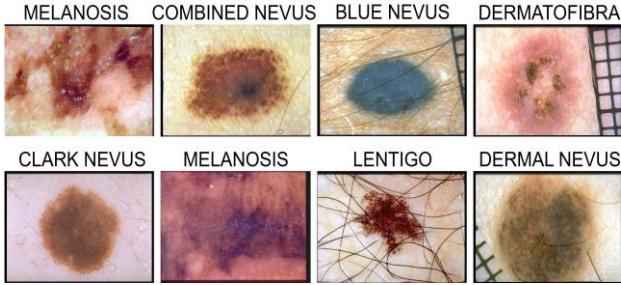


Fig. 3: Sample images from random classes of Derm7pt dataset

4.2 Implementation Details

The experiments utilize a CONV4 backbone consisting of four sequential blocks, each containing 64 filters of size 2×2 , followed by ReLU activations, batch normalization, and 2×2 max-pooling with stride 2. We train the network with Stochastic Gradient Descent (SGD), setting the learning rate to 0.1 and the momentum to 0.9, under a standard few-shot learning framework. All experiments were carried out on an NVIDIA A100-SXM4 GPU with 40 GB of memory. The implementation was done in Python 3.10.4 using the PyTorch framework.

4.3 Evaluation

Few-shot learning (FSL) performance is evaluated using classification accuracy. Meta-training consists of 2000 episodes, repeated for E_{Train} epochs. Meta-testing includes 600 tasks or randomly sampled episodes per dataset, with average accuracy reported for Derm7pt [6] and BloodMNIST [1] dataset.

5. RESULTS AND ANALYSIS

5.1 Comparison with SOTA Methods

We performed a comparative evaluation of proposed model's

performance against multiple state-of-the-art (SOTA) methods. The values in **bold** are the highest and values in underline are the second highest

5.1.1 Derm7pt. Figure 4 presents a comparative analysis of proposed method against various state-of-the-art approaches on the Derm7pt dataset. The bar charts depict 1-shot and 5-shot accuracies across multiple methods, where each bar corresponds to a reported SOTA model, and the error bars represent performance variance. As shown in the 1-shot chart, most methods cluster between 56–62%, with WRN-28-10 and DenseNet121 variants reaching around 61%, and Conv6 approaches plateauing below 63%. In contrast, proposed method's Conv4-based model (red bar) clearly rises above the group, achieving 67.18%, which is an approximate gain of 5% over the strongest WRN-28-10 baseline and more than 6% over Conv6. The 5-shot chart illustrates an overall upward shift in performance, with WRN-28-10 methods peaking at 79.83% and DenseNet121 around 79.18%, while Conv6 approaches remain below 77%. Proposed method's model maintains strong competitive performance at 78.20%, comparable to the heavy WRN-28-10 and DenseNet121 backbones.

Despite employing a lightweight Conv4 backbone with only 50,688 parameters, the proposed method consistently outperforms much larger models in the more challenging 1-shot setting while sustaining competitive results in the 5-shot scenario. Moreover, as highlighted in the Baseline [19] work, such compact architectures are advantageous for feature extraction offering faster inference, lower memory requirements, and improved generalization making them particularly well-suited for resource-constrained clinical deployment.

5.1.2 BloodMNIST. Figure 5 presents a comparative analysis of the proposed method against various state-of-the-art approaches on the BloodMNIST dataset. The bar charts show 1-shot, 5-shot, and 10-shot accuracies, with each bar representing a reported SOTA model and the error bars indicating variance. In the 1-shot setting, most approaches fall between 55–63%, with ResNet-50 methods performing strongly. Multi-ProtoNets achieves 63.10% and GDSL with ResNet-25 reaches 63.47%. The proposed Conv4-based model (red bar) delivers 63.39%, nearly matching the best heavy ResNet-50 baselines despite using only 50k parameters. In the 5-shot chart, overall performance rises, with LDP-Net (ResNet50) achieving 73.31% and Latent Augmentation reaching 70.22%. Yet, the proposed Conv4-based method clearly surpasses all, attaining 77.17% and outperforming both ResNet-50 and ResNet25 models.



