

SSD-KD: A Self-supervised Diverse Knowledge Distillation Method for Lightweight Skin Lesion Classification Using Dermoscopic Images

Yongwei Wang^{a,1}, Yuheng Wang^{b,c,d,e,f,*}, Tim K. Lee^{c,d,e,f}, Chunyan Miao^{g,*} and Z. Jane Wang^{b,c}

^aJoint NTU-WeBank Research Center on Fintech, Nanyang Technological University, Singapore

^bDepartment of Electrical and Computer Engineering, University of British Columbia, Vancouver, BC, Canada

^cSchool of Biomedical Engineering, University of British Columbia, Vancouver, BC, Canada

^dDepartment of Dermatology and Skin Science, University of British Columbia, Vancouver, BC, Canada

^ePhotomedicine Institute, Vancouver Coast Health Research Institute, Vancouver, BC, Canada

^fCancer Control Research Program, BC Cancer, Vancouver, BC, Canada

^gSchool of Computer Science and Engineering, Nanyang Technological University, Singapore

ARTICLE INFO

Keywords:

Skin cancer detection
Knowledge distillation
Deep learning
Dermoscopy


ABSTRACT

Skin cancer is one of the most common types of malignancy, affecting a large population and causing a heavy economic burden worldwide. Over the last few years, computer-aided diagnosis has been rapidly developed and make great progress in healthcare and medical practices due to the advances in artificial intelligence, particularly with the adoption of convolutional neural networks. However, most studies in skin cancer detection keep pursuing high prediction accuracies without considering the limitation of computing resources on portable devices. In this case, the knowledge distillation (KD) method has been proven as an efficient tool to help improve the adaptability of lightweight models under limited resources, meanwhile keeping a high-level representation capability. To bridge the gap, this study specifically proposes a novel method, termed SSD-KD, that unifies diverse knowledge into a generic KD framework for skin diseases classification. Our method models an intra-instance relational feature representation and integrates it with existing KD research. A dual relational knowledge distillation architecture is self-supervisedly trained while the weighted softened outputs are also exploited to enable the student model to capture richer knowledge from the teacher model. To demonstrate the effectiveness of our method, we conduct experiments on ISIC 2019, a large-scale open-accessed benchmark of skin diseases dermoscopic images. Experiments show that our distilled lightweight model can achieve an accuracy as high as 85% for the classification tasks of 8 different skin diseases with minimal parameters and computing requirements. Ablation studies confirm the effectiveness of our intra- and inter-instance relational knowledge integration strategy. Compared with state-of-the-art knowledge distillation techniques, the proposed method demonstrates improved performances. To the best of our knowledge, this is the first deep knowledge distillation application for multi-diseases classification on the large-scale dermoscopy database. Our codes and models are available at: <https://github.com/enkiwang/Portable-Skin-Lesion-Diagnosis>.

1. Introduction

Skin cancer is among the most common human malignancies globally, especially among the fair-skinned population. Over the past decades, skin cancer has pervaded different cultures and caused a huge economic burden in the healthcare maintenance. Meanwhile, both malignant melanoma (MM), which causes most deaths in skin cancers, and keratinocyte carcinomas (squamous cell carcinoma and basal cell carcinoma) have a promising prognosis if they are detected and well treated at the early stage (Apalla, Lallas, Sotiriou, Lazaridou and Ioannides, 2017). Therefore, early diagnosis plays an essential role in skin cancer's efficient management and successful treatment.

*Corresponding author.

 yuhengw@ece.ubc.ca (Yuheng Wang); ascymiao@ntu.edu.sg

(Chunyan Miao)

ORCID(s): 0000-0001-5252-7005 (Yuheng Wang); 0000-0002-0157-2081 (Tim K. Lee); 0000-0002-3791-0249 (Z. Jane Wang)

¹Yongwei Wang did part of the work at UBC, and he has moved to NTU.

²Yongwei Wang and Yuheng Wang contributed equally to this paper.

Benefiting from the unprecedented advances in optical imaging techniques, huge amounts of high-quality skin images were collected in multiple modalities. Dermoscopy, a noninvasive digital imaging technique, has been widely used for skin cancer diagnosis since it allows the in vivo evaluation of colors and structures of the epidermis that are hard to be visualized by the naked eye (Celebi, Codella and Halpern, 2019; Barata, Celebi and Marques, 2018). In traditional skin cancer diagnosis, expert dermatologists can inspect the skin lesions following valid algorithms such as ABCD rules (Nachbar, Stolz, Merkle, Cognetta, Vogt, Landthaler, Bilek, Braun-Falco and Plewig, 1994), 7-point Checklist (Argenziano, Fabbrocini, Carli, De Giorgi, Sammarco and Delfino, 1998) using dermoscopy images. However, such time consuming and expensive diagnostic methods cause a great burden on clinical diagnosis. Therefore, automatic skin cancer detection is in urgent need to relieve dermatologists' heavy work pressure in clinical diagnosis.

Convolutional neural networks (CNN) have been widely used in computer-aided diagnosis (CAD) for skin cancer detection and achieved remarkable performance, because of its excellent and robust capability in feature extraction

and categorical classification (Esteva, Kuprel, Novoa, Ko, Swetter, Blau and Thrun, 2017; Litjens, Kooi, Bejnordi, Sotio, Ciompi, Ghafoorian, Van Der Laak, Van Ginneken and Sánchez, 2017; Brinker, Hekler, Utikal, Grabe, Schadendorf, Klode, Berking, Steeb, Enk and Von Kalle, 2018). In recent years, more and more research efforts have been made to develop advanced deep CNN architectures to improve their diagnostic performance for skin cancer detection. Esteva et al. (Esteva et al., 2017) involved the pre-trained Inception V3 models in skin lesions' classification and presented dermatologist-leveled prediction results for the first time. Abbas et al. (Abbas and Celebi, 2019) proposed the DermoDeep framework for the classification of melanoma and nevus that consisted visual-features fusion and deep learning without pre- and post-processing steps. Kawahara et al. (Kawahara, BenTaieb and Hamarneh, 2016) found that the features extracted from the ImageNet pre-trained model outperformed general handcrafted features in skin lesion classification tasks. With the further development of deep learning technology, more novel structures have also been proposed and being used in this task. The integration strategies of various deep learning models and the interpretive analysis of the predictions they produce have also been extensively studied. For example, Gessert et al. (Gessert, Nielsen, Shaikh, Werner and Schlaefer, 2020) proposed an ensemble framework for skin lesion classification using multi-resolution Efficient Net architectures with a search strategy. Barata et al. (Barata, Celebi and Marques, 2021) came up with an explainable framework that can provide insightful information about deep learning models' decisions by incorporating the taxonomies information and special attention blocks. Meanwhile, the exploration of the combined effect of multiple modalities and the introduction of clinical diagnostic knowledge have further promoted the development of deep learning in the diagnosis of skin diseases. Kawahara et al. (Kawahara, Daneshvar, Argenziano and Hamarneh, 2018) proposed a multi-task deep convolutional neural network containing clinical, dermoscopic images and patient metadata. Their approach incorporated clinical rule-based knowledge into deep learning and demonstrated benchmark results for 7-point criteria and diagnostics. Based on this study, Bi et al. (Bi, Feng, Fulham and Kim, 2020) further considered the feature information that multi-modal input can share in the training phase of deep learning architecture and achieved further improvement in melanoma detection. Pacheco et al. (Pacheco and Krohling, 2021) proposed a MetaBlock module to improve classifiers' performance by utilizing the clinical metadata. Wang et al. (Wang, Cai, Louie, Wang and Lee, 2021) proposed a novel method that introduced the inter-dependencies between different criteria using a constrained classifier chain, which made the resulting predictions more acceptable to physicians. Tang et al. (Tang, Yan, Nan, Xiang, Krammer and Lasser, 2022) came up with a new algorithm that fused both feature-level and decision-level information from multiple modalities and achieved state-of-the-art performance in skin cancer diagnosis.

