


Attention-based Dynamic Subspace Learners for Medical Image Analysis

Sukesh Adiga V , Jose Dolz , and Herve Lombaert

Abstract—Learning similarity is a key aspect in medical image analysis, particularly in recommendation systems or in uncovering the interpretation of anatomical data in images. Most existing methods learn such similarities in the embedding space over image sets using a single metric learner. Images, however, have a variety of object attributes such as color, shape, or artifacts. Encoding such attributes using a single metric learner is inadequate and may fail to generalize. Instead, multiple learners could focus on separate aspects of these attributes in subspaces of an overarching embedding. This, however, implies the number of learners to be found empirically for each new dataset. This work, Dynamic Subspace Learners, proposes to dynamically exploit multiple learners by removing the need of knowing *a priori* the number of learners and aggregating new subspace learners during training. Furthermore, the visual interpretability of such subspace learning is enforced by integrating an attention module into our method. This integrated attention mechanism provides a visual insight of discriminative image features that contribute to the clustering of image sets and a visual explanation of the embedding features. The benefits of our attention-based dynamic subspace learners are evaluated in the application of image clustering, image retrieval, and weakly supervised segmentation. Our method achieves competitive results with the performances of multiple learners baselines and significantly outperforms the classification network in terms of clustering and retrieval scores on three different public benchmark datasets. Moreover, our method also provides an attention map generated directly during inference to illustrate the visual interpretability of the embedding features. These attention maps offer a proxy-labels, which improves the segmentation accuracy up to 15% in Dice scores when compared to state-of-the-art interpretation techniques.

Index Terms—Deep Metric Learning, Clustering, Image Retrieval, and Weakly Supervised Segmentation

I. INTRODUCTION

Learning the similarity between arbitrary images is a fundamental problem in many key areas of computer vision such as image retrieval [1], [2], [3], recommender system [4], [5], duplicate detection [6], clustering [7], or zero-shot learning [8]. In this context, metric learning is commonly used for measuring similarities by learning a distance function over objects [9], [10]. Recently, deep metric learning (DML) has been raised as a powerful approach to learn these similarities [11]. More

specifically, the goal of DML is to learn an embedding space where images from the same classes are encouraged to be close to one another. In contrast, images belonging to different classes are pushed away in the embedding space. In recent DML approaches, the loss function can be typically expressed in Euclidean distances or cosine similarities between pairs or tuples of images in the embedding space. Well-known losses employed in DML include: contrastive loss [12], triplet loss [13], lifted structure loss [14], N-pairs loss [1], margin loss [15], angular loss [16], or ProxyNCA loss [2]. In addition to novel learning objectives, recent efforts are also devoted to designing efficient sample-mining [15], or sample weighting [17] strategies.

Most of these methods learn the embedding mapping function with a single metric learner. However, medical images have complex distributions consisting of different object attributes such as color, shape, size, or artifacts. Thus, learning the complex similarity associated with these different object attributes may be inadequate with only one single-learner. A few attempts have been made towards leveraging multiple metric learners to address this complexity [18], [19], [20]. For example, Kim *et al.* [19] ensemble multiple learners, whereas a *divide-and-conquer* strategy is used in [20] by splitting the manifold into several embedding subspaces. One main limitation of these approaches is a need to empirically find the optimal number of learners, which requires a new validation for every new setting, including every use of a new dataset. Furthermore, the sizes of the embedding subspaces associated with each learner might differ since learning the various sets of object attributes requires varying degrees of modeling complexity.

Despite the popularity of DML, surprisingly few works attempt to visually explain which regions contribute to the similarity between images in embedding networks [21]. These visualizations are of pivotal importance since they provide an efficient mechanism to understand the predictions of the model. Recent efforts have been devoted to the interpretability of deep neural networks, resulting in a variety of different approaches [22], [23], [24], [25], [26]. Among these methods, GradCAM [24] has been widely employed to explain deep classification models. This method uses gradients to highlight the discriminative regions of an image. Nevertheless, since the gradients are not available during testing, directly applying this strategy in embedding networks is not feasible [27]. Integrating interpretability in embedding networks requires either attaching an additional classification branch [28] or employing

All authors are with the Department of Computer and Software Engineering at ETS Montreal, Montreal, QC H3C 1K3, Canada. (Corresponding author: Sukesh Adiga V, Email: sukesh.adiga-vasudeva.1@ens.etsmtl.ca.)

multiple images simultaneously [29], [30]. Needless to say, interpretability is of particular interest in medical imaging, as visual explanations of predictions directly impact the diagnosis, therapy planning, and follow-up of many diseases. Thus, existing DML approaches may be inadequate to visually uncover what constitutes similarities among a complex set of medical images.

Motivated by these gaps and the scarcity of the DML literature in medical imaging, we propose a novel attention-based dynamic subspace learners approach. The underlying metric learning method is inspired by the idea of a *divide-and-conquer* strategy. More specifically, we propose to follow the approach of [20] in order to capture different object attributes, each of them processed with an independent subspace learner. These subspace learners having variable sizes are learned dynamically as and when the network accuracy is plateauing during training. Thereby avoids the need to find *a priori* the number of subspace learners while retaining the state-of-the-art performance. Furthermore, the visual interpretation of the embedding is addressed by integrating an attention module after feature extraction layers, encouraging the learners to focus on the discriminative areas of target objects. Consequently, the learning process provides a visual insight of which image region considerably contributes to the clustering of image sets in the form of pixel-wise interpretable predictions.

Our Contribution: We contribute a novel approach to the state-of-the-art method in deep metric learning and illustrate its application in medical image analysis. More precisely, we propose a training strategy that (i) explores the dynamic learning of an embedding, (ii) overcomes the empirical search of an optimal number of subspaces in approaches based on multiple metric learners, and (iii) produces compact subspaces of variable size to attend different object attributes. Furthermore, the integration of an attention module in our dynamic learner approach focuses the attention of each independent learner on the discriminative regions of an object of interest. This attention mechanism provides the added benefit of visually interpreting relevant embedded features. The evaluation of our proposed method is conducted by extensive experiments on three publicly available benchmarks: ISIC19 [31], [32], MURA [33], and HyperKvasir [34]. The performance is evaluated on clustering and image retrieval tasks, showing that the proposed method achieves competitive results with the state-of-the-art without requiring the grid searches over optimal numbers of learners. We also demonstrate that the attention maps produced by our method can be used as proxy-labels to train deep segmentation models. In particular, we evaluate our approach on ISIC18 [35], [31] in a weakly supervised segmentation task and show improvements to the visual attention and class activation maps obtained from recent state-of-the-art methods, including the method specifically designed for skin lesion detection [36].

II. RELATED WORK

A. Deep Metric Learning

Metric learning is a widely explored research field in the learning community [37], [9]. The seminal work of Siamese

Networks [37] represents the first attempt to use neural networks for feature embedding. Its concept is to employ two identical neural networks that learn a contrastive embedding from a pair of images. With the advent of deep learning, deep metric learning (DML) has gained popularity, becoming a mainstay in many modern computer vision problems, such as image retrieval [38], person re-identification [39], or few-shot learning [40]. In DML, the images are mapped into a manifold space via deep neural networks. Euclidean or cosine distances can then be directly used as a metric distance between two images in this mapped space. Typical losses employed in DML include contrastive [12] or triplet loss [41]. The contrastive loss [12] encourages images from the same class to stay closer –in the learned manifold– while pushing away samples from different classes, which should be separated by a given fixed distance. Nevertheless, forcing the same distance for all pairs of images can discourage any potential distortion in the embedded space. In contrast, this assumption is relaxed in triplet loss [12], which only imposes that negative pairs of images should be further away than positive pairs.

In the same direction as our work, [19] and [20] have leveraged the use of multiple learners to diversify the learning space towards different object attributes. While [19] propose an ensemble of multiple learners driven by attention, a *divide and conquer* strategy is employed in [20], which promotes the discovery of multiple subspaces. For example, Sanakoyeu *et al.* [20] explicitly splits the embedded space into a predefined number of learners with fixed size subspaces. Then, each learner independently learns a part of an embedding space, i.e., a subspace, from a portion of clustered data, and the final embedding is later refined from multiple learners. Even though this strategy leads to improvements over its single-learner counterpart, a grid search is needed to find an optimal number of learners with each new dataset. Furthermore, the size of the embedding space is uniform across the learners, whereas some attributes, such as color, might require smaller embeddings to encode the information than other attributes, such as shape.

