Seeing More with Less: Meta-Learning and Diffusion Models for Tumor Characterization in Low-data Settings

Eva Pachetti[1,2,✉][0000−0002−1321−9285] and Sara Colantonio[1][0000−0003−2022−0804]

[1] Institute of Information Science and Technologies, National Research Council, Pisa, Italy
eva.pachetti@isti.cnr.it
[2] Department of Information Engineering, University of Pisa, Pisa, Italy

Abstract. While deep learning excels in many areas, its application in medicine is hindered by limited data, which restricts model generalizability. Few-shot learning has emerged as a potential solution to this problem. In this work, we leverage the strengths of meta-learning, the primary framework for few-shot learning, along with diffusion-based generative models to enhance few-shot learning capabilities. We propose a novel method that jointly trains a diffusion model and a feature extractor in an episodic-based manner. The diffusion model learns conditional generation based on each episode's support samples. After updating its parameters, it generates additional support samples for each class. The augmented support set is used to train a feature extractor within a prototypical meta-learning framework. Notably, we propose a weighted prototype computation based on the distance between each generated sample and the original class prototype, i.e., derived solely from the original support samples. Evaluations on two tumor characterization tasks (prostate cancer aggressiveness and breast cancer malignity assessment) demonstrate our approach's effectiveness in improving prototype representation and boosting classification performance. Find our code at: https://github.com/evapachetti/meta_diffusion.

Keywords: Few-shot learning · Meta-learning · Diffusion models

1 Introduction

Deep learning (DL) models in medical imaging hold significant promise for clinical applications. However, insufficient data makes the training phase challenging, often resulting in inaccurate and unreliable models for real-world use. A promising solution to this data scarcity problem is few-shot learning (FSL) [5,6], which aims to enable practical DL model training in data-limited scenarios, such as medical imaging [2,13,28]. One of the most popular frameworks for addressing FSL is meta-learning. Meta-learning involves two main objectives: an inner objective related to a specific classification task (or *episode*) and an outer objective

where the model learns to distinguish data more generally. This approach, also known as *episodic* training, enhances the model's ability to learn robust features, improving generalization, an essential property in data-scarce domains.

Within meta-learning, a popular technique is the prototype-based approach [19]. In each episode, the idea is to calculate a prototype for each class based on the few training examples available in that episode (known as the *support* set). New images within the episode (*query* samples) are then classified by measuring their distance to each class prototype. Various techniques have been proposed to enhance prototype informativeness. Cao et al. [1] leveraged the similarity between classes to calibrate prototypes of new classes using base classes learned beforehand. He et al. [9] introduced a transformer-based module to extract more informative prototypes. Zhang et al. [26] tackled the problem of low-informative prototypes utilizing prior knowledge and pre-trained features obtained from complete prototypes of well-represented classes.

Another strategy employed in literature to enhance prototype informativeness involves using generative models. Zhang et al. [27] proposed a prototype meta-hallucination approach to generate more informative prototypes by hallucinating additional support samples. Specifically, they meta-trained a Variational Autoencoder (VAE) [16] to learn the distribution of inter-sample differences and synthesized newly labeled samples by fusing the sampled inter-sample difference and each given support sample. However, VAEs often struggle with capturing the full complexity of images due to their reliance on a Gaussian-distributed latent space, leading to blurry or imprecise image generation [3,12] and, thus, to less meaningful prototypes.

In recent years, the emergence of diffusion-based models has resulted in generative models exhibiting unprecedented capabilities [15,17]. Indeed, diffusion models overcame the limitations provided by generative models such as VAEs and Generative Adversarial Networks (GANs) [7], including the need for estimation of intractable the normalizing constant of the probability function, the need for network constraints, and the training instability, leading to the generation of more realistic and informative samples that closely resemble actual data.