These remarkable achievements using varieties of CNN architectures have a great impact on the progress of automated skin cancer diagnosis. However, these research works related to CNNs also indicate that sophisticated models with more neural network layers and blocks tend to have better predictive performance with sufficient training data and proper training procedures. These works further suggest that, CNNs necessitate a huge number of parameters and massive floating point operations to obtain a satisfactory performance. For example, ResNet50 (He, Zhang, Ren and Sun, 2016), as one of the most well-known and effective deep learning models, has about 25.6 millions of parameters, which requires 98 MB and 4.11 GFlops to process the algorithm. Although these complex and powerful deep learning models can allow us to obtain better prediction results, they also require significant computing resources (e.g., memories and hardware resources). These strict requirements may impede their practicability in restrained environments and the embedded devices that have limited computation power and memories. Therefore, exploring portable and efficient networks with comparable performance for skin cancer auto-detection that can work on embedded equipment or limited resources has both practical and theoretical significance.

Knowledge Distillation (KD), proposed as an extended model compression method, can transfer the learned knowledge from a complicated model (the teacher) to a simple model (the student) by sharing the soft labels during the training phase, thus improving the representation power of the lightweight CNNs preferably (Hinton, Vinyals, Dean et al., 2015; Tung and Mori, 2019). In recent years, many novel and practical deep learning models based on KD ideas have been proposed. Romero et al. (Romero, Ballas, Kahou, Chassang, Gatta and Bengio, 2015) proposed the FitNet, an approach that added an additional requirement of predicting the outputs of the intermediate layers to the traditional KD algorithm and made it suitable for thin and deep networks. Yim et al. (Yim, Joo, Bae and Kim, 2017) measured a flow of feature correlations among adjacent layers in residual blocks and transferred them to a student network, then Liu et al. (Liu, Cao, Li, Yuan, Hu, Li and Duan, 2019) extended it to a dense cross-layer transfer scenario. Park et al. (Park, Kim, Lu and Cho, 2019) built a dubbed relational KD deep learning framework that considered both distance-wise and angle-wise distillation losses and improved the prediction ability of student models significantly. Peng et al. (Peng, Jin, Liu, Li, Wu, Liu, Zhou and Zhang, 2019) proposed a correlation congruence distillation scheme to model the relation between samples in the kernel form. Xu et al. (Xu, Liu, Li and Loy, 2020) further found that self-supervision signals can effectively transfer the hidden information from the teacher to the student via their proposed SSKD architecture, which substantially benefited the scenarios with few-shot and noisy-labels. Ding et al. (Ding, Wang, Xu, Wang and Welch, 2021c) proposed an data augmentation-based knowledge distillation framework for classification and regression tasks via synthetic samples (Ding, Wang, Xu,

Welch and Wang, 2021d; Ding, Wang, Wang and Welch, 2021b).

Meanwhile, KD has been proven effective in the compression of deep learning network structures on the basis of ensuring model prediction and learning capabilities among varieties of areas (Gou, Yu, Maybank and Tao, 2021; Wang and Yoon, 2021). In particular, KD shows great potential for applications in computer-aided diagnostics. For example, Qin et al. (Qin, Bu, Liu, Shen, Zhou, Gu, Wang, Wu and Dai, 2021) proposed an efficient architecture by distilling knowledge from well-trained medical image segmentation networks to train a lightweight network and obtained significant improvements for the segmentation capability while retaining the runtime efficiency. Abbasi et al. (Abbasi, Hajabdollahi, Khadivi, Karimi, Roshandel, Shirani and Samavi, 2021) implemented knowledge distillation on unlabeled medical data via an unsupervised learning manner and showed a significant improvement in the prediction of diabetic retinopathy. Chen et al. (Chen, Gao, Li and Shen, 2022) compared multiple deep learning models' performance regarding cervical cells classification using knowledge distillation and showed the importance of model selection. At the same time, knowledge distillation methods have also shown great potential in the field of skin disease diagnosis although there still exists quite limited attention in this area. Back et al. (Back, Lee, Shin, Yu, Yuk, Jong, Ryu and Lee, 2021) used an ensemble learning strategy to help a student network learn from multiple teacher networks progressively in the portable diagnosis of herpes zoster skin disease using clinical images, and achieved robust performances. To our best knowledge, however, there are still no studies on the classification of multiple skin diseases based on large-scale dermoscopy image datasets.

Despite that these innovative studies have advanced the application of KD in the field of computer-aided diagnosis, they still mainly rely on the sharing and transferring the subject-level category information between the teacher network and the student network. In detail, existing skin diagnosis KDs have largely ignored the inter and intra relations within input samples, which may preserve important knowledge to diagnose skin diseases from the teacher network. This work explores the intra-instance relational knowledge, as well as its integrated analysis with traditional inter-instance relational knowledge and dark knowledge, to enhance the effectiveness of information transfer from the teacher network to the student network. Consequently, the proposed diverse knowledge-aware KD can effectively improve the predictive capabilities of the student model.

In this work, we develop a novel knowledge distillation framework, termed SSD-KD, that integrates diverse conventional knowledge with a novel type of intra-instance relational knowledge to improve the auto-diagnosis of multiple types of skin diseases based on lightweight deep learning models. Specifically, we formulate a dual relational knowledge distillation architecture introducing different relational representations and softened network outputs to distill and transfer diverse knowledge from the teacher model to student

model. Moreover, we also incorporate the self-supervised knowledge distillation strategy into the framework to enable the student model capture richer structured knowledge from the self-supervision predictions of the teacher model. Additionally, to make the proposed SSD-KD framework more suitable for the skin lesions auto-detection task, which always suffers from data imbalance problems, we replace the general cross-entropy loss function in conventional KDs with an improved weighted version to benefit the skin disease categories with less subjects. We have evaluated our proposed method in the publicly available skin lesion datasets ISIC 2019, which contains over 20000 dermoscopy images and 8 types of skin diseases. The comparison with state-of-the-art methods and ablation studies are also performed to demonstrate the effectiveness of our proposed framework.

The main contributions of this study can be summarized as follows:

- We propose a novel self-supervised diverse knowledge distillation method for lightweight multi-class skin diseases classification using dermoscopy images, named SSD-KD. We replace the traditional single relational modeling block with dual relational blocks in our method, which assists the student model capture both intra- and inter-instance relational information from the teacher model. Hence, our diverse knowledge-integrated KD can further improve the transferable ability of KD and the performance for skin disease classification.
- We employ the self-supervision based auxiliary learning strategy in our knowledge distillation framework, which enables richer structured knowledge to be transferred, and stabilizes the training dynamics of the intra- and inter-instance relational distillation scheme.
- We utilize the weighted cross entropy loss function and combine it with our SSD-KD structure to assist the lightweight student model better mimic the behavior of the teacher model in capturing additional information from the diseases with limited subjects.
- The results demonstrate that our carefully designed method can help the student model achieve comparable or even better performance than the teacher model in the multi-class skin lesion classification task. From a practical clinical point of view, our proposed method can further advance the application of remarkable deep learning models in scenarios with limited computing resources.