Fig. 4: Performance comparison of few-shot learning methods on Derm7pt dataset for 2-way classification under 1-shot and 5-shot settings. Accuracy (%) is reported along with model backbones and parameter sizes

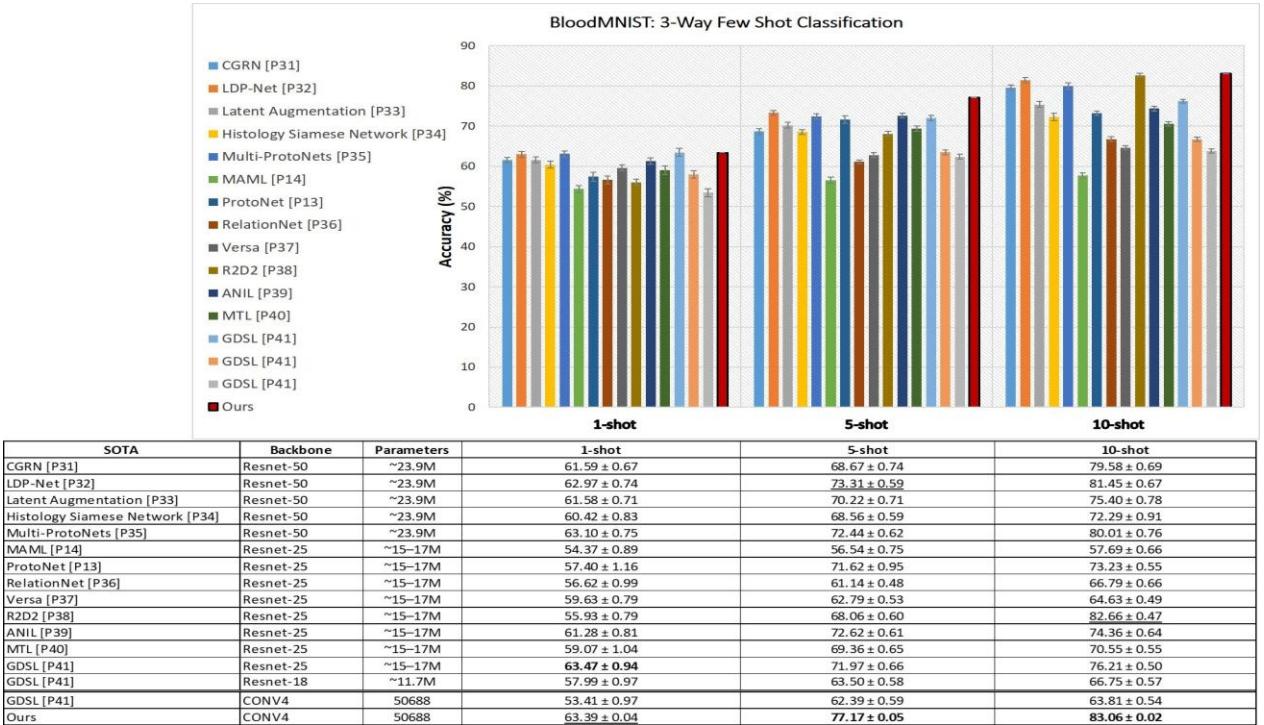


Fig. 5: Performance comparison of few-shot learning methods on the BloodMNIST dataset for 3-way classification under 1-shot, 5-shot, and 10-shot settings. Accuracy (%) is reported along with model backbones and parameter sizes

The 10-shot results highlight a similar trend: ResNet50 methods plateau around 79–80% (e.g., Multi-ProtoNets at 80.01%), while some ResNet-25 approaches such as R2D2 reach 82.66%. In contrast, the proposed Conv4 model continues to improve and achieves the best overall accuracy of 83.06%, exceeding both heavy and mid-weight backbones.