In this work, we leverage the power of Denoising Diffusion Probabilistic Models (DDPMs) [20,10] to enhance prototype informativeness in a prototypical meta-learning framework for few-shot medical image classification. Specifically, we propose a training method integrating real and synthetic data within a joint episodic training process between the DDPM and a feature extractor. Support samples are provided to train a conditional DDPM during each episode, which then hallucinates additional data samples for each class. These synthetic samples supplement the limited actual data available for prototype construction. Furthermore, we implement a dynamic weighting mechanism to ensure the synthetic samples enhance rather than confuse the model. This method calculates weights for the synthetic samples based on their proximity to the original data-derived prototype. By integrating actual and synthetic data with dynamic weighting, we aim to forge more reliable prototypes, thereby enhancing classification performance in low-data scenarios.

2 Methods

2.1 Proposal

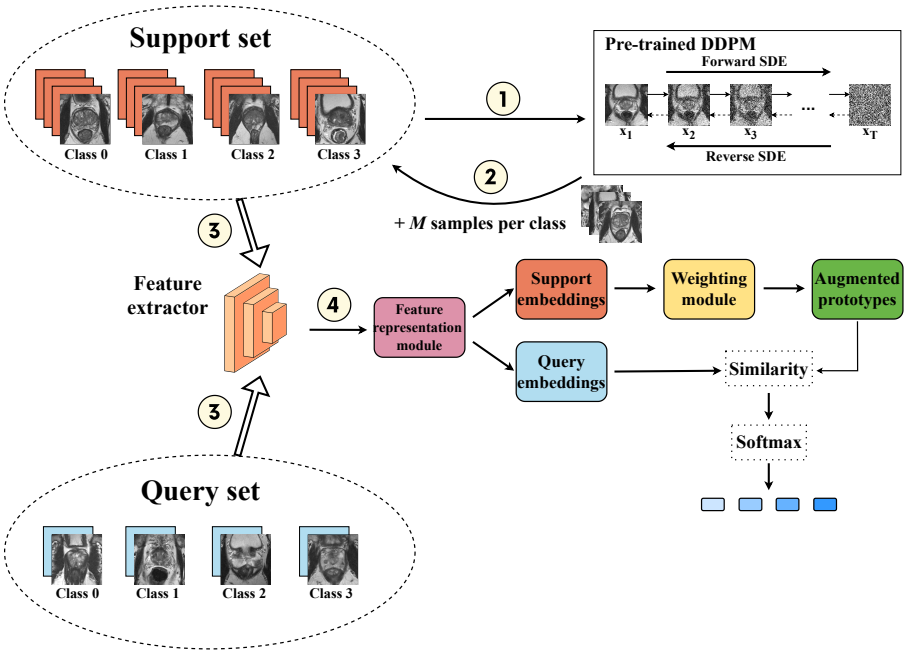

Fig. 1: Illustration of the proposed approach for a single episode. Each episode consists of a support set and a query set. All support samples are provided to the DDPM for conditional training, guided by the class embedding. The DDPM parameters are updated to generate M additional support samples for each class. The augmented support and query sets are fed through a feature extractor to compute feature maps. These maps are then processed by a dedicated feature representation module, whose function depends on the chosen prototypical framework (e.g., average pooling for ProtoNet). The weighting module (Fig. 2) adjusts the augmented support embeddings based on their similarity to the original prototype. The original and generated support embeddings, weighted accordingly, are averaged to compute the augmented prototype for each class. A distance calculation between query embeddings and augmented prototypes follows. Finally, classification is achieved using a softmax distribution over these computed similarities.