B. Metric Learning in Medical Image Analysis

Despite the interest in other domains, metric learning, and more particularly DML, remains almost unexplored in medical imaging. In the pre-deep learning era, related work includes [42], which employed a distance metric learning in a traditional boosting framework in a medical image retrieval scenario. More recently, [43] investigates the use of DML to model the similarity relationship between lesions in the context of radiology images, where a triplet loss is employed to learn the lesion embeddings. Gupta *et al.* [44] also resorts to the triplet loss to learn the underlying manifold space for the task of Mitotic classification, whose embedded features are subsequently used as input for a Support Vector Machine classifier. Recently, a combination of cross-entropy loss with a contrastive loss or triplet loss is used to classify whole slide images in digital pathology [45], [46]. In [47], a triplet loss is used to learn a representation of source domain images, which is later used for target domain classification under

the few-shot learning paradigm. In [48], DML is used to pre-train a model in the application of digital pathology classification, where authors use a ProxyNCA loss for learning transferable features. To enhance the embedding, [49], [50] has integrated a multi-similarity loss to DML in the context of chest radiography and liver histopathology image, respectively. Nevertheless, most of these methods are developed with the goal of classification tasks and do not effectively leverage the geometrical information of the underlying embedding space.

C. Weakly Supervised Segmentation

Weakly supervised segmentation (WSS) has emerged as an alternative to alleviate the need for large amounts of *pixel-level* labelled data. These labels can come in the form of *image-level* labels [51], scribbles [52], points [53], bounding boxes [54] or direct losses [55]. Among them, image-level labels are easier and inexpensive to obtain [53]. Particularly, class activation maps (CAM) [56] have gained popularity in identifying saliency regions based on image labels. It is achieved by associating feature maps of the last layers and weighting their activation using a global average pooling (GAP) layer. However, generated saliency maps are typically spread around the target object, only focusing on the most discriminant areas. This limits its usability as pixel-level supervision for semantic segmentation. To enhance the generated saliency regions, some alternatives based on back-propagation (GradCAM [24]) or super-pixels (SP-CAM [57]) have been proposed. Nevertheless, these methods demand additional gradients computations [24] or supervisions [57].

The literature on WSS in medical imaging remains scarce. While few methods resort to direct losses, hence requiring additional priors, such as the target size [58], [55], other approaches rely on stronger forms of supervision, for instance, using bounding boxes [54] or scribbles [59]. Tackling WSS from a perspective of *image-level* labels typically involves visual features, which has not been thoroughly investigated [60], [61], [62], [63]. For example, Nguyen *et al.* [62] has proposed a CAM-based approach for the segmentation of uveal melanoma. In their method, the CAMs generated by the classification network are further refined by an active shape model and conditional random fields [64]. More recently, CAMs derived from image-level labels have been combined with attention scores to refine lesion segmentation in brain images [65]. By doing so, they have demonstrated a performance improvement compared to the vanilla version of CAMs. Nevertheless, these methods typically integrate CAM/GradCAM with complex models to enhance the performance of a final segmentation.

III. METHODOLOGY

A. Overview

An overview of the proposed approach is depicted in Fig. 1. The main idea is to split the embedding space into multiple subspaces (K) such that the original embedding space can be learned by refining its subspaces. Contrary to [20], the embedding space is split dynamically, which removes the need to search for the optimal number of learners K in

each scenario. The whole process is divided into two iterative steps. First, input images are mapped into the lower dimension embedding space using the entire embedding layer E (d -dimension), where they are clustered into different groups. Second, the clustered data is consequently assigned to an individual subspace learner, where their corresponding images are used to train each subspace. These two steps are repeated at regular intervals, as well as each time a new learner is added. The key idea is that each subspace learner learns a part of the embedding space from a subgroup of images instead of learning a whole embedded representation vector. Finally, all subspaces are combined to generate a full embedding space. Furthermore, an attention module is integrated within the learning process to guide the learning of distance metrics. The following sections describe the deep metric learning formulation, present the proposed dynamic subspace metric learning and attention module.

B. Deep Metric learning Formulation

Let the training dataset be defined as $\mathbf{X} = \{(x_i, y_i)\}_{i=1}^N$, where the i -th image is denoted as $x_i \in \mathbb{R}^m$, and $y_i \in \{1, 2, \dots, C\}$ is its corresponding class label. C defines the total number of classes. The goal of deep metric learning is to learn an embedding function $f_\theta(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}^d$, which discriminatively maps semantically similar images (same class) in the input space \mathbb{R}^m onto metrically close points in the learned manifold \mathbb{R}^d . Similarly, semantically dissimilar images (different class) in \mathbb{R}^m should be mapped metrically far in \mathbb{R}^d . The parameters θ of the mapping function are typically learned by a convolutional neural network. Formally, the distance metric $d(x_i, x_j) : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ between two images in the embedding space \mathbb{R}^d can be defined as:

$$d(x_i, x_j) = \|f_\theta(x_i) - f_\theta(x_j)\|, \quad (1)$$

where $\|\cdot\|$ denotes the Euclidean norm. This distance can be minimized in different ways, depending on the loss function employed. In this work, we resort to the Margin loss [15]:

$$l_{margin}(x_i, x_j) = [\alpha + \mu_{ij}(d(x_i, x_j) - \beta)]_+, \quad (2)$$

where β is the boundary between the similar and dissimilar pairs, α is a separation margin, and $\mu_{ij} \in \{-1, 1\}$ indicates whether the images in the pair are similar ($\mu_{ij} = 1$) or different ($\mu_{ij} = -1$). Note that any other metric learning loss function can be employed with our approach.

C. Dynamic Subspace Learners

The complexity of the original problem can be solved by dividing the problem into smaller sub-problems, which are easier to solve. We follow the approach in [20], where the embedding space \mathbb{R}^d and the data is split into multiple groups. Specifically, splitting of the embedding space is conducted by slicing the \mathbb{R}^d space, i.e., the last dense layer of the network, into K sub-vectors of the same size, d/K . Furthermore, data is clustered into K groups based on their pairwise distance in the embedding space \mathbb{R}^d , for instance, using K-means. Then, a set of K independent learners is used to learn over

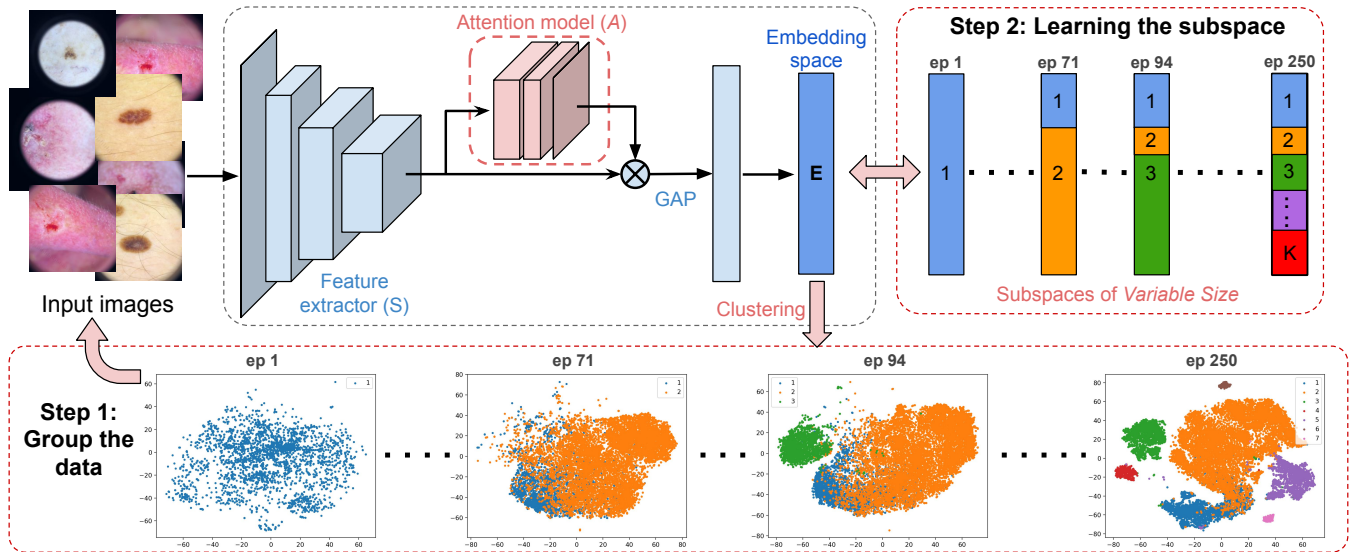


Fig. 1: Overview of our proposed attention-based dynamic subspace learners - The embedding space is dynamically divided into the subspaces of varying sizes during training. Suppose there are K subspaces at a particular training time; the data are first grouped into K groups in the full embedding space (step 1, from epoch 1 to 250) and assign each subgroup of data to an individual subspace learner. Each learner then only attends the data from its subgroup in the learning stage (step 2). In inference time, our method uses the entire embedding space E to map an image. Best viewed in color.