Overall, while deep models like ResNet-50 (23.9M parameters) and ResNet-25 (15–17M parameters) perform competitively, their gains diminish at higher shots. Lightweight Conv4 methods like GDSL suffer sharp drops in accuracy (53.41% at 1-shot and

only 63.81% at 10-shot). The proposed Conv4-based approach not only closes this gap but also establishes new performance highs, demonstrating that extremely compact models can exploit scarce support examples more effectively than heavier architectures, making them highly practical for resource-constrained clinical deployment.

t-SNE Visualization

To qualitatively assess the discriminative power of learned prototypes, we visualize the embedded representations of support

and query samples for a representative 3-way 10-shot task on BloodMNIST using t-SNE. As shown in Figure 6 (a) conventional Prototypical Network produces clusters that are both dispersed and partially overlapping indicating prototype ambiguity and weaker inter-class separation. In contrast, 6 (b) Selective Feature Representation in Prototypical Networks approach yields markedly more compact and well-separated clusters with each class forming a distinct, tightly grouped manifold. This clear improvement stems from proposed method's channel-wise filtering, which suppresses noisy or redundant activations and retains only the most informative feature maps for prototype construction. Consequently, the resulting embeddings exhibit sharper decision boundaries and reduced intra-class variance, providing compelling gains in few-shot accuracy achieved by proposed method.

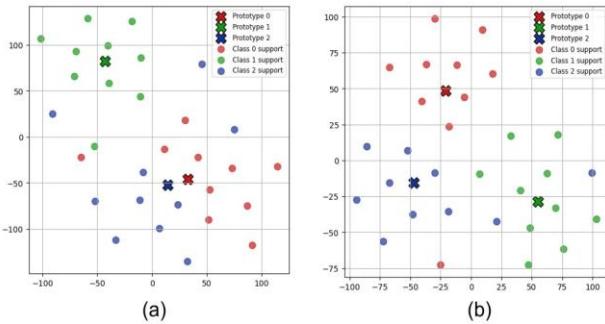


Fig. 6: t-SNE Visualization of (a) ProtoNet vs (b) Proposed approach on 3-way 10-shot classification in BloodMNIST dataset

6. CONCLUSION AND FUTURE DIRECTIONS

In this work, we introduced a lightweight yet effective refinement to prototypes of prototypical few-shot learning by Selective Feature Representation, which explicitly filters support embeddings by selecting only the top- M feature maps per class using an ensemble of Global Average Pooling, Max Pooling, and Variance criteria. Built upon a compact CONV4 backbone, proposed approach achieves state-of-the-art performance on two challenging medical imaging benchmarks, yielding 67.18% (1-shot) on Derm7pt and up to 77.17% (5-shot) and 83.06% (10-shot) on BloodMNIST, outperforming much larger architectures. Looking ahead in future, although proposed selective feature refinement strategy shows strong potential, there are several directions to further improve and broaden its impact. First, to select more varied and discriminative feature maps, future work will focus on employing additional statistical measures or alternative techniques alongside the current ensemble of Global Average Pooling, Max Pooling, and Variance. This would allow the identification of the most descriptive and representative feature maps that not only form stronger class prototypes but also highlight the most salient regions of medical images, thereby improving interpretability and classification accuracy. Second, we plan to extend this framework to a wider range of medical imaging datasets, particularly those that face resource constraints, limited annotations, and severe class imbalance. Validating proposed method's under such challenging conditions will help demonstrate its robustness and practical usefulness in real-world clinical scenarios where high-quality labelled data is scarce. Together, these directions aim to make selective feature refinement a more general, adaptive, and clinically valuable solution for few-shot medical image analysis. Overall, selective feature refinement represents a promising direction for enhancing prototype quality in few-shot medical imaging while maintaining efficiency and scalability.