We provide a detailed overview of our proposed strategy in Fig. 1. As a first step, to provide the DDPM with prior knowledge, we pre-train it on an unlabeled dataset to perform unconditioned generation. This step aims to equip

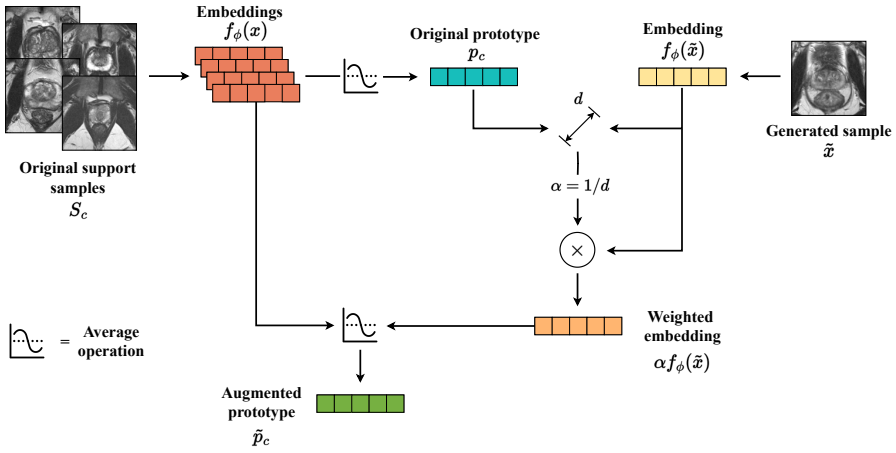

Fig. 2: Representation of the weighting module for augmented prototype estimation. For clarity, we depicted the weighting module functioning for a single class and a single generated sample.

the model with an understanding of the invariant features in the anatomical structures regardless of the specific class. Following this pre-training step, we perform a joint episodic training between the DDPM and a convolutional feature extractor. Within each episode, consisting of a support set (containing N images per class) and a query set, we utilize the entire support set as input to the pre-trained DDPM to train it on performing conditioned generation leveraging the class embedding, as proposed by Ho et al. [11]. After updating its parameters, we employ the DDPM to generate M additional support samples for each class. These synthetic samples are added to the support set and used to train the feature extractor.

Especially during the initial training phase, the DDPM's synthetic samples may not match the quality and detail of the real support images. This discrepancy can result in less accurate prototypes. To address this challenge, we propose a weighting module (see Fig. 2) that implements a dynamic weighting approach when constructing prototypes. We define $S_c = \{(\boldsymbol{x}_1, y_1), \ldots, (\boldsymbol{x}_N, y_N)\}$ as the support set of class c, containing only the original support samples, and $\widetilde{S}_c = \{(\tilde{\boldsymbol{x}}_1, y_1), \ldots, (\tilde{\boldsymbol{x}}_M, y_M)\}$ the support set containing only generated samples of that class. Given a feature extractor f_ϕ, we calculate the real prototype, i.e., built on the original support samples only as follows:

$$\boldsymbol{p}_c = \frac{\sum_{i=1}^{|S_c|} f_\phi(\boldsymbol{x}_i)}{|S_c|}. \tag{1}$$

The weight for a generated support sample $\tilde{\boldsymbol{x}}_j^c$ of class c is provided by calculating the reciprocal of the Euclidean distance between the generated support sample

embedding $f_\phi(\tilde{\boldsymbol{x}}_j^c)$ and the class prototype computed on the original support samples $\boldsymbol{p}_c$:

$$\alpha_j^c = \frac{1}{d(f_\phi(\tilde{\boldsymbol{x}}_j^c), \boldsymbol{p}_c))}. \tag{2}$$

Finally, the prototype for the class c is calculated as follows:

$$\tilde{\boldsymbol{p}}_c = \frac{\sum_{i=1}^{|S_c|} f_\phi(\boldsymbol{x}_i) + \sum_{j=1}^{|\tilde{S}_c|} \alpha_j^c f_\phi(\tilde{\boldsymbol{x}}_j)}{|S_c| + \sum_{j=1}^{|\tilde{S}_c|} \alpha_j^c}. \tag{3}$$

After calculating the final prototypes for each class, we measure the Euclidean distance between each query sample and the corresponding prototypes. Finally, we apply the softmax function to these distances to predict the class labels for the query samples.