The rest of the paper is organized as follows: the detailed introduction of the proposed method is presented in the second section, followed by the experiment setup and results analysis in the third section. The discussion and the future works are shown in the fourth section while the conclusion of the study is put in the last section.

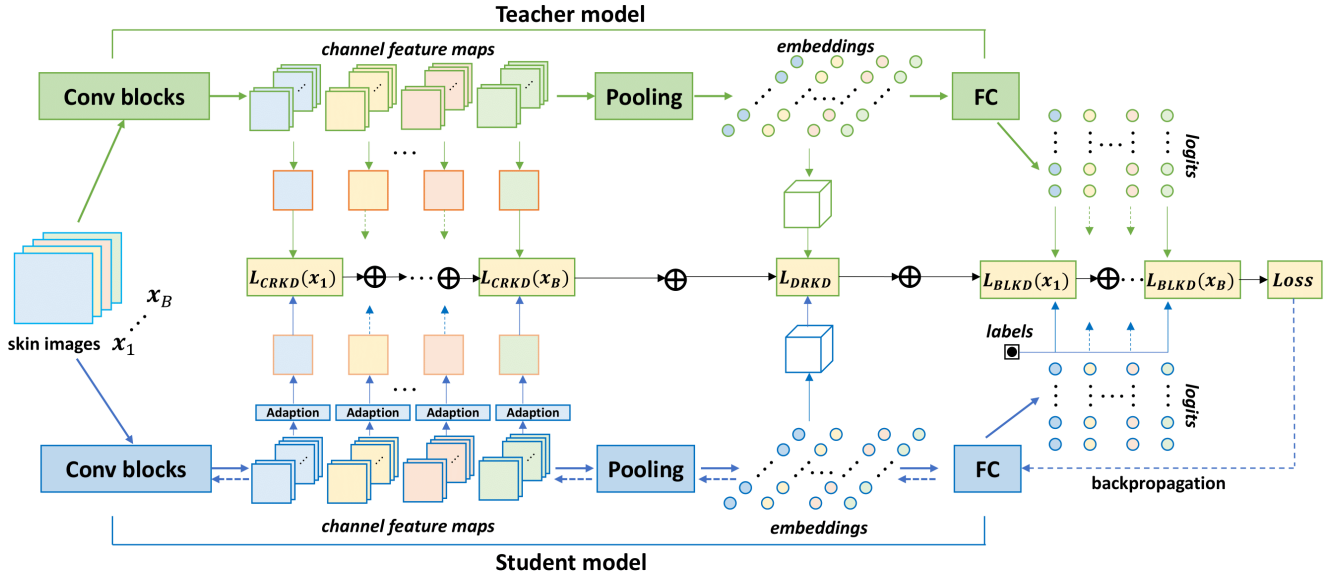


Figure 1: An illustration of the proposed knowledge distillation framework that distills and transfers diverse knowledge to the student model for skin cancer diagnosis. The upper branch (in green) denotes the pretrained teacher classification model, and the lower branch (in blue) denotes the untrained student classification model. Both models consist of a set of convolutional blocks to extract feature maps of a batch of skin images, followed by average pooling operations to produce embeddings, and a fully connected (FC) layer to yield logits for skin disease classification. The three types of knowledge and corresponding loss functions are quantified based on channel feature maps, embeddings and logits, respectively. In particular, our novel channel relational distillation loss is defined in the feature map stream. Finally, the diverse knowledge-integrated distillation loss is computed and backpropagated (along dashed reverse arrow lines) to train the student network. The self-supervised learning strategy can be conveniently incorporated to further boost the training of the student model.

2. Methodology

In this section, we describe the proposed method that aims to develop an effective knowledge distillation framework for skin disease detection. The overall pipeline of the proposed method is shown in Figure 1. In the proposed framework, we distill and transfer three types of knowledge from a large and pretrained teacher model to a portable student model that will be trained from scratch. The first type of knowledge adopts a customized version of the conventional knowledge distillation scheme (Hinton et al., 2015) that defines the “dark knowledge” via logits (i.e., the output vector from the last layer of a CNN model). The second type of knowledge utilizes the inter-instance relational knowledge (Park et al., 2019) to model the relationship of a batch of training samples at the penultimate layer. In addition, we define the intra-instance relational knowledge between channel feature maps of each sample at the last convolutional layer. This type of knowledge further exploits the inherent structures within model responses to facilitate the training of the student model.

For convenience, we introduce some notations to be used in the following section. Let us denote the overall training set of skin images as \mathcal{X} with groundtruth labels as \mathcal{Y} . Following the general routine in deep learning, we split the training set $(\mathcal{X}, \mathcal{Y})$ into non-overlapping subsets randomly. The number of skin images (i.e., batch size) in each subset is denoted as B , i.e., $\{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_B, \mathbf{y}_B)\} \subset \mathcal{X}$. For an input sample \mathbf{x}_i ($i = 1, \dots, B$), we denote its

logits vector by $\mathbf{l}(\mathbf{x}_i)$, the embedding vector by \mathbf{e}_i (i.e., $\mathbf{e}(\mathbf{x}_i)$ for brevity), and channel feature maps by $\mathbf{f}(\mathbf{x}_i)$, respectively. Here $\mathbf{l}(\mathbf{x}_i) = [l_1(\mathbf{x}_i), \dots, l_C(\mathbf{x}_i)]$, where C represents the number of unique classes of skin images; $\mathbf{f}(\mathbf{x}_i) = [f_1(\mathbf{x}_i), \dots, f_K(\mathbf{x}_i)]$ with $f_k(\mathbf{x}_i) \in \mathbb{R}^{H_f \times W_f}$ ($k = 1, \dots, K$), where H_f, W_f denote the height and width of a feature map, respectively.

2.1. Logit-based dark knowledge

Hinton et al. (2015) proposed a logit-based knowledge distillation method (a.k.a BLKD). Essentially, BLKD utilizes the (softened) logits of a teacher model as informative knowledge, which is then transferred to assist the training of a student model.

BLKD includes two loss terms: 1) a term that matches the softened logits of a teacher model with those of a student one, and 2) a regular classification term that matches the predictions of a student model with the groundtruth label. For an image \mathbf{x}_i , denote its softened logits from a model as $\mathbf{p}(\mathbf{x}_i, T)$ at a temperature T , where $\mathbf{p}(\mathbf{x}_i, T) = [p_1(\mathbf{x}_i, T), \dots, p_C(\mathbf{x}_i, T)]$. Here $p_c(\mathbf{x}_i, T)$ denotes the soft probability that \mathbf{x}_i belongs to a class c . With the softmax function, $p_c(\mathbf{x}_i, T)$ can be computed via the logit $l_c(\mathbf{x}_i)$,

$$p_c(\mathbf{x}_i, T) = \frac{\exp(l_c(\mathbf{x}_i)/T)}{\sum_{j=1}^C \exp(l_j(\mathbf{x}_i)/T)} \quad (1)$$

where a higher temperature T provides softer probabilities for class predictions, and tends to indicate richer informative

structures between classes. In BLKD, the loss of the first term measures the distribution differences of soft probabilities from a teacher (i.e., $p^t(\mathbf{x}_i, T)$) and a student (i.e., $p^s(\mathbf{x}_i, T)$) using the Kullback–Leibler (KL) divergence,

$$\begin{aligned} L_{KD} &= \sum_{i=1}^B D_{KL}(p^t(\mathbf{x}_i, T) || p^s(\mathbf{x}_i, T)) \\ &= \sum_{i=1}^B \sum_{c=1}^C p_c^t(\mathbf{x}_i, T) \log \frac{p_c^t(\mathbf{x}_i, T)}{p_c^s(\mathbf{x}_i, T)} \end{aligned} \quad (2)$$

where L_{KD} denotes the loss of the first term in BLKD that are computed on a batch of images.