each subspace by using a fraction of the input data, thereby reducing the complexity of the original problem. Nevertheless, a major bottleneck is finding an optimal number of subspaces K to learn an effective embedding, which must be found empirically for every new dataset. Moreover, the subspace is divided equally, which is ineffective as not all the object attributes require the same size to encode the information.

Contrary to [20], our proposed learning strategy finds an optimal embedding by dynamically splitting the embedding space and associating with a metric learner during training. To construct each subspace, we group highly contributing neurons of the embedding layer E , which is repeated until network convergence. Initially, the entire embedding space is learned with all the data, with an initial single learner $K = 1$. As the learning progresses, the accuracy of the model starts to reach an initial plateau. At this stage, we compute the score of each neuron (e_i) in the embedding layer, similarly to the pruning strategy as in [66]. In particular, the low-scoring neurons are pruned such that the performance drop of the model is minimal, i.e., $|\Delta f_\theta(e_i)| = |f_\theta(X, e_i = 0) - f_\theta(X, e_i)|$. By using Taylor expansion, as in [66], the scoring of each neuron e_i can be reduced to:

$$s(e_i) = |\Delta f_\theta(e_i)| = \left| f_\theta(X, e_i) - \frac{\partial f_\theta}{\partial e_i} e_i - f_\theta(X, e_i) \right| = \left| \frac{\partial f_\theta}{\partial e_i} e_i \right| \quad (3)$$

Thus, the scoring of neurons is simplified to multiplying the activation and the gradient output in the embedding layer. This score $s(e_i)$ is computed for each training example separately, and is consequently averaged across all training data and normalized to $[0, 1]$. The neurons having high normalized scores are subsequently grouped to form a new subspace. Particularly, the neurons having more than 50% of the confidence

score, i.e., $s(e_i) > 0.5$, are grouped as a new subspace. The current metric learner (\mathcal{L}_k) is later assigned to this group of neurons. The remaining neurons of the embedding layer, e_r , are eventually reset, similar to the pruning technique [66] and assign a new metric learner as in Eq. 4. After adding this new learner, the training data is clustered by mapping into the entire embedding space using K-means with the updated K ($K = 2$ for the second iteration). Note that the entire embedding space here is a combination of all the subspaces. Each learner is eventually assigned a subgroup of data from the clustering, resulting in each learner being trained with a fraction of the input data. The addition of a new learner is repeated with the remaining neurons e_r when the network performance reaches a new plateau, until convergence. In the end, it results in K mapping functions, $\mathbf{f} = [f^1, f^2, \dots, f^K]$, where each mapping function f^k will project the images \mathbb{R}^m into the corresponding subspace of \mathbb{R}^{d_k} , each with a variable size.

All learners are trained jointly by resorting to the margin loss [15], which for each learner can be defined as:

$$\mathcal{L}_k^{f_{\theta_k}^k}(x_i, x_j) = \sum_{(x_i, x_j) \sim B} [\alpha + \mu_{ij}(d_{f_{\theta_k}^k}(x_i, x_j) - \beta)]_+, \quad (4)$$

where $(x_i, x_j) \sim B$ is the current mini-batch (uniformly sampled from each data group) having both positive and negative classes, and $d_{f_{\theta_k}^k}$ is the distance metric (similar to Eq. 1) for the k -th learner. Once individual learners are trained, these are merged to compose the entire embedding space, which is refined with the entire training set. Furthermore, assuming that the learned embedding space is improving over time, we re-cluster the images at every T_c epochs by mapping all the images using the entire embedding space E . An outline of the proposed method is presented in Algorithm 1.

Algorithm 1: Dynamic Subspace Learner Pseudocode

Inputs: X, X_{test} : Training and test data
 θ : backbone network parameters
 E : Embedding space
 T_c, T_p : clustering and network plateau threshold

Initialize: $K \leftarrow 1$, number of learner
 $B \leftarrow 0$, Best epoch
 $ep \leftarrow 1$, current epoch
 $e_r \leftarrow E$, remaining embedding space
 $RC \leftarrow \text{True}$, re-clustering flag

while *Not converged* **do**

if RC **then** ▷ Re-cluster the data

$E \leftarrow \text{ConcatEmbedding}(\{e_1, e_2, \dots, e_{K-1}, e_r\})$
 $\text{emb} \leftarrow \text{ComputeEmbedding}(X, \theta, E)$
 $\{C_1, C_2, \dots, C_K\} \leftarrow \text{ClusterData}(\text{emb}, K)$
 $\{e_1, e_2, \dots, e_{K-1}, e_r\} \leftarrow \text{SplitEmbedding}(E, K)$
 $RC \leftarrow \text{False}$

repeat ▷ Train all learners

$C_k \sim \{C_1, C_2, \dots, C_K\}$
 $b \leftarrow \text{GetBatch}(C_k)$
 $L_k \leftarrow \text{FPass}(b, \theta, f^k)$
 $\theta, f^k \leftarrow \text{BPass}(L_k, \theta, f^k)$

until *epoch completed*

$ep \leftarrow ep + 1$
 $E \leftarrow \text{ConcatEmbedding}(\{e_1, e_2, \dots, e_{K-1}, e_r\})$
 $RC \leftarrow (ep \bmod T_c == 0)$

if $\text{Evaluate}(X_{test}, \theta, E, ep) > B$ **then** ▷ Is best

$B \leftarrow ep$

else if $ep \geq (B + T_p)$ **then** ▷ Is network plateaued

$K \leftarrow K + 1$ ▷ Update new learner
 $\{e_{K-1}, e_r\} \leftarrow \text{splitLearner}(\{e_r\})$ ▷ using Eq.3
 $\{e_1, \dots, e_{K-1}, e_r\} \leftarrow \text{SplitEmbedding}(E, K, e_{K-1})$
 $\text{reset}(e_r)$
 $RC \leftarrow \text{True}$

$E \leftarrow \text{ConcatEmbedding}(\{e_1, e_2, \dots, e_{K-1}, e_r\})$
 $\theta, E \leftarrow \text{FineTune}(X, \theta, E)$

Output: θ, E

D. Attentive Dynamic Subspace Learners

Deep attention is raising as an efficient mechanism to focus the learning on the objects of interest in a wide range of applications, such as person re-identification [67], object classification [68], or medical image segmentation [69], [70]. Inspired by these advances, we introduce an attention module to learn attentive features, with the goal of enhancing the learning of the embedding space. For a given input image x_i , feature extractor $S(\cdot)$ produces a feature maps $s_i = S(x_i) \in \mathbb{R}^{c \times m \times n}$, where m, n denote the spatial dimension of the feature map and c the number of channels. The attention map produced by the attention module $A(\cdot)$ can be then defined as $a_i = A(s_i) \in \mathbb{R}^{m \times n}$. The generated attention map is multiplied with each feature map $a_i \odot s_i$, where \odot is the element-wise product, resulting in the set of attentive features. Last, the attentive features are combined to produce a c -dimensional vector by using global average pooling (GAP), which are mapped into the manifold space using a dense layer (Fig. 1).