2.2 Model details

Generative model. Traditional DDPMs corrupt data through a finite number of noise steps and train a sequence of probabilistic models to reverse each step of this noise corruption. Unlike the classical approach, we adopt a continuous diffusion approach proposed by Song et al. [21]. This method employs a continuous noise-perturbation process modeled by stochastic differential equations (SDEs). A reverse-time SDE describes the inverse diffusion process from noise to image, and samples can be generated by solving this equation using numerical SDE solvers [21]. In our experiments, we described the diffusion process using a sub-variance-preserving SDE [21] and employed the Euler-Maruyama numerical solver [14] for image generation.

Prototypical framework. We evaluated our approach leveraging three well-known prototypical meta-learning frameworks: Prototypical Network (ProtoNet) [19], Meta Deep Browninan Distance Covariance (Meta DeepBDC) [24] and Covariance Network (CovNet) [23]. These methods all share the common principle of constructing a class prototype by averaging the embeddings of the support samples belonging to that class. Classification is then performed by measuring the distance between the embedding of each query sample and all class prototypes. However, they diverge in their approach to feature representation: ProtoNet utilizes a basic first-order representation via average pooling, CovNet leverages a second-order representation with the covariance matrix, and Meta DeepBDC goes beyond pairwise relationships by considering the joint distribution between features through the computation of the BDC matrix.

2.3 Experiments

We evaluated our approach on two tumor characterization tasks: prostate cancer aggressiveness (PI-CAI dataset [18]) and breast cancer lesion classification

(BreakHis dataset [22]). For prostate cancer, we classified tumors based on four (2-5) prognostic scores defined by the International Society of Urological Pathology (ISUP) [4]. We pre-trained our DDPM on all available benign lesions (11202 images) from PI-CAI dataset. Then, we performed episodic supervised learning using all cancerous lesions (1611 training, 200 validation, 238 testing from a total of 2049). In breast cancer classification, we used images from various magnifications (40X, 100X, 200X; 6090 images total) for DDPM unconditioned pre-training and focused on 400X magnification images (1475 training, 165 validation, 183 testing from a total of 1819) for episodic training, performing binary classification (benign vs. malignant).

We evaluated our approach using *k-shot* classification tasks ($k \in \{1, 2, 3, 4, 5\}$) for each prototypical framework. We investigated the effect of generating one, two or three synthetic support samples per class to improve prototype informativeness. The training process involved 100 epochs, each with 50 meta-training and 50 meta-validation episodes. We assessed performance using mean and standard deviation (STD) of Area Under the ROC Curve (AUROC) across 50 meta-testing episodes. We employed a ResNet-18 for feature extraction, with a learning rate of 10^{-4} and weight decay of 10^{-2}, trained using AUC margin loss [25] and the Proximal Epoch Stochastic optimizer [8]. For the DDPM, we utilized a learning rate of 20^{-4} and 400000 steps for training, with 1000 noise scales (variance from 0.1 to 20) during image generation. The training phase was conducted using the negative log-likelihood loss function and the Adam optimizer.

3 Results and Discussion

Our experimental results (Table 1 and Table 2) demonstrate that generating synthetic support samples improves performance (AUROC) across all three prototypical methods for both classification tasks. On the prostate cancer aggressiveness classification task, ProtoNet benefits the most, achieving an AUROC increase of over 11% (from 0.634 to 0.749) in the *5-shot* setting with two synthetic samples per class. On the other hand, for breast lesion classification, CovNet exhibits the most considerable absolute improvement (15%, from 0.514 to 0.664) in the *4-shot* setting with three synthetic samples, even though ProtoNet achieves the highest AUROC (0.785) with three synthetic samples in the *5-shot* setting.