The predictions of a model can be obtained by simply setting $T = 1$ in Eq. (1). Then, as the second loss term in BLKD, a cross entropy loss can be computed between model predictions and groundtruth labels. Unlike general computer vision tasks, however, the second loss term is not well suited for our skin cancer diagnosis problems. This is mainly due to the data imbalance issue in skin disease application, namely, we may only have few skin images for some diseases. Therefore, in this case, we turn to the weighted cross entropy (WCE). Denote by w_c the weight for a class c , where w_c can be computed as the inverse of the frequency for a disease class c over the training dataset. Denote also the one-hot encoded label of \mathbf{x}_i as $\mathbf{y}(\mathbf{x}_i) = [y_c(\mathbf{x}_i), \dots, y_C(\mathbf{x}_i)]$. Then, the WCE loss is expressed as,

$$L_{WCE} = \sum_{i=1}^B \sum_{c=1}^C -w_c \cdot y_c(\mathbf{x}_i) \log p_c^s(\mathbf{x}_i, T=1) \quad (3)$$

Then, the customized overall loss term of BLKD equals,

$$L_{BLKD} = (1 - \lambda_{kd}) \cdot L_{WCE} + \lambda_{kd} \cdot L_{KD} \quad (4)$$

where λ_{kd} is a hyperparameter to balance the two loss terms in BLKD.

2.2. Inter-instance relational knowledge

An inter-instance relational knowledge was proposed to model the structural relations among a batch of training data (Park et al., 2019), termed data relational knowledge distillation (DRKD). Though there exist different variants of inter-instance knowledge (e.g., a kernelized form in (Peng et al., 2019)), this study selects to use the popular DRLK due to its effectiveness and computational efficiency. To deal with the data scarcity challenge in skin cancer diagnosis, we also utilize such knowledge to additionally supervise the training of the student model.

DRKD maps the embeddings of different instances (within a batch) into relational representations, then it minimizes the structural differences of such representations between the teacher and student models, such that the student can better mimic the teacher. Given instances $\{\mathbf{x}_1, \dots, \mathbf{x}_B\}$ and their embeddings denoted by $\{e_1, \dots, e_B\}$, Park et al. (2019) defined two types of inter-instance relational knowledge: the distance-wise relation $\psi_d(e_i, e_j)$ between any

two instances, and the angle-wise relation $\psi_a(e_i, e_j, e_k)$ between any three instances. Here the potential functions $\psi_d(e_i, e_j)$ and $\psi_a(e_i, e_j, e_k)$ are computed based on network embeddings e_i, e_j, e_k via the normalized Euclidean distance and angular distance, respectively. The loss term of DRKD is defined as,

$$\begin{aligned} L_{DRKD} &= \lambda_d \cdot \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{X}_B^2} l(\psi_d^t(e_i, e_j), \psi_d^s(e_i, e_j)) \\ &\quad + \lambda_a \cdot \sum_{(\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k) \in \mathcal{X}_B^3} l(\psi_a^t(e_i, e_j, e_k), \psi_a^s(e_i, e_j, e_k)) \end{aligned} \quad (5)$$

where \mathcal{X}_B represents a training subset comprised of a batch of instances, l denotes the Huber function (Huber, 1992); λ_d, λ_a denote the weight for the distance-wise loss and the angle-wise loss (Park et al., 2019), respectively.

2.3. Intra-instance relational knowledge

To further improve the performance of a student model, we propose to exploit the intra-instance relational knowledge using channel feature maps for each instance. We term this novel type of knowledge distillation method as channel relational knowledge distillation (CRKD). This work only extracts channel relations of feature maps at the last convolutional layer, since these feature maps contain more fine-grained and critical features for making final predictions. It is worth mentioning that the instance-level congruence term in (Peng et al., 2019) actually refers to the KL divergence of logits between a teacher and a student (i.e., as in Eq. 2), which is clearly different from our formulation.

As illustrated in Figure 1, given an input image \mathbf{x}_i , a (teacher or student) model extracts high-level features $f_k(\mathbf{x}_i)$ for $k = 1, \dots, K$. Then, CRKD can be performed following two sequential steps: the channel adaptation step, and the intra-instance relation generation step. Since the number of channels of a student model may not be equal to that of a teacher model, it is infeasible to align the intra-instance relations from two models directly. Therefore, for the student model, we adopt an adaptation module to match its channel number with that of the teacher. Specifically, we utilize a 1×1 convolution (Lin, Chen and Yan, 2014) that linearly projects the channel feature maps of the student model to a space that shares a same channel number with the teacher model (Romero et al., 2015).

In the second step, we will model the relation between an arbitrary (vectorized) feature map pair of an instance \mathbf{x}_i as $(\text{Vec}(f_k(\mathbf{x}_i)), \text{Vec}(f_{k'}(\mathbf{x}_i)))$ utilizing the un-normalized cosine similarity metric,

$$r(f_k(\mathbf{x}_i), f_{k'}(\mathbf{x}_i)) = \langle \text{Vec}(f_k(\mathbf{x}_i)), \text{Vec}(f_{k'}(\mathbf{x}_i)) \rangle \quad (6)$$

Denote the overall intra-instance relational matrix as $\mathbf{R} \in \mathbb{R}^{K \times K}$, then \mathbf{R} can be generated with entries $\mathbf{R}_{k,k'}$ as,

$$\mathbf{R}_{k,k'}(\mathbf{x}_i) = r(f_k(\mathbf{x}_i), f_{k'}(\mathbf{x}_i)), k, k' \in \{1, \dots, K\} \quad (7)$$

A more computationally efficient way is adopt the Grammian matrix (Johnson, Alahi and Fei-Fei, 2016; Ding, Wang,

Yuan, Jiang, Wang, Huang and Wang, 2021a) to generate this relational matrix. To encourage the student learn the intra-instance relations of the teacher, we define the CRKD loss term over a batch of samples,

$$L_{CRKD} = \sum_{i=1}^B \frac{1}{KH_fW_f} \left\| \mathbf{R}^t(\mathbf{x}_i) - \mathbf{R}^s(\mathbf{x}_i) \right\|_F \quad (8)$$

where F denotes the Frobenius norm of matrices; H_f, W_f refer to the height and width of a feature map of the teacher model, respectively.

The intra-instance relation formulation above appears to resemble the angle-wise relation in DRKD, but here we are considering a two-tuple relation scenario. Another major distinction is that we are modeling the relation between any two channel feature maps of one instance rather than the embeddings from different instances. Our CRKD also differs from (Yim et al., 2017) and (Liu et al., 2019) that aim to measure the correlations across different network layers which tend to involve considerably more computational costs and instability issues. Besides, without the channel adaptation module as in our method, they require a similar network architecture (e.g., from a same architecture family) for the teacher and student models. This strict requirement, however, restrains their applicability in our considered scenario to establish more portable skin lesion diagnosis models.

A physical interpretation of our intra-instance relational modeling is to extract the texture information of an input image \mathbf{x}_i in the feature space. The extracted textures have been demonstrated crucial in computer vision tasks, such as image recognition (Geirhos, Rubisch, Michaelis, Bethge, Wichmann and Brendel, 2018) and visual tracking attacks (Ding et al., 2021a). Since textural features also contain important discriminative knowledge in the skin lesion diagnosis task, this type of relation and knowledge can provide auxiliary favorable cues to facilitate the training of our student model.