E. Attention maps for Weakly Supervised Segmentation

The attention maps obtained by our proposed method can serve as proxy *pixel-level* labels to train a segmentation network in a fully-supervised manner. Specifically, the input image x_i and corresponding attention map a_i are used as a training pair. To differentiate foreground pixels from the background pixels in a_i , we threshold the attention maps with T_s (i.e., pixels in a_i greater than T_s are set to 1, 0 otherwise) before training the segmentation network. The network is trained with binary cross-entropy as a loss function, which is computed over pixel-wise softmax probabilities, defined as:

$$\mathcal{L}_{BCE}(x, a) = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^2 a_i^c \cdot \log(F_{\theta_s}(x_i^c)) \quad (5)$$

where F_{θ_s} is a segmentation network parameterized by θ_s . Note that the learning objective that trains a segmentation network is same in both the fully and weakly supervised scenario. However, the main difference lies in the labels employed in the cross-entropy term. In particular, while the former resorts to given segmentation masks, e.g., y , the latter leverages the obtained attention masks as pseudo-labels, i.e., a .

IV. EXPERIMENTS**A. Experimental Setting**

The performance of the proposed attention-based dynamic subspace learners (ADSL) is compared to other deep metric learning methods applied in medical imaging [45], [46], [47], [44], [43], which resort to contrastive or triplet loss. To assess the effectiveness of the dynamic learner training strategy, we compare it with the divide and conquer approach (DCML) [20]. Since we use class labels information, we compare with the classification network trained using a cross-entropy loss. For a fair evaluation, the backbone architecture and hyperparameters are fixed across the different methods. In addition, experiments across all the models and datasets are run three times, and their average performances are reported. Note that the baselines based on triplet and contrastive loss rely on single-learner, whereas models based on the *divide-and-conquer* strategy and our method employ multiple learners.

To assess the performance of our approach in terms of segmentation, we benchmark the resulting attention maps against the popular GradCAM [24] from the classification networks. We include a recent Attention Residual Learning (ARL) approach in [36] since it has been similarly proposed in the context of skin lesion analysis. We also include a recently proposed weakly supervised segmentation method, Embedded Discriminative Attention Mechanism (EDAM) [71], applied for the natural image. Lastly, we include as an upper bound the results obtained by UNet [72] that was trained on the provided pixel-level masks. Note that the model architecture and hyperparameters are fixed across the different methods. Nevertheless, the ARL model employs a carefully modified ResNet50 backbone with soft-attention blocks in each layer. It is noteworthy to mention that it also uses an offline multi-scale patch extraction strategy, resulting in extra images during




Fig. 2: Impact of number of learners K in DCML [20] - Each line indicates the NMI (top) and Recall@1 (bottom) scores across the three datasets. The default loss function employed is margin loss, whereas models with a triplet loss are explicitly mentioned. Best seen in color.

training. Whereas, the EDAM model employs a collaborative multi-head attention module after the feature extraction layer to directly generate the discriminative activation masks.

Datasets: The performance of the proposed method, in terms of clustering and image retrieval, is evaluated on three diverse medical imaging datasets: skin lesion from the ISIC 2019 Challenge [31], [32], musculoskeletal radiographs from the MURA dataset [33], and gastrointestinal tract images from the HyperKvasir dataset [34]. To assess the segmentation performance, we resort to the skin lesion dataset from the ISIC 2018 Challenge [35], [31].

a) ISIC19: This dataset consists of 25,331 images across 8 different categories. In our experiments, following the standard procedure in DML, we split our dataset into independent training and testing sets. Specifically, 20,000 images were used for training and the remaining 5,331 for testing.

b) MURA: It consists of 40,561 images from 9,045 normal and 5,818 abnormal musculoskeletal radiography studies across seven standard upper extremity types. We configure this as 14 categories (7 normal and 7 abnormal) to represent the data in a manifold. We use the provided split of 36,808 images for training and 3,197 images for testing.

c) HyperKvasir: This dataset consists of 110,079 images, of which 10,662 images are labeled across 23 different classes of findings. We randomly split the data into 8,567 images for training and the remaining 2,095 images for testing.

d) ISIC18: This dataset is composed of 2,594 images and their corresponding pixel-level masks. The segmentation dataset is randomly split into three sets: training (1,042), validation (520), and testing (1,038). We leverage the attention maps and GradCAMs generated on the ISIC19 dataset (25,331 images) as proxy-labels to train the segmentation networks. In contrast, the training set is used to train the upper-bound model, i.e., fully-supervised.

Evaluation Metrics: We follow the evaluation protocol typically employed in deep metric learning [20], [14]. In particular, we employ the normalized mutual information (NMI) to assess the clustering performance using K-means and the Recall score (with $k = 1$ and 4) to evaluate the image retrieval quality. To assess the segmentation performance, we employ the common

Dice score coefficient.

Implementation details: As in [20], we use ResNet50 [73] as the backbone architecture. The feature extractor layers consist of the first three residual blocks of ResNet50, used as input to the attention module. The attention module consists of three convolution layers with 3×3 kernel and filters size of $\{128, 32, 1\}$, with a ReLU activation between each convolutional layer. Last, a sigmoid activation is integrated into the final layer to produce the activation map. An input image size of 224×224 is used for all our experiments. All models are trained using the Adam optimizer [74] with batch size of $B = 32$. In each mini-batch, 8 images per class are sampled to ensure a class-balanced scenario and experiments are trained for 300 epochs. The last 50 epochs are fine-tuned with full embedding. The re-clustering parameter is set to $T_c = 2$ as in [20] and the network plateau threshold is empirically set to $T_p = 10$. The margin loss parameters are set to $\alpha = 0.2$, $\beta = 1.2$, as in [15]. Last, since most DML approaches [15], [20] employ an embedding space of size $d = 128$, we use the same latent dimension in all our experiments. The PyTorch implementation of our work is publicly available here: https://github.com/adigasu/Dynamic_subspace_learners.

Regarding the segmentation task, we use UNet [72] architecture with an initial kernel size of 32 with two convolution layers and a depth of 3. It is trained with Adam optimizer with batch sizes of 16. For each method, the threshold parameter T_s is set to maximize the Dice score on the initial maps of the validation set (Fig. 6).

B. Clustering and image retrieval results

Impact of number of learners K : One of the motivations of this work is to remove the need to empirically searching for the optimal number of learners. To validate this hypothesis, we first study the performance of DCML [20] by varying the number of subspace learners (K). Figure 2 depicts the results of this experiment across the three datasets and under two different loss functions: margin and triplet loss. In these plots, it can be observed that the optimal K value significantly differs across datasets and metrics. Thus, this limitation of the DCML approach results in extra time-consuming steps to fine-tune

the model in each dataset. In contrast, the proposed method (dotted line) eliminates the need of manually defining K by dynamically exploring the manifold, yet achieves on par results with the best performing DCML setting.

We also report the average K values obtained from our method over three runs, as well as the DCML (best) K in Table I. The table shows that the K value has no relation to the number of ground-truth classes. The dynamically obtained K in our method is driven by image content, not by the number of ground-truth classes, which explains their uncorrelated values.

Dataset	#classes	ADSL - Avg. K	DCML - Best K
ISIC19	8	7	6
MURA	14	4.67	1
HyperKvasir	23	4.33	2

TABLE I: Comparison of the obtained K values from our method and the DCML best K values with respect to the number of ground-truth classes.

Comparison to prior literature: We now compare our method with recent prior work as baselines, whose results are reported in Tables II-IV. As the performance of DCML varies with K , we report only the best and worst models. Note that the DCML with a single-learner, i.e., $K = 1$, is equivalent to a margin loss method [15]. We also report the performance of the embedding space learned by the classification network. From the Tables II-IV, we observe that the proposed method consistently achieves the best results in terms of NMI across the three datasets while performing on par with the best setting of the DCML approach on image retrieval metrics. As shown previously, it is important to note that the performance of DCML heavily depends on the value of K . For instance, the difference between the worst and best DCML configuration in NMI score can be up to 5% on the ISIC19 dataset. Compared to single-learner approaches, our method brings 5 and 2% improvements in NMI and Recall score on the ISIC19 dataset and up to a 1% improvement in both scores on the MURA and HyperKvasir datasets. This highlights the potential of exploring embeddings via multiple subspaces.