Figure 3 visually explores the relationship between the number of synthetic support samples and AUROC performance. Interestingly, the results indicate that performance gains are not always monotonic. This implies that adding more synthetic samples may not always yield the best results. For example, the CovNet model on the BreakHis dataset achieves its peak AUROC with a single synthetic support sample in the *1-shot* setting. Adding more synthetic samples, in this case, actually leads to decreased performance. This behavior may be due to low-quality image generation (e.g., images lacking features of the desired class), which the weighting method failed to mitigate. However, this phenomenon is also evident in baseline configurations (without adding synthetic data), where increasing the number of support samples (e.g., from *1-shot* to *2-shot*) does

Table 1: Mean and STD (in brackets) of AUROC across 50 meta-test episodes on the PI-CAI dataset (prostate cancer aggressivenes classification).

Framework	K-shot	Baseline	+1 support	+2 support	+3 support
ProtoNet	*1-shot*	0.527 (0.067)	0.578 (0.087)	0.564 (0.065)	**0.628 (0.094)**
	2-shot	0.570 (0.080)	0.599 (0.088)	0.597 (0.072)	**0.631 (0.066)**
	3-shot	0.590 (0.074)	0.595 (0.133)	0.610 (0.104)	**0.634 (0.067)**
	4-shot	0.609 (0.067)	0.594 (0.084)	0.677 (0.088)	**0.698 (0.083)**
	5-shot	0.634 (0.095)	0.678 (0.091)	**0.749 (0.089)**	0.616 (0.087)
Meta DeepBDC	*1-shot*	0.579 (0.061)	0.580 (0.058)	**0.660 (0.124)**	0.658 (0.084)
	2-shot	0.600 (0.062)	0.598 (0.049)	0.563 (0.094)	**0.660 (0.046)**
	3-shot	0.595 (0.068)	0.627 (0.070)	0.637 (0.082)	**0.646 (0.090)**
	4-shot	0.612 (0.044)	0.538 (0.093)	0.656 (0.065)	**0.664 (0.077)**
	5-shot	0.632 (0.099)	0.646 (0.076)	0.662 (0.039)	**0.667 (0.073)**
CovNet	*1-shot*	0.503 (0.041)	0.527 (0.057)	0.552 (0.064)	**0.570 (0.048)**
	2-shot	0.529 (0.045)	0.558 (0.052)	0.560 (0.080)	**0.585 (0.059)**
	3-shot	0.575 (0.075)	0.588 (0.043)	**0.652 (0.074)**	0.619 (0.065)
	4-shot	0.580 (0.032)	0.599 (0.081)	0.619 (0.087)	**0.623 (0.068)**
	5-shot	0.602 (0.053)	0.610 (0.044)	0.626 (0.083)	**0.631 (0.080)**

Table 2: Mean and STD (in brackets) of AUROC across 50 meta-test episodes on the BreakHis dataset (breast lesion classification).

Framework	K-shot	Baseline	+1 support	+2 support	+3 support
ProtoNet	*1-shot*	0.588 (0.197)	0.572 (0.163)	0.618 (0.212)	**0.655 (0.205)**
	2-shot	0.651 (0.186)	0.674 (0.252)	0.678 (0.160)	**0.762 (0.151)**
	3-shot	0.644 (0.176)	0.679 (0.210)	**0.718 (0.163)**	0.698 (0.178)
	4-shot	0.665 (0.163)	0.679 (0.199)	0.740 (0.165)	**0.762 (0.140)**
	5-shot	0.724 (0.151)	0.732 (0.145)	0.778 (0.159)	**0.785 (0.158)**
Meta DeepBDC	*1-shot*	0.524 (0.194)	0.514 (0.212)	**0.544 (0.208)**	0.539 (0.216)
	2-shot	0.524 (0.204)	0.536 (0.213)	0.562 (0.198)	**0.576 (0.198)**
	3-shot	0.589 (0.187)	0.591 (0.179)	0.609 (0.196)	**0.633 (0.170)**
	4-shot	0.534 (0.181)	0.537 (0.205)	0.575 (0.203)	**0.632 (0.178)**
	5-shot	0.594 (0.181)	0.625 (0.176)	0.578 (0.190)	**0.666 (0.180)**
CovNet	*1-shot*	0.506 (0.143)	**0.536 (0.131)**	0.495 (0.226)	0.534 (0.199)
	2-shot	0.516 (0.103)	0.560 (0.179)	0.590 (0.187)	**0.618 (0.191)**
	3-shot	0.512 (0.062)	0.513 (0.083)	0.579 (0.140)	**0.643 (0.224)**
	4-shot	0.514 (0.154)	0.578 (0.175)	0.611 (0.158)	**0.664 (0.175)**
	5-shot	0.532 (0.103)	0.549 (0.084)	0.589 (0.168)	**0.642 (0.187)**