2.4. Diverse knowledge integration

To further boost the student training process, we integrate diverse levels of knowledge and transfer them to the student model. As depicted in Figure 1, our diverse knowledge distillation (D-KD) loss can be expressed as,

$$L_{D-KD} = \lambda_{blkd} \cdot L_{BLKD} + \lambda_{drkd} \cdot L_{DRKD} + \lambda_{crkd} \cdot L_{CRKD} \quad (9)$$

where $\lambda_{blkd}, \lambda_{drkd}, \lambda_{crkd}$ denote the hyperparameter to balance the loss terms for BLKD, DRKD, and CRKD, respectively.

In addition, we can incorporate the self-supervised training strategy (Xu et al., 2020) into the diverse knowledge-integrated distillation loss L_{D-KD} to additionally improve the training dynamics of the student model. Self-supervised learning (SSL) is an effective strategy to learn generic visual representations from unlabeled data for downstream tasks (e.g., image classification, semantic segmentation) in the

Algorithm 1: The algorithm of our method.

Data: A training dataset $(\mathcal{X}, \mathcal{Y})$, a pretrained teacher model, an untrained student model, batch size B , the maximum number of training epochs N , hyperparameters $\lambda_{blkd}, \lambda_{drkd}, \lambda_{crkd}, \lambda_{sskd}$, a minibatch stochastic gradient descent optimizer, an early stop epoch K_{stop} .

```

1 Initialize the student model using random initialization, set the
  early stopping flag as False;
2 Split  $(\mathcal{X}, \mathcal{Y})$  into  $\lceil N/B \rceil$  non-overlapping batches randomly;
3 for  $n = 1$  to  $N$  do
4   Reset the knowledge distillation loss  $L$  as 0;
5   for  $iter = 1$  to  $\lceil N/B \rceil$  do
6     Draw a batch of skin images  $\{\mathbf{x}_1, \dots, \mathbf{x}_B\}$ ;
7     Input the image batch into the teacher model, computes
      intra, inter-instance relational knowledge, and
      logit-based knowledge;
8     Input the image batch into the student model, computes
      intra, inter-instance relational knowledge, and
      logit-based knowledge;
9     Compute the knowledge distillation loss using Eq. (10)
      for SSD-KD;
10    Update the student model using the optimizer;
11  end
12  If loss not decrease within  $K_{stop}$  epochs: stop the program;
13  Else: continue the program;
14 end
15 Return: An optimized portable student model.
```

computer vision field (Chen, Kornblith, Norouzi and Hinton, 2020; He, Fan, Wu, Xie and Girshick, 2020; Jing and Tian, 2020). Moreover, a very recent work (Liu, HaoChen, Gaidon and Ma, 2022) shows that this learning strategy can well alleviate the data imbalance issue, which further validated our motivation to incorporate the self-supervision knowledge.

A typical SSL scheme is to employ a contrastive loss (Van den Oord, Li and Vinyals, 2018) that maximizes the agreement in feature space of positive sample pairs generated from different augmentation transformations on a same image (He et al., 2020; Jing and Tian, 2020). For the first time, Xu et al. (Xu et al., 2020) introduced SSL into knowledge distillation (SSKD), and conducted a thorough study on the effectiveness of this learning strategy in the KD setting. To perform the self-supervision, SSKD pre-trains a teacher model with the contrastive loss, then transfers the self-supervised predictions from the teacher to the student. Following (Xu et al., 2020), we integrate the self-supervised training strategy into our D-KD framework. Then, the loss function for the self-supervised diverse knowledge distillation (SSD-KD) is,

$$L_{SSD-KD} = L_{D-KD} + \lambda_{sskd} \cdot L_{SSKD} \quad (10)$$

where λ_{sskd} denotes the hyperparameter for the self-supervision regularizer.

In Algorithm 1, we describe the details of the proposed method.

3. Experiments and results

In this section, we conduct experiments to empirically demonstrate the effectiveness of the proposed method. We will describe in detail our experimental setup, evaluation

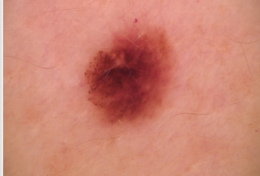
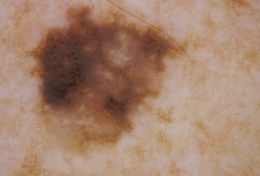
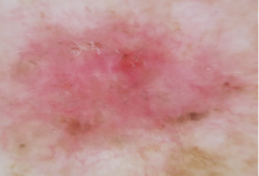
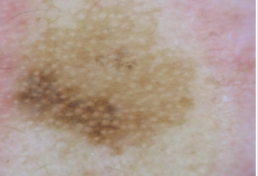
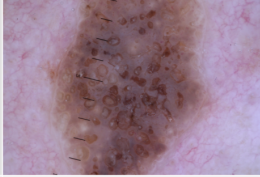
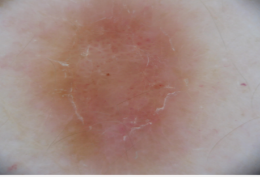
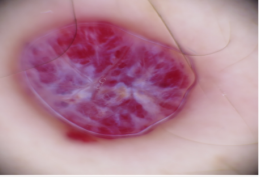
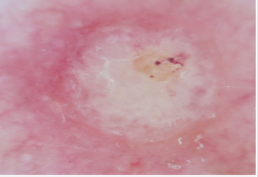
Types	Malignant Melanoma (MM)	Melanocytic Nevus (MN)	Basal Cell Carcinoma (BCC)	Actinic Keratosis (AK)
Example Images				
Description	A type of skin cancer that develops from the pigment-producing cells known as melanocytes.	Benign nests of melanocytes that typically appear as small brown spots.	A type of common skin cancer that often appears as a translucent bump on the skin.	A rough, scaly patch that often appears in the sun exposure body sites.
Number of Images	4522	12875	3323	867
Types	Benign Keratosis (BKL)	Dermatofibroma (DF)	Vascular Lesion (VASC)	Squamous Cell Carcinoma (SCC)
Example Images				
Description	A common benign skin growth that often appears in older population. Generally are brown, black or light tan.	Benign fibrous growth on the skin that could appear in various colors.	Represent a number of skin abnormalities that usually caused by vascular malformations.	The second most common form of skin cancer that develops in the squamous cells.
Number of Images	2624	239	253	628

Figure 2: Example images along with their detailed information from the ISIC 2019 dataset.

metrics, and then present the comparison results and our analysis.

3.1. Dataset

The dataset we used in this study is ISIC 2019 (Tschandl, Rosendahl and Kittler, 2018; Codella, Gutman, Celebi, Helba, Marchetti, Dusza, Kalloo, Liopyris, Mishra, Kittler et al., 2018; Combalia, Codella, Rotemberg, Helba, Vilaplana, Reiter, Carrera, Barreiro, Halpern, Puig et al., 2019), the large-scale publicly available dermoscopy imaging database, that contains 25331 dermoscopic images acquired from different resources along with their diagnosis labels as their ground truth. Part of the images also have corresponding meta information, such as gender, age, anatomy site, etc. The ISIC 2019 contains 8 different skin diseases including malignant melanoma (MM), melanocytic nevus (MN), basal cell carcinoma (BCC), actinic keratosis (AK), benign keratosis (BKL), dermatofibroma (DF), vascular lesion (VASC) and squamous cell carcinoma (SCC). An overview of details for each diseases along with the example images is presented in Figure 2.