Furthermore, the comparison with the conventional classification network shows that our method consistently outperforms its accuracy up to 10% in terms of NMI scores on ISIC19 and up to 4% NMI score on MURA and HyperKvasir datasets and up to 4% and 1.5% in terms of Recall scores on the ISIC19 and MURA datasets. The averaged NMI and R@1 results of the proposed method slightly outperform the best DCML configuration, which is consistent across all the datasets. The standard deviation of our method is smaller in all cases for all metrics compared to the DCML. Overall, our method shows better robustness with respect to the state-of-the-art methods in the learning manifold space. The performance of our method is in line with the recent literature [75], [76].

Ablation study on the use of attention: Adding an attention module brings additional value to our model in terms of interpretability. Nevertheless, to assess whether this improvement is also reflected in the model performance, we compare our model to its non-attention counterpart, denoted as Dynamic




Fig. 3: Impact of the embedding size - Each bar indicates the NMI (top) and Recall@1 (bottom) scores on ISIC19 dataset. Compare to the best model of DCML method, our method produces better NMI and Recall scores for most cases.

Subspace Learners (DSL). Results from this study are reported in Table V, which shows that adding attention typically leads to a boost on the model performance. In particular, the attentive model brings 0.5 and 0.3% improvement as average over the three datasets for the NMI and R@4 metrics, respectively, while achieves on par results for R@1. Additionally, the attention module minimally increases the model memory by 5 MB (includes parameters, forward and backward pass size) when compared to non-attention counterpart, which is arguably negligible with respect to the overall model size (607 MB) in case of deployment.

Impact of the embedding size: We also evaluate the effect of representing the embedding space with different sizes. In particular, we assess the clustering and image retrieval performance on the ISIC19 dataset by fixing the embedding dimension size to 64, 128, 256, and 512. Figure 3 shows that increasing the embedding size results in a performance improvement, which is reflected in both NMI and recall metrics. Nevertheless, beyond a 256-dimension embedding, the performance of both models typically decreases.

Qualitative Analysis: To show the inter and intra-class representation power in the embedding space across different models, we visualize a t-SNE mapping [77] on the ISIC19 test set (Fig. 4). The classification network fails to discover clear boundaries across classes in the embedding space (Fig. 4b). This could be because of the cross-entropy loss when coupled with softmax, does not explicitly guarantee the minimization of intra-class variance or maximization of inter-class variance, which results in suboptimal discriminative features [78]. The single metric learner, i.e., DCML $K = 1$ (Fig. 4c), improves the class boundaries when compared to the classification network, yet they fail to possess compact clusters. On the other hand, inter-class discrimination is visually enhanced when resorting to multiple learners, i.e., DCML $K = 6$ (Fig. 4d) and our approach (Fig. 4e). Further, we can also observe that the proposed model yields more compact clusters than the DCML approach, which might be due to the freedom of our

Method	NMI (\uparrow)	R@1 (\uparrow)	R@4 (\uparrow)	Avg. of NMI + R@1 (\uparrow)
Classification network	45.41 \pm 1.95	77.85 \pm 0.86	90.54 \pm 0.51	61.63 \pm 1.40
Contrastive loss	31.47 \pm 0.39	78.13 \pm 0.59	91.13 \pm 0.08	54.80 \pm 0.49
Triplet loss	50.97 \pm 0.61	79.84 \pm 0.49	91.70 \pm 0.26	65.41 \pm 0.55
DCML (worst NMI, K = 1)	50.53 \pm 1.01	82.84 \pm 0.39	91.51 \pm 0.43	66.69 \pm 0.70
DCML (best NMI, K = 6)	55.08 \pm 0.83	82.29 \pm 0.56	91.73 \pm 0.36	68.69 \pm 0.70
ADSL (free from K, ours)	55.14 \pm 0.87	82.39 \pm 0.11	92.11 \pm 0.27	68.77 \pm 0.49

TABLE II: Quantitative evaluation on ISIC19 test set - The NMI, Recall, and average scores from the different methods. Our method is emphasized with light gray, whereas best and second-best results are highlighted with bold and underline.

Method	NMI (\uparrow)	R@1 (\uparrow)	R@4 (\uparrow)	Avg. of NMI + R@1 (\uparrow)
Classification network	71.09 \pm 1.25	74.21 \pm 0.27	92.59 \pm 0.40	72.65 \pm 0.76
Contrastive loss	74.28 \pm 0.53	71.65 \pm 0.53	92.07 \pm 0.36	72.97 \pm 0.53
Triplet loss	74.41 \pm 0.27	74.51 \pm 0.78	92.95 \pm 0.33	74.46 \pm 0.53
DCML (worst NMI, K = 10)	72.88 \pm 0.40	73.55 \pm 0.16	91.17 \pm 0.19	73.22 \pm 0.28
DCML (best NMI, K = 1)	74.67 \pm 0.35	75.36 \pm 0.79	92.89 \pm 0.18	75.02 \pm 0.57
ADSL (free from K, ours)	74.88 \pm 0.09	75.52 \pm 0.18	92.25 \pm 0.42	75.20 \pm 0.15

TABLE III: Quantitative evaluation on MURA test set - The NMI, Recall, and average scores from the different methods. Our method is emphasized with light gray, whereas best and second-best results are highlighted with bold and underline.

Method	NMI (\uparrow)	R@1 (\uparrow)	R@4 (\uparrow)	Avg. of NMI + R@1 (\uparrow)
Classification network	80.13 \pm 2.34	85.66 \pm 0.39	94.42 \pm 0.39	82.90 \pm 1.87
Contrastive loss	83.89 \pm 0.15	78.52 \pm 0.86	93.44 \pm 0.48	81.21 \pm 0.51
Triplet loss	82.24 \pm 0.19	83.44 \pm 0.34	93.92 \pm 0.22	82.84 \pm 0.27
DCML (worst NMI, K = 1)	83.31 \pm 0.19	84.79 \pm 0.59	94.05 \pm 0.26	84.05 \pm 0.39
DCML (best NMI, K = 2)	84.40 \pm 0.52	85.46 \pm 0.31	94.19 \pm 0.28	84.93 \pm 0.42
ADSL (free from K, ours)	84.18 \pm 0.12	85.82 \pm 0.27	94.24 \pm 0.41	85.00 \pm 0.20

TABLE IV: Quantitative evaluation on HyperKvasir test set - The NMI, Recall, and average scores from the different methods. Our method is emphasized with light gray, whereas best and second-best results are highlighted with bold and underline.




Fig. 4: Visualization of ISIC19 test set in embedding space using t-SNE - Each class is indicated by its individual color. When compared to a standard classification network, DCML $K = 1$ (a single-learner) improves the separation between classes. The multi-learner methods, DCML $K = 6$ and our method, further improve the separation between classes, while our method has the advantage of being free from the number of learners K . Best seen in color.

Dataset	Method	NMI (\uparrow)	R@1 (\uparrow)	R@4 (\uparrow)
ISIC19	DSL	54.11	82.74	91.95
	ADSL	55.14	82.39	92.11
MURA	DSL	74.21	75.85	92.26
	ADSL	74.88	75.52	92.25
HyperKvasir	DSL	84.44	85.36	93.54
	ADSL	84.18	85.82	94.24
Average	DSL	70.92	81.32	92.58
	ADSL	71.40	81.24	92.87

TABLE V: Impact of attention module - Per-dataset and average results of the proposed model with (ADSL) and without (DSL) the attention module. Best results are highlighted with bold for each dataset as well as for the average.

model to explore the manifold.