not always lead to improved performance. For example, in Table 2, the Meta DeepBDC model demonstrates a decrease in AUROC of approximately 5%when transitioning from a *3-shot* to a *4-shot* setting. This could be due to redundant information or poor-quality samples.

To qualitatively assess our proposed approach's performance, we compare real and synthetic MRI images of prostate cancer in Figure 4, representing each classification category. To the untrained eye, the synthetic images closely resem-

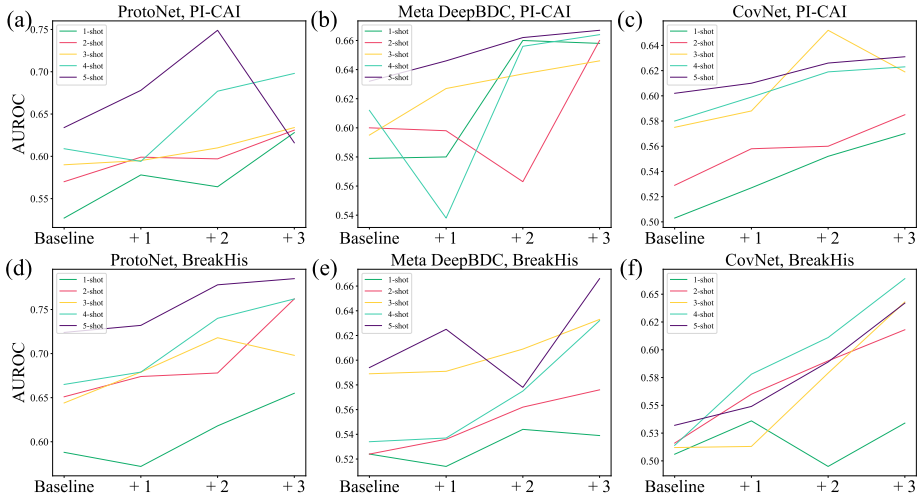

Fig. 3: Illustration of classification performance in terms of AUROC as the number of synthetic support samples per class increases. Best viewed in color. (a) ProtoNet with PI-CAI; (b) Meta DeepBDC with PI-CAI; (c) CovNet with PI-CAI; (d) ProtoNet with BreakHis; (e) Meta DeepBDC with BreakHis; (f) CovNet with BreakHis.

ble the shape and features of the prostate, though they tend to show slightly more noise compared to real images. A comprehensive evaluation by a radiology expert is necessary to verify the presence of features associated with each classification. Nonetheless, our preliminary findings suggest that incorporating these synthetic images into the class prototype computation enhances classification performance, positioning our approach as a promising task-agnostic method for tumor characterization in data-limited scenarios. In future work, we will investigate how improvements in image quality and feedback from radiology experts can further boost the effectiveness of our approach.