3.2. Implementation details

In the experiment part, we selected ResNet50 as the teacher model while MobileNetV2 was chosen as the student model. ResNet50 is one of the most powerful convolutional neural network (CNN) architectures and has demonstrated outstanding representation and generalization abilities in medical imaging analytic area (Raghu, Zhang, Kleinberg and Bengio, 2019; Zhang, Xie, Wu and Xia, 2019; Xie, Niu, Liu, Chen, Tang and Yu, 2021). However, ResNet50 increases the complexity of architecture and requires huge computational power during training. Therefore, ResNet50 is a suitable choice for teacher model in our project. On the other hand, MobileNetV2, as an advanced version of MobileNet, benefits from its reduced network size and number of parameters, and it has been proven as an effective feature extractor for object detection for mobile applications (Sandler, Howard, Zhu, Zhmoginov and Chen, 2018; Toğaçar, Cömert and Ergen, 2021; Srinivasu, SivaSai, Ijaz, Bhoi, Kim and Kang, 2021). In this project, the carefully selected teacher model and student model not only have huge differences in model size, but also process of dramatic distinctions in the backbone structure. Unlike other knowledge distillation studies that only focus on the same structure with

different number of layers, our task is more challenging and this setup can better validate the practicality and applicability of the KD methods in the field of dermoscopy imaging.

We implemented our proposed SSD-KD method using PyTorch (Paszke, Gross, Massa, Lerer, Bradbury, Chanan, Killeen, Lin, Gimelshein, Antiga et al., 2019) on one NVIDIA Tesla V100 GPU card (32 GB memory) card at the UBC ARC Sockeye, which is a high-performance computing platform available to UBC researchers across all disciplines. For a total of 25331 images from the ISIC2019, following general practices, we randomly selected 10% of the images to compose a test set, while the rest images were also randomly divided into training and validation with a ratio of 8: 2. Before feeding images into the deep learning models, we uniformly resized all the images to 224×224 and applied data augmentation using common image processing operations, including horizontal and vertical flips, adjustments in brightness, contrast, saturation, image scaling, and random noise. The teacher and student models were pre-trained on ImageNet (Deng, Dong, Socher, Li, Li and Fei-Fei, 2009) with the weights gathered from torchvision API and then fine-tuned for 150 epochs using the mini-batch SGD optimizer with an initial learning rate of 0.001, momentum of 0.9, and the weight decay of 0.001. For hyperparameters in D-KD and SSD-KD, we used default values for parameters within modules BLKD (Hinton et al., 2015), DRKD (Park et al., 2019), and SSKD (Xu et al., 2020). Then for D-KD, we set λ_{blkd} , λ_{drkd} , λ_{crkd} as 1, 1, and 1000, respectively. The parameters were tuned on the validation dataset with grid search. In SSD-KD, since the SSKD module already included BLKD in it (Xu et al., 2020), we therefore treated SSKD as a whole module and used their default parameters. We then set λ_{drkd} , λ_{crkd} , λ_{sskd} as 1, 1000 and 1, respectively.

In the training phase, we used a conditional reduction strategy for the adjustment of the learning rate. If the model did not improve for 10 consecutive epochs, the learning rate value was reduced by a rate of 0.1. Early stopping was also used for every 15 consecutive epochs if the model did not improve. In the initial experiments, we found that increasing the batch size within a certain range will significantly improve the learning ability of the model. Therefore, while considering the computational cost, on the premise of obtaining a better teacher model, we set the batch size as 128. Additionally, because of the severe data imbalance problems between different skin diseases, we used the weighted cross-entropy as the loss function. The weights are determined according to the frequency of diagnosis labels and thus assigning more attention to the classes that less appeared in the training dataset.

3.3. Evaluation metrics

We employed four widely used metrics to evaluate the prediction performance for multi-classes skin diseases of our proposed method. The metrics include Accuracy (ACC), Balanced Accuracy (BACC), mean Average Precision (mAP), and Area Under the Curve (AUC).

3.4. Comparison of the teacher and student models

To clearly indicate the differences between the teacher model and the student model in various indicators, we compared the classification accuracy along with the computational costs in detail. In Table 1, we measured the computational costs by four commonly used complexity standards, including model size, number of parameters, time for every inference step, and the flops. Both teacher model and student model were trained without any knowledge of distillation. As shown in Table 1, the teacher model achieved much better performance in skin lesion classification with an accuracy of 82.0% , while the lightweight student model only got 75.4% in the same task. However, the model size of the teacher model reached 98M and the number of parameters was 25.6M, which was over 6 times that of the student model. In addition, the time for each inference step of the teacher model was 58.2ms, while the student model only used half of this time. The flops, which were used to measure the number of operations of the network, indicated that the teacher model owned over 10 times of complexity in terms of the operation numbers. The above data comparison showed that the teacher model can indeed achieve an essential improvement in the model's predictive ability, but it is also accompanied by a greater computational requirement. This is well aligned with our assumptions for this project.

3.5. Comparison with other KD methods

To demonstrate the effectiveness of our proposed method, several popular representative knowledge distillation methods were listed for comparison. All these listed methods followed exactly the same training settings as introduced above to make sure the comparison was fair and the results were listed in Table 2, Figure 3, and Figure 4. As we can see, all knowledge distillation methods greatly improved the prediction performance of the student model, where the accuracy improvement ranged from 5% to 9% and the balanced accuracy improvement ranged from 4% to 8%. The increase in mean average precision even reached a maximum of 18%. These results were comparable to the teacher model, and some parts were even higher than the teacher model. It proved that the KD method can improve the representative ability of the lightweight student model with the help of the teacher model in the classification task for multiple skin diseases via the dermoscopy imaging technique. Among all the methods listed in Table 2, we also found that our proposed method achieved the best performance in ACC, BACC, and mAP, and the second-best in AUC. Especially in the balanced accuracy, the obvious improvement indicated that our method is more suitable for the classification of multiple skin diseases with unbalanced sample numbers. In Figure 3 and Figure 4, we presented the detailed confusion matrix of prediction and the ROC curves along with the AUC values for each diseases, which also supported our above analysis.

Table 1

The comparison of the prediction accuracy and computation costs on ISIC 2019 between the teacher model(ResNet50) and the student model (MobileNetV2).

Model	Input Size	ACC	Model Size	Number of parameters	Time per inference step	GFlops
ResNet50	224 × 224	0.820	98MB	25.6M	58.2ms	4.11
MobileNetV2	224 × 224	0.754	16MB	3.5M	22.9ms	0.31

Table 2

Comparison of single teacher/models and other state-of-the-art knowledge distillation methods. Best results are highlighted in bold for each column.

Teacher	Student	KD Method	ACC	BACC	AUC	mAP
ResNet50	N/A	N/A	0.820	0.822	0.975	0.740
MobileNetV2	N/A	N/A	0.754	0.767	0.959	0.625
ResNet50	MobileNetV2	BLKD (Hinton et al., 2015)	0.838	0.824	0.976	0.751
ResNet50	MobileNetV2	FitNet (Romero et al., 2015)	0.807	0.806	0.971	0.700
ResNet50	MobileNetV2	DRKD (Park et al., 2019)	0.816	0.804	0.972	0.744
ResNet50	MobileNetV2	SSKD (Xu et al., 2020)	0.840	0.840	0.978	0.759
ResNet50	MobileNetV2	SSKD+DRKD+CRKD (ours)	0.846	0.843	0.977	0.796

3.6. Ablation studies

As introduced in the methodology section, our proposed method was based on BLKD, while adding the distilled knowledge between instances and within the channels, along with a self-supervised learning strategy. Therefore, it was essential to explore the roles and effects of different components of the overall framework. In Table 3, we conducted ablation studies to compare the performance of different methods with different modules. In the training phase, we also fixed all the training optimization parameters and settings to make sure the comparison was fair.