Qualitative evaluation in terms of image retrieval is assessed in Fig. 5, where a given random query with its five nearest neighbors, found using both DCML and our method, are shown. Additionally, we overlay the contour of our attention maps (having probability above 0.5) from the proposed method over their respective retrieved image. First, our method indeed retrieves images having similar lesions and colors from the ISIC19 dataset. In radiography wrist images, both DCML and our method have similar retrieval errors. Finally, retrieval images from the HyperKvasir dataset have similar image semantics in terms of texture and probe length using our method when compared to DCML. The coherence of image retrievals indicates that the intra- and inter-class similarities have been captured by our method and thereby demonstrates

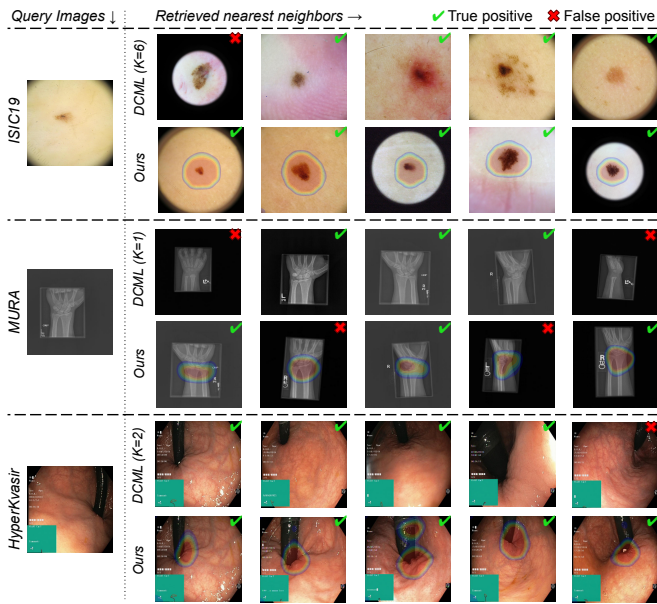


Fig. 5: Performance of image retrieval on test sets - Each query image and its five nearest neighbors in ascending order of distance are shown (left to right) from the DCML (best K) and our method with an overlay of our attention maps (probability above 0.5).

the robustness of our learned embedding. Moreover, our attention maps mainly concentrate on the lesion in the skin images, the wrist in the radiography images, and the probe contact region in the endoscopic images, demonstrating that our model decision are consistent over all retrievals.

C. Weakly Supervised Segmentation results

Table VI reports the results of the segmentation experiments. In this table, *Init maps* are used to denote the raw visual salient regions from either GradCAM or attention maps. *Refined* refers to the performance of the segmentation network trained on the *Init maps*. First, we can observe that segmentation results obtained by raw attention maps and GradCAMs are considerably low, with Dice values around 40%. This is likely due to the well-known fact that both are highly discriminative, resulting in over-segmented regions. The Attention Residual Learning (ARL) significantly outperforms these baselines, whose improvement could be due to the use of attentive residual blocks and additional multiscale data augmentation. The attention maps from the recent Embedded Discriminative Attention Mechanism (EDAM) method perform at a similar level when compared to ARL. Last, the attention maps from the proposed approach bring a significant boost compared to all the other methods. In particular, our model outperforms the baselines by nearly 30% and the recent ARL model by 13%. These results are typically consistent if we employ the initial maps as proxy-labels to train a segmentation network. In this case, raw attention maps or GradCAMs barely improve or even decrease the initial segmentation performance. In contrast, ARL, EDAM, and the proposed method reach higher Dice values, with about 1%, 3.5%, and 3% of increase,

Method	Init maps	Refined
Attention *	38.45	33.43
Attention †	38.52	38.38
GradCAM *	41.55	40.76
GradCAM †	39.80	41.27
ARL [36] ◊	56.78	57.60
EDAM [71] ◊	51.99	55.50
ADSL (ours) ◊	69.23	72.42
Full-supervision (upperbound)	-	86.15

TABLE VI: Performance of weakly supervised segmentation - “Initial maps” and “Refined” are Dice scores (in %) on the ISIC18 test set for different methods. Our method yields the best results (in bold). *, † and ◊ are from ResNet50, ResNet101 and modified ResNet50, respectively, indicating the used architecture in each visual map.

respectively. This represents a difference of 15% in Dice with respect to ARL. On the other hand, by only using image-level information, the proposed model bridges the gap with a fully-supervised network, with only 14% of difference. This suggests that the proposed model generates reliable segmentations.

Ablation study of threshold T_s on the raw visual maps: We evaluate the effect of threshold values T_s on the Dice score for raw visual maps from attention maps and GradCAMs, as shown in Fig 6. First, the attention maps and GradCAMs from the classification network have an almost flat Dice score of around 40% until $T_s = 0.4$, succeeded by a gradual decrease. The ARL and EDAM have a gradually increasing Dice score until $T_s = 0.4$ and $T_s = 0.6$ with a maximum score of 57.33% and 50.89%, respectively, followed by a gradual decrease. Our method outperforms the baselines for all threshold values in Dice scores with a maximum dice score of 69.0%, showing the robustness of the attention maps derived from our method. This study assists in setting a threshold value T_s for each method before training the segmentation network.

Qualitative Performance Evaluation: Visual results of the different methods are shown in Fig. 7. In this figure, *Init maps* (row 1 and 3) are raw visual salient regions from either GradCAM or attention maps shown as heatmaps, whereas *Refined* (row 2 and 4) refers to the performance of the segmentation network trained using *Init maps* as a proxy-labels. The attention maps (row 1 and 3) produced by the classification network spread all over the image, capturing some discriminative regions on the target lesion. GradCAMs spread around the target, highlighting discriminative regions of the lesion but failing to capture the whole context. The saliency map produced by the ARL method is focused on the target lesion. The attention maps obtained by the recent EDAM method spread around the target lesion, including the artifact regions, and fail to capture the target object context. In contrast, the attention maps derived from our approach better capture the attentive region, which mostly cover the lesion regions. The results show that our proposed approach generates superior attention maps compared to attention maps or GradCAMs from classification networks.

The results obtained by training a segmentation network on the initial salient regions (row 1 and 3) are depicted in




Fig. 6: Threshold T_s selection - Each line indicates the Dice scores of initial maps on the ISIC18 validation set for different methods. Our method outperforms the baselines for all T_s values. * and † are obtained by classification networks using ResNet50 and ResNet101, respectively.

row 3 and 4. These images demonstrate the feasibility of our method to weakly generate pixel-level labels that are usable for training segmentation networks.

V. DISCUSSION AND CONCLUSION

This paper presents a novel attention-based dynamic subspace metric learning approach for medical image analysis. The proposed algorithm leverages recent advances in deep metric learning using multiple metric learners. Our contribution improves the state-of-the-art method [20] with dynamic exploitation of subspace learners to learn the embedding space. Specifically, our novel training strategy overcomes the empirical search of the optimal number of subspace learners parameter while achieving competitive results in clustering and image retrieval tasks. Performance is extensively evaluated on three publicly available benchmark datasets: skin lesions, musculoskeletal radiography, and endoscopic images. Results demonstrate that our dynamic learner approach achieves the best results in clustering performance across all three datasets. Compared to the single-learner method, our method brings a maximum of 5 and 2% improvements in clustering and image retrieval scores on the ISIC19 dataset. Furthermore, our method significantly outperforms the classification network in all the datasets with a maximum of 10% and 4% improvements in clustering and retrieval scores on the ISIC19 dataset. Overall, the proposed method slightly outperforms in averaged results and has a smaller standard deviation when compared to the state-of-the-art methods in multiple metric learning. Our experiments have shown consistency across all the datasets, demonstrating the robustness of our method. Qualitative results show that the proposed method produces compact clustering and coherent image retrievals.

The addition of the attention module to our subspace learners provides the visual interpretability of the learned embedding space in terms of attention maps and improves the

clustering metrics. Our method offers new tools in multiple metric learners approaches, notably dynamically learning the number of learners and providing attention maps to hint at salient information caught by the learners. Studying the clinical usability of these tools remains to be explored. Nevertheless, A recent study [75] shows that the use of a retrieval network, in a single learner, yields an improvement of 9.2% in the decision accuracy of dermatologists. Our method indeed suggests that multiple learners capture a data embedding that yields a higher accuracy in clustering and retrieval tasks over single-learner methods, while additionally offering visual saliency from our attention mechanism.

The attention maps produced by our proposed method can serve as proxy pixel-level labels to train a segmentation network. The segmentation results outperform a state-of-the-art method, Attention Residual Learning (ARL) [36], as well as the recent Embedded Discriminative Attention Mechanism (EDAM) [71] by a margin of 15% and 17% in Dice scores, respectively, on the skin lesion dataset. The qualitative results demonstrate that the produced attention maps and their segmentation masks focus on the target lesion, demonstrating the effectiveness and robustness of our method. These attention maps produced in our subspace learning approach could therefore be potentially beneficial to a broader range of weakly supervised tasks, where the feature space remains challenging to represent using a single metric model within a specific task.