4 Conclusion

In this work, we combined the strengths of meta-learning and diffusion-based generative models to tackle tumor characterization in low-data scenarios. Our novel approach involves jointly training a diffusion model and a feature extractor within an episodic framework, utilizing the generation of synthetic support samples to enhance the creation of more informative prototypes. Additionally, we introduced a dynamic weighting mechanism that adjusts based on the similarity between generated samples and the original prototype of the same class, helping to mitigate the impact of generating poor or uninformative images, particularly during the early training stages. Preliminary experiments across various prototypical frameworks and two tumor characterization tasks demonstrate the ef-

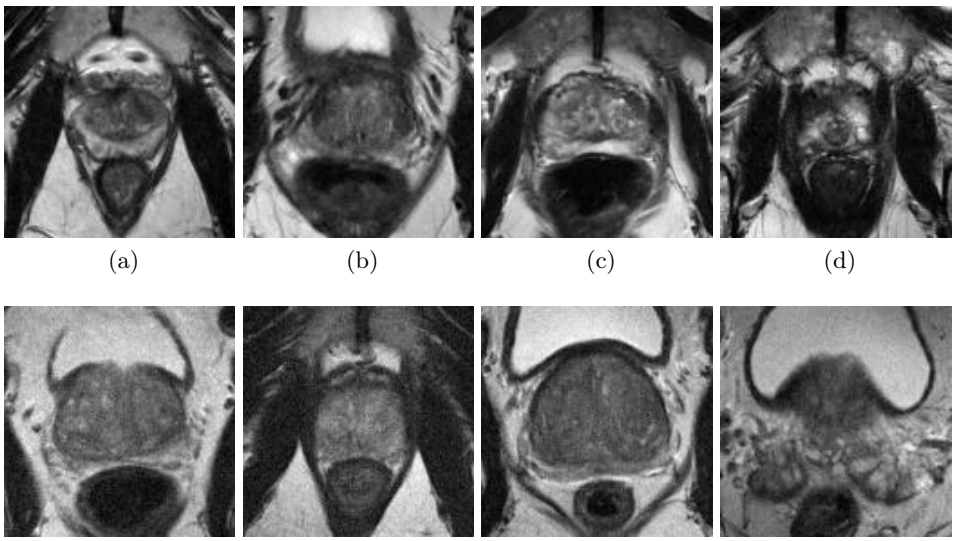

Fig. 4: Examples of real and synthetic MRI images of prostate cancer categorized by ISUP grade. (a) ISUP 2 real image; (b) ISUP 3 real image; (c) ISUP 4 real image; (d) ISUP 5 real image; (e) ISUP 2 synthetic image; (f) ISUP 3 synthetic image; (g) ISUP 4 synthetic image; (h) ISUP 5 synthetic image.

fectiveness of our method in improving classification performance. Future work will focus on enhancing image generation quality and incorporating feedback from radiology experts to better assess the presence of class-related features in synthetic images.

Acknowledgments. This study has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 952159 (ProCAncer-I) and from the Regional Project PAR FAS Tuscany—NAVIGATOR. The funders had no role in the design of the study; the collection, analysis, and interpretation of data; or writing the manuscript.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Cao, J., Yao, Z., Yu, L., Ling, B.W.K.: Wpe: Weighted prototype estimation for few-shot learning. Image and Vision Computing **137**, 104757 (2023). https://doi.org/10.1016/j.imavis.2023.104757
2. Dai, Z., Yi, J., Yan, L., Xu, Q., Hu, L., Zhang, Q., Li, J., Wang, G.: Pfemed: Few-shot medical image classification using prior guided feature enhancement. Pattern Recognition **134**, 109108 (2023)