Firstly, we removed self-supervision mechanism and trained BLKD along with DRKD and CRKD separately as shown in the first two rows of the Table 3. Compared with BLKD, we found that both DRKD and CRKD methods consistently improved the accuracy, balanced accuracy, AUC, mAP values. Based on this result, we drew preliminary conclusions that the incorporation of relational-based KD can benefit the predictive ability of the student model by introducing more knowledge. Next, we involved the self-supervision in BLKD, which is noted as SSKD in the Table 2 and Table 3. DRKD and CRKD methods were also considered individually in this scenario. According to the

third and forth rows of the comparison, we noticed that the self-supervised method further improved the results while DRKD and CRKD methods were included, especially the ACC and BACC. This results demonstrated that while introducing DRKD and CRKD method to knowledge distillation, the adoption of the SSKD module can achieve further improvement of representative ability of the student model. Finally, we carried out another two sets of experiments that introduced both DRKD and CRKD modules and trained them with and without the self-supervision strategy to further validate our hypothesis. As shown in the last two rows of the Table 3, our jointly optimized knowledge distillation methods achieved the best performance with the highest prediction accuracy and balanced accuracy, which agreed with our preliminary conclusions mentioned above. Each component of our architecture contributed to the improved performance of the lightweight student model, thus making the proposed method more effective.

3.7. Visualizations of class activation maps

In addition to the analysis and comparisons, we further observed the prediction performance of the test set samples before and after using our knowledge distillation method. We noticed that the skin lesion classes of many samples in

Table 3

Comparison of different modules equipped in SSD-KD for ablation studies. Best results are highlighted in bold for each column.

Teacher	Student	KD Method	ACC	BACC	AUC	mAP
ResNet50	MobileNetV2	BLKD+DRKD	0.839	0.827	0.977	0.755
ResNet50	MobileNetV2	BLKD+CRKD	0.839	0.824	0.977	0.773
ResNet50	MobileNetV2	SSKD+DRKD	0.841	0.840	0.978	0.769
ResNet50	MobileNetV2	SSKD+CRKD	0.842	0.839	0.977	0.773
ResNet50	MobileNetV2	BLKD+DRKD+CRKD (D-KD)	0.844	0.839	0.978	0.788
ResNet50	MobileNetV2	SSKD+DRKD+CRKD (SSD-KD)	0.846	0.843	0.977	0.796



Figure 3: Visualizations of confusion matrices from skin lesion classification models that are trained: without KD (as shown in subfigures (a) – (b)), and with different KDs (as shown in subfigures (c) – (l)).

the test set were correctly predicted by the teacher model and the student model with knowledge distillation, but are diagnosed as the wrong classes by the student model without the knowledge distillation. For example, 17 melanomas were predicted as benign keratoses by the student model without any knowledge distillation but they were classified correctly after the implementation of our KD method. To show how the KD methods helped the learning of the student model, we visualized the class activation mapping (CAM) from the

teacher model, the student model before the implementation of KD, and the student model after using KD applying Grad-CAM (Selvaraju, Cogswell, Das, Vedantam, Parikh and Batra, 2017), which was commonly used to locate discriminative regions for object detection and classification tasks. As presented in Figure 5, the located discriminative areas of student model after the KD implementation was more align with the teacher model's activation maps; Hence, it enhanced the performance of the student model.