ACKNOWLEDGMENTS

This research work was partly funded by the Canada Research Chair on Shape Analysis in Medical Imaging, the Natural Sciences and Engineering Research Council of Canada (NSERC), and the Fonds de Recherche du Quebec (FQRNT). We would like to acknowledge Compute Canada for providing computing resources used for this work.

REFERENCES

- [1] K. Sohn, "Improved deep metric learning with multi-class N-pair loss objective," in *Advances in Neural Information Processing Systems*, 2016, pp. 1857–1865.
- [2] A. Y. Movshovitz, A. Toshev, T. K. Leung, S. Ioffe, and S. Singh, "No fuss distance metric learning using proxies," in *International Conference on Computer Vision*. IEEE, 2017, pp. 360–368.
- [3] X. He, Y. Zhou, Z. Zhou, S. Bai, and X. Bai, "Triplet-center loss for multi-view 3D object retrieval," in *Computer Vision and Pattern Recognition*. IEEE, 2018, pp. 1945–1954.
- [4] J. Ma, C. Zhou, P. Cui, H. Yang, and W. Zhu, "Learning disentangled representations for recommendation," in *Advances in Neural Information Processing Systems*, 2019.
- [5] N. Liu, Q. Tan, Y. Li, H. Yang, J. Zhou, and X. Hu, "Is a single vector enough? exploring node polysemy for network embedding," in *International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 932–940.
- [6] S. Zheng, Y. Song, T. Leung, and I. Goodfellow, "Improving the robustness of deep neural networks via stability training," in *Computer Vision and Pattern Recognition*. IEEE, 2016, pp. 4480–4488.
- [7] I. Ziko, E. Granger, and I. Ben Ayed, "Scalable laplacian K-modes," in *Advances in Neural Information Processing Systems*, 2018, pp. 10041–10051.
- [8] Z. Zhang and V. Saligrama, "Zero-shot learning via joint latent similarity embedding," in *Computer Vision and Pattern Recognition*. IEEE, 2016, pp. 6034–6042.
- [9] K. Q. Weinberger, J. Blitzer, and L. K. Saul, "Distance metric learning for large margin nearest neighbor classification," in *Advances in Neural Information Processing Systems*, 2006, pp. 1473–1480.




Fig. 7: Visual results of segmentation - (*Init maps*) Saliency map obtained by different methods and (*Refined*) their segmentation results. * and † are obtained by GradCAM on classification networks using ResNet50 and ResNet101, respectively.

- [10] B. Kulis et al., “Metric learning: A survey,” *Foundations and trends in machine learning*, vol. 5, no. 4, pp. 287–364, 2012.
- [11] M. Kaya and H. Ş. Bilge, “Deep metric learning: A survey,” *Symmetry*, vol. 11, no. 9, p. 1066, 2019.
- [12] R. Hadsell, S. Chopra, and Y. LeCun, “Dimensionality reduction by learning an invariant mapping,” in *Computer Vision and Pattern Recognition*, vol. 2. IEEE, 2006, pp. 1735–1742.
- [13] J. Wang et al., “Learning fine-grained image similarity with deep ranking,” in *Computer Vision and Pattern Recognition*. IEEE, 2014, pp. 1386–1393.
- [14] H. Oh Song, Y. Xiang, S. Jegelka, and S. Savarese, “Deep metric learning via lifted structured feature embedding,” in *Computer Vision and Pattern Recognition*. IEEE, 2016, pp. 4004–4012.
- [15] C.-Y. Wu, R. Manmatha, A. J. Smola, and P. Krahenbuhl, “Sampling matters in deep embedding learning,” in *International Conference on Computer Vision*. IEEE, 2017, pp. 2840–2848.
- [16] J. Wang, F. Zhou, S. Wen, X. Liu, and Y. Lin, “Deep metric learning with angular loss,” in *International Conference on Computer Vision*. IEEE, 2017, pp. 2593–2601.
- [17] X. Wang, X. Han, W. Huang, D. Dong, and M. R. Scott, “Multi-similarity loss with general pair weighting for deep metric learning,” in *Computer Vision and Pattern Recognition*. IEEE, 2019, pp. 5022–5030.
- [18] H. Lombaert, D. Zikic, A. Criminisi, and A. Nicholas, “Laplacian forests: semantic image segmentation by guided bagging,” in *Medical Image Computing and Computer Assisted Intervention*. Springer, 2014, pp. 496–504.
- [19] W. Kim, B. Goyal, K. Chawla, J. Lee, and K. Kwon, “Attention-based ensemble for deep metric learning,” in *European Conference on Computer Vision*, 2018, pp. 736–751.
- [20] A. Sanakoyeu, V. Tschernezki, U. Buchler, and B. Ommer, “Divide and conquer the embedding space for metric learning,” in *Computer Vision and Pattern Recognition*. IEEE, 2019, pp. 471–480.
- [21] B. Hu, B. Vasu, and A. Hoogs, “X-MIR: Explainable medical image retrieval,” in *Winter Conference on Applications of Computer Vision*, 2022, pp. 440–450.
- [22] M. D. Zeiler and R. Fergus, “Visualizing and understanding convolutional networks,” in *European Conference on Computer Vision*. Springer, 2014, pp. 818–833.
- [23] P. W. Koh and P. Liang, “Understanding black-box predictions via influence functions,” in *International Conference on Machine Learning*, vol. 70. JMLR.org, 2017, pp. 1885–1894.
- [24] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” in *International Conference on Computer Vision*. IEEE, 2017, pp. 618–626.
- [25] R. Chen, H. Chen, J. Ren, G. Huang, and Q. Zhang, “Explaining neural networks semantically and quantitatively,” in *International Conference on Computer Vision*. IEEE, 2019, pp. 9187–9196.
- [26] S. Belharbi, J. Rony, J. Dolz, I. B. Ayed, L. McCaffrey, and E. Granger, “Deep interpretable classification and weakly-supervised segmentation of histology images via max-min uncertainty,” *IEEE Transactions on Medical Imaging*, vol. 41, no. 3, pp. 702–714, 2021.
- [27] L. Chen, J. Chen, H. Hajimirsadeghi, and G. Mori, “Adapting Grad-CAM for embedding networks,” in *Winter Conference on Applications of Computer Vision*, 2020.
- [28] M. Zheng, S. Karanam, Z. Wu, and R. J. Radke, “Re-identification with consistent attentive siamese networks,” in *Computer Vision and Pattern Recognition*. IEEE, 2019, pp. 5735–5744.
- [29] A. Stylianou, R. Souvenir, and R. Pless, “Visualizing deep similarity networks,” in *Winter Conference on Applications of Computer Vision*. IEEE, 2019, pp. 2029–2037.
- [30] S. Zhu, T. Yang, and C. Chen, “Visual explanation for deep metric learning,” *IEEE Transactions on Image Processing*, vol. 30, pp. 7593–7607, 2021.
- [31] N. Codella et al., “Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (ISIC),” *arXiv preprint arXiv:1902.03368*, 2019.
- [32] M. Combalia et al., “BCN20000: Dermoscopic lesions in the wild,” *arXiv preprint arXiv:1908.02288*, 2019.
- [33] P. Rajpurkar et al., “MURA: Large dataset for abnormality detection in musculoskeletal radiographs,” *Medical Imaging with Deep Learning*, 2018.
- [34] H. Borgli et al., “HyperKvasir, a comprehensive multi-class image and video dataset for gastrointestinal endoscopy,” *Scientific Data*, vol. 7, no. 1, pp. 1–14, 2020.
- [35] P. Tschandl, C. Rosendahl, and H. Kittler, “The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions,” *Scientific data*, vol. 5, p. 180161, 2018.
- [36] J. Zhang, Y. Xie, Y. Xia, and C. Shen, “Attention residual learning for skin lesion classification,” *IEEE Transactions on Medical Imaging*, vol. 38, no. 9, pp. 2092–2103, 2019.
- [37] J. Bromley, I. Guyon, Y. LeCun, E. Säckinger, and R. Shah, “Signature