3. Dosovitskiy, A., Brox, T.: Generating images with perceptual similarity metrics based on deep networks. Advances in neural information processing systems **29** (2016)
4. Egevad, L., Delahunt, B., Srigley, J.R., Samaratunga, H.: International society of urological pathology (isup) grading of prostate cancer–an isup consensus on contemporary grading (2016). https://doi.org/10.1111/apm.12533
5. Fei-Fei, L., Fergus, R., Perona, P.: One-shot learning of object categories. IEEE transactions on pattern analysis and machine intelligence **28**(4), 594–611 (2006)
6. Fink, M.: Object classification from a single example utilizing class relevance metrics. Advances in neural information processing systems **17** (2004)
7. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in neural information processing systems **27** (2014)
8. Guo, Z., Yan, Y., Yuan, Z., Yang, T.: Fast objective & duality gap convergence for non-convex strongly-concave min-max problems with pl condition. Journal of Machine Learning Research **24**, 1–63 (2023)
9. He, F., Li, G., Si, L., Yan, L., Li, F., Sun, F.: Prototypeformer: Learning to explore prototype relationships for few-shot image classification. arXiv (2023). https://doi.org/arXiv:2310.03517
10. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) Advances in Neural Information Processing Systems. vol. 33, pp. 6840–6851. Curran Associates, Inc. (2020)
11. Ho, J., Salimans, T.: Classifier-free diffusion guidance. In: NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications (2021)
12. Jabbar, A., Li, X., Omar, B.: A survey on generative adversarial networks: Variants, applications, and training. ACM Computing Surveys (CSUR) **54**(8), 1–49 (2021)
13. Jiang, H., Gao, M., Li, H., Jin, R., Miao, H., Liu, J.: Multi-learner based deep meta-learning for few-shot medical image classification. IEEE Journal of Biomedical and Health Informatics **27**(1), 17–28 (2022)
14. Kloeden, P.E., Platen, E., Kloeden, P.E., Platen, E.: Stochastic differential equations. Springer (1992)
15. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: International conference on machine learning. pp. 8821–8831. Pmlr (2021)
16. Rezende, D.J., Mohamed, S., Wierstra, D.: Stochastic backpropagation and approximate inference in deep generative models. In: International conference on machine learning. pp. 1278–1286. PMLR (2014)
17. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models (2021)
18. Saha, A., Bosma, J., Twilt, J., van Ginneken, B., Yakar, D., Elschot, M., Veltman, J., Fütterer, J., de Rooij, M., et al.: Artificial intelligence and radiologists at prostate cancer detection in mri—the pi-cai challenge. In: Medical Imaging with Deep Learning, short paper track (2023)
19. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 30. Curran Associates, Inc. (2017)
20. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: Bach, F., Blei, D. (eds.)

Proceedings of the 32nd International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 37, pp. 2256–2265. PMLR, Lille, France (07–09 Jul 2015)
21. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: International Conference on Learning Representations (2020)
22. Spanhol, F.A., Oliveira, L.S., Petitjean, C., Heutte, L.: A dataset for breast cancer histopathological image classification. Ieee transactions on biomedical engineering **63**(7), 1455–1462 (2015). https://doi.org/10.1109/TBME.2015.2496264
23. Wertheimer, D., Hariharan, B.: Few-shot learning with localization in realistic settings. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6558–6567 (2019)
24. Xie, J., Long, F., Lv, J., Wang, Q., Li, P.: Joint distribution matters: Deep brownian distance covariance for few-shot classification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 7972–7981 (2022)
25. Yuan, Z., Yan, Y., Sonka, M., Yang, T.: Large-scale robust deep auc maximization: A new surrogate loss and empirical studies on medical image classification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3040–3049 (2021). https://doi.org/10.1109/ICCV48922.2021.00303
26. Zhang, B., Li, X., Ye, Y., Huang, Z., Zhang, L.: Prototype completion with primitive knowledge for few-shot learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3754–3762 (2021)
27. Zhang, L., Zhou, F., Wei, W., Zhang, Y.: Meta-hallucinating prototype for few-shot learning promotion. Pattern Recognition **136**, 109235 (2023). https://doi.org/10.1016/j.patcog.2022.109235
28. Zhu, Y., Cheng, Z., Wang, S., Zhang, H.: Learning de-biased prototypes for few-shot medical image segmentation. Pattern Recognition Letters (2024)