7. REFERENCES

- [1] Andrea Acevedo, Anna Merino, Santiago Alferez, Angel' Molina, Laura Boldu, and Jos' e Rodellar. A dataset of microscopic peripheral blood cell images for development of automatic recognition systems. *Data in Brief*, 30:105474, 2020.
- [2] Arijit Bera, Subhransu Kumar, and Ashwin Machanavajjhala. Exploring medical image mining using deep convolutional networks. In *SIAM International Conference on Data Mining*, pages 543–551, 2018.
- [3] Runyang Cao, Yu Zhang, Song Bai, and Philip Torr. Comet: Concept-based prototypical networks for few-shot learning. In *International Joint Conference on Artificial Intelligence*, pages 2108–2114, 2021.
- [4] Yujia Chen, Zhenguo Li, and Ruijia Wang. Feature-level meta-learning with attention for few-shot classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 512–520, 2022.
- [5] Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. In *Nature*, pages 115–118, 2017.
- [6] J. Kawahara et al. Seven-point checklist and skin lesion classification. *IEEE JBHI*, 2019.
- [7] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Modelagnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135, 2017.
- [8] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. In *International Conference on Machine Learning*, pages 1180–1189, 2016.
- [9] Xia He, Xue Yuan, Kai Zhang, Hu Li, and Wei Liu. Robust re-weighting prototypical networks for few-shot classification. In *Proceedings of the 6th International Conference on Robotics and Artificial Intelligence (ICRAI '20)*, pages 125–132, Singapore, 2020. ACM.
- [10] Timothy M. Hospedales, Antreas Antoniou, Paul Micahelis, and Amos Storkey. Meta-learning in neural networks: A survey. In *International Conference on Learning Representations*, 2020.
- [11] Jieneng Hu and et al. Feature disentanglement and filtering for few-shot learning. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.
- [12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012.
- [13] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. In *Proceedings of the 28th International Conference on Neural Information Processing Systems*, pages 1–9, 2015.
- [14] Qi Li, Hao Zhang, and Wei Wang. Learning to select informative features for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11835–11844, 2023.
- [15] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi,

Arnaud A. A. Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen A. W. M. van der Laak, Bram van Ginneken, and Clara I. Sanchez. A survey on deep learning in medical image analysis. In Medical Image Computing and Computer-Assisted Intervention – MICCAI, pages 295–307, 2017.

[16] Xiaolong Liu, Yifan Wang, and Lei Zhang. Weighted prototypical networks for few-shot learning. In AAAI Conference on Artificial Intelligence, pages 9319–9326, 2019.

[17] Yun Liu, Geert Litjens, and Bram van Ginneken. Deep learning for content-based medical image retrieval: A performance study. In 2016 IEEE International Conference on Image Processing (ICIP), pages 1275–1279, 2016.

[18] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I. Jordan. Learning transferable features with deep adaptation networks. In International Conference on Machine Learning, pages 97–105, 2015.

[19] Xu Luo, Hao Wu, Ji Zhang, Lianli Gao, Jing Xu, and Jingkuan Song. A closer look at few-shot classification again. In International Conference on Machine Learning, pages 23103–23123. PMLR, 2023.

[20] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In ICLR Workshop, 2017.

[21] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. In Advances in Neural Information Processing Systems, NeurIPS 2019, pages 1917–1928, 2019.

[22] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In Advances in Neural Information Processing Systems, pages 4077–4087, 2017.

[23] Koichi Suzuki, Katsuyuki Mori, and Kunio Doi. Medical image sorting and retrieval by combining shape and texture features. In SPIE Medical Imaging, pages 978–985, 2016.

[24] Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one-shot learning. In Advances in Neural Information Processing Systems, pages 3630–3638, 2016.

[25] Yongqiang Wang, Yikang Li, Kai Chen, Shiji Song, and Ying Huang. A survey on few-shot learning. In CVPR Workshops, pages 2866–2875, 2020.

[26] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In Advances in Neural Information Processing Systems, pages 3320–3328, 2014.

[27] Zhongzheng Yue, Qianru Sun, and Qi Yu. Compact feature learning for prototypical few-shot classification. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), pages 10454–10463, 2023.