- verification using a "Siamese" time delay neural network," in *Advances in Neural Information Processing Systems*, 1994, pp. 737–744.
- [38] M. Opitz, G. Waltner, H. Possegger, and H. Bischof, "BIER-boosting independent embeddings robustly," in *International Conference on Computer Vision*. IEEE, 2017, pp. 5189–5198.
- [39] S. Liao, Y. Hu, X. Zhu, and S. Z. Li, "Person re-identification by local maximal occurrence representation and metric learning," in *Computer Vision and Pattern Recognition*. IEEE, 2015, pp. 2197–2206.
- [40] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Advances in Neural Information Processing Systems*, 2017, pp. 4077–4087.
- [41] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *Computer Vision and Pattern Recognition*. IEEE, 2015, pp. 815–823.
- [42] L. Yang *et al.*, "A boosting framework for visuality-preserving distance metric learning and its application to medical image retrieval," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 1, pp. 30–44, 2008.
- [43] K. Yan *et al.*, "Deep lesion graphs in the wild: relationship learning and organization of significant radiology image findings in a diverse large-scale lesion database," in *Computer Vision and Pattern Recognition*. IEEE, 2018, pp. 9261–9270.
- [44] K. Gupta, D. Thapar, A. Bhavsar, and A. K. Sao, "Deep metric learning for identification of mitotic patterns of HEP-2 cell images," in *Computer Vision and Pattern Recognition Workshops*. IEEE, 2019.
- [45] E. W. Teh and G. W. Taylor, "Metric learning for patch classification in digital pathology," 2019.
- [46] P. Pati, A. Foncubierta-Rodríguez, O. Goksel, and M. Gabrani, "Reducing annotation effort in digital pathology: A co-representation learning framework for classification tasks," *Medical Image Analysis*, p. 101859, 2020.
- [47] M. Sikaroudi, A. Safarpour, B. Ghogh, S. Shafiei, M. Crowley, and H. R. Tizhoosh, "Supervision and source domain impact on representation learning: A histopathology case study," in *Engineering in Medicine & Biology Society*. IEEE, 2020, pp. 1400–1403.
- [48] E. W. Teh and G. W. Taylor, "Learning with less data via weakly labeled patch classification in digital pathology," in *International Symposium on Biomedical Imaging*. IEEE, 2020, pp. 471–475.
- [49] A. Zhong *et al.*, "Deep metric learning-based image retrieval system for chest radiograph and its clinical applications in COVID-19," *Medical Image Analysis*, 2021.
- [50] P. Yang *et al.*, "Liver histopathological image retrieval based on deep metric learning," in *Bioinformatics and Biomedicine*. IEEE, 2019, pp. 914–919.
- [51] G. Papandreou, L.-C. Chen, K. Murphy, and A. L. Yuille, "Weakly-and semi-supervised learning of a DCNN for semantic image segmentation," in *International Conference on Computer Vision*. IEEE, 2015.
- [52] D. Lin, J. Dai, J. Jia, K. He, and J. Sun, "ScribbleSup: Scribble-supervised convolutional networks for semantic segmentation," in *Computer Vision and Pattern Recognition*. IEEE, 2016, pp. 3159–3167.
- [53] A. Bearman, O. Russakovsky, V. Ferrari, and L. Fei-Fei, "What's the point: Semantic segmentation with point supervision," in *European conference on computer vision*. Springer, 2016, pp. 549–565.
- [54] M. Rajchl *et al.*, "DeepCut: Object segmentation from bounding box annotations using convolutional neural networks," *IEEE Transactions on Medical Imaging*, vol. 36, no. 2, pp. 674–683, 2016.
- [55] H. Kervadec, J. Dolz, M. Tang, E. Granger, Y. Boykov, and I. B. Ayed, "Constrained-CNN losses for weakly supervised segmentation," *Medical Image Analysis*, vol. 54, pp. 88–99, 2019.
- [56] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Computer Vision and Pattern Recognition*. IEEE, 2016, pp. 2921–2929.
- [57] S. Kwak, S. Hong, B. Han *et al.*, "Weakly supervised semantic segmentation using superpixel pooling network," in *Association for the Advancement of Artificial Intelligence*, 2017, pp. 4111–4117.
- [58] Z. Jia, X. Huang, I. Eric, C. Chang, and Y. Xu, "Constrained deep weak supervision for histopathology image segmentation," *IEEE Transactions on Medical Imaging*, vol. 36, no. 11, pp. 2376–2388, 2017.
- [59] Y. B. Can, K. Chaitanya, B. Mustafa, L. M. Koch, E. Konukoglu, and C. F. Baumgartner, "Learning to segment medical images with scribble-supervision alone," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Springer, 2018, pp. 236–244.
- [60] X. Feng, J. Yang, A. F. Laine, and E. D. Angelini, "Discriminative localization in CNNs for weakly-supervised segmentation of pulmonary nodules," in *Medical Image Computing and Computer-Assisted Intervention*. Springer, 2017, pp. 568–576.
- [61] Q. Meng *et al.*, "Weakly supervised estimation of shadow confidence maps in fetal ultrasound imaging," *IEEE Transactions on Medical Imaging*, vol. 38, no. 12, pp. 2755–2767, 2019.
- [62] H.-G. Nguyen *et al.*, "A novel segmentation framework for uveal melanoma in magnetic resonance imaging based on class activation maps," in *Medical Imaging with Deep Learning*, 2019, pp. 370–379.
- [63] F. Dubost *et al.*, "Weakly supervised object detection with 2D and 3D regression neural networks," *Medical Image Analysis*, vol. 65, p. 101767, 2020.
- [64] P. Krähenbühl and V. Koltun, "Efficient inference in fully connected CRFs with gaussian edge potentials," in *Advances in Neural Information Processing Systems*, 2011, pp. 109–117.
- [65] K. Wu, B. Du, M. Luo, H. Wen, Y. Shen, and J. Feng, "Weakly supervised brain lesion segmentation via attentional representation learning," in *Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019, pp. 211–219.
- [66] P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz, "Pruning convolutional neural networks for resource efficient inference," *International Conference on Learning Representations*, 2017.
- [67] W. Li, X. Zhu, and S. Gong, "Harmonious attention network for person re-identification," in *Computer Vision and Pattern Recognition*. IEEE, 2018, pp. 2285–2294.
- [68] F. Wang *et al.*, "Residual attention network for image classification," in *Computer Vision and Pattern Recognition*. IEEE, 2017, pp. 3156–3164.
- [69] A. Sinha and J. Dolz, "Multi-scale self-guided attention for medical image segmentation," *IEEE Journal of Biomedical and Health Informatics*, 2020.
- [70] J. Schlemper *et al.*, "Attention gated networks: Learning to leverage salient regions in medical images," *Medical image analysis*, vol. 53, pp. 197–207, 2019.
- [71] T. Wu *et al.*, "Embedded discriminative attention mechanism for weakly supervised semantic segmentation," in *Computer Vision and Pattern Recognition*. IEEE/CVF, 2021, pp. 16 765–16 774.
- [72] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention*. Springer, 2015, pp. 234–241.
- [73] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Computer Vision and Pattern Recognition*. IEEE, 2016, pp. 770–778.
- [74] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations*, vol. 5, 2015.
- [75] C. Barata and C. Santiago, "Improving the explainability of skin cancer diagnosis using CBIR," in *Medical Image Computing and Computer-Assisted Intervention*. Springer, 2021, pp. 550–559.
- [76] S. Allegretti, F. Bolelli, F. Pollastri, S. Longhitano, G. Pellacani, and C. Grana, "Supporting skin lesion diagnosis with content-based image retrieval," in *International Conference on Pattern Recognition*. IEEE, 2021, pp. 8053–8060.
- [77] L. v. d. Maaten and G. Hinton, "Visualizing data using t-SNE," *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.
- [78] W. Liu, Y. Wen, Z. Yu, and M. Yang, "Large-margin softmax loss for convolutional neural networks," in *International Conference on Machine Learning*, vol. 2, no. 3, 2016, p. 7.