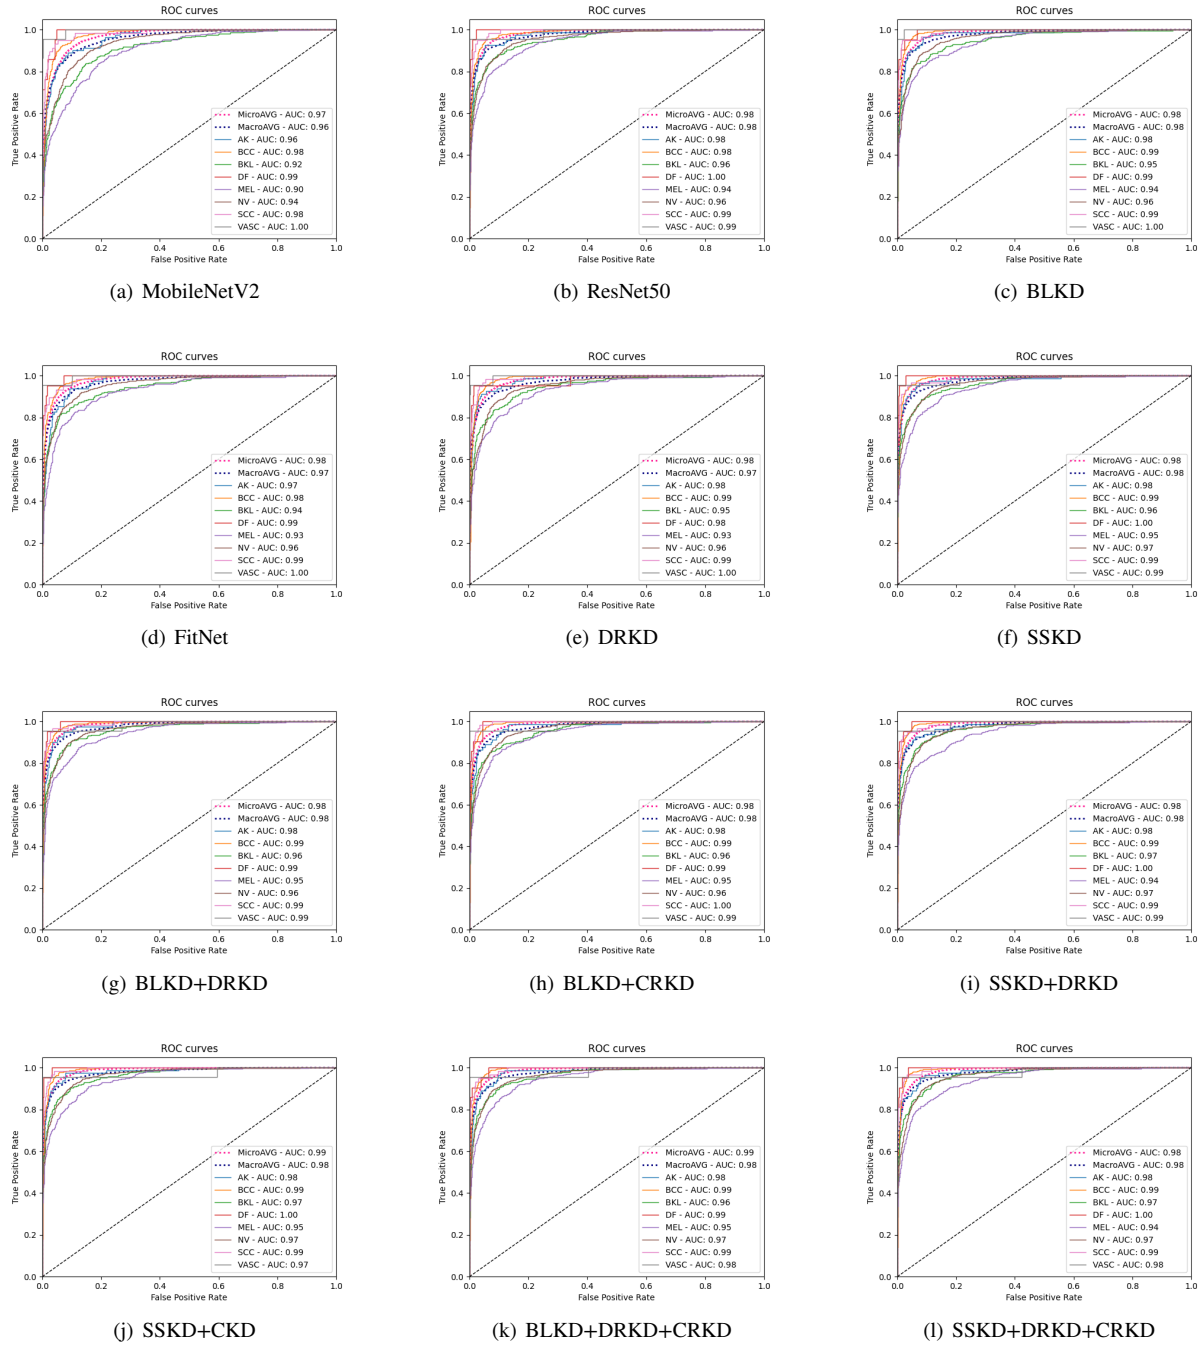


Figure 4: Visualization of ROC curves from skin lesion classification models that are trained: without KD (as shown in subfigures (a) – (b)), and with different KDs (as shown in subfigures (c) – (l)).

4. Discussion and future works

The multi-class classification task has made great progress in recent years due to the in-depth research of medical imaging and deep learning technology. However, excellent model capabilities are often accompanied by complex structures and high training and storage costs, which limit their practical applications in the scenarios with restricted computational resources, such as portable devices. Correspondingly, although the small lightweight model can

greatly reduce the computational cost, it cannot extract more representative features, so that the predictive ability of the model is degraded. For example, in this paper, both ResNet50 and MobileNetV2 are trained on the large-scale dataset of ISIC. However, the prediction accuracy of ResNet50 can reach 82%, while the lightweight MobileNet can only reach 75%. It can be seen that using the KD methods to improve the learning performance of lightweight models




Figure 5: Visualization of class activation mapping (CAM). The example images are correctly predicted by the teacher model and the student model only after implementing our proposed method. (a) is the original dermoscopy image while the other three subfigures (b)-(d) denote their corresponding CAM output from different models.

while maintaining the complexity advantage is an important direction in both theoretical and application levels.

Considering the morphological characteristics of skin lesions, we added a number of modules based on the traditional KD framework, including the widely used DRKD and SSKD, as well as CRKD and weighted soften knowledge that were specifically designed for skin lesions classification via dermoscopy imaging. While the number of skin cancer lesions is far less than that of benign skin lesions, in our proposed method, we integrated weighted softened output and generic feature representations in a self-supervised manner, and also improved the ability of a lightweight model to obtain more useful information from classes with fewer subjects. Therefore, our proposed method can better help the lightweight student model to extract features, thereby achieving improved results. In addition, when we performed ablation studies, we also noticed that although the KD methods improved the student model substantially, the improvement brought by adding different modules to BLKD was not very obvious. To verify that the presented results are sufficiently convincing, we tried ten experiments under the same conditions and calculated the standard deviation of all calculated metrics. The standard deviation values were all less than 0.01, demonstrating the robustness of the trends shown by our results. In fact, further tuning of the hyperparameters should make the model perform better when adding a combination of different modules, but in this paper we fixed the hyperparameter values of all modules for a completely fair comparison.

Additionally, after a series of pre-experiments, we selected the two most representative models, based on their distinct characteristics, as the teacher model and the student model. As the shown in table 1, our criteria for selecting models was to ensure that the teacher model has excellent performance, while the complexity of the student model is suitable for practical portable applications. Although we achieved our expectation of improving the predictive power of the student model, we strongly believed that this was only a starting point for this line of research. In our future works, we plan to further introduce other architectures, such as multi-teacher KD and KD with the help of the teaching assistant mechanism, so as to improve the generalizability of the obtained results. Moreover, our proposed method can also be applied to other medical problems with similar

scenarios after fine-tuning, helping to improve lightweight models in other areas.

5. Conclusion

In this paper, we proposed an efficient knowledge distillation framework, called SSD-KD, which aimed to address the challenging task of multi-class classification of dermoscopic skin lesion imaging by lightweight deep learning models. Different from existing KD methods, SSD-KD considered both intra- and inter-instance based knowledge to guide the lightweight model capture richer structural knowledge that are important for skin lesion diagnosis. Additionally, the self-supervised auxiliary training strategy and weighted cross entropy loss were also involved to further boost its performance.

We conducted experiments to evaluate our proposed method on the large-scale open-accessed dermoscopic skin lesion imaging dataset ISIC 2019, and we presented detailed comparisons with the existing KD methods. We also applied a CAM method to visualize the changes of activation maps generated by the models. Experiments demonstrated the effectiveness of our proposed SSD-KD with a promising result in multi-diseases classification. With low model complexity and storage requirements while improving predictive performance, our method could greatly assist in the promotion of deep learning techniques to clinical practices.

CRediT authorship contribution statement

Yongwei Wang: Conceptualization, Methodology, Data Curation, Software, Writing - Original Draft, Writing - Review & Editing. **Yuheng Wang:** Conceptualization, Methodology, Formal analysis, Validation, Writing - Original Draft, Writing - Review & Editing. **Tim K. Lee:** Supervision, Project administration, Funding acquisition, Writing - Review & Editing. **Chunyan Miao:** Supervision, Project administration, Funding acquisition, Writing - Review & Editing. **Z. Jane Wang:** Supervision, Project administration, Funding acquisition, Writing - Review & Editing.

References

- Abbas, Q., Celebi, M.E., 2019. Dermodeep-a classification of melanoma-nevus skin lesions using multi-feature fusion of visual features and deep neural network. *Multimedia Tools and Applications* 78, 23559–23580.

- Abbasi, S., Hajabdollahi, M., Khadivi, P., Karimi, N., Roshandel, R., Shirani, S., Samavi, S., 2021. Classification of diabetic retinopathy using unlabeled data and knowledge distillation. *Artificial Intelligence in Medicine* 121, 102176.
- Apalla, Z., Lallas, A., Sotiriou, E., Lazaridou, E., Ioannides, D., 2017. Epidemiological trends in skin cancer. *Dermatology practical & conceptual* 7, 1.
- Argenziano, G., Fabbrocini, G., Carli, P., De Giorgi, V., Sammarco, E., Delfino, M., 1998. Epiluminescence microscopy for the diagnosis of doubtful melanocytic skin lesions: comparison of the abcd rule of dermoscopy and a new 7-point checklist based on pattern analysis. *Archives of dermatology* 134, 1563–1570.
- Back, S., Lee, S., Shin, S., Yu, Y., Yuk, T., Jong, S., Ryu, S., Lee, K., 2021. Robust skin disease classification by distilling deep neural network ensemble for the mobile diagnosis of herpes zoster. *IEEE Access* 9, 20156–20169.
- Barata, C., Celebi, M.E., Marques, J.S., 2018. A survey of feature extraction in dermoscopy image analysis of skin cancer. *IEEE journal of biomedical and health informatics* 23, 1096–1109.
- Barata, C., Celebi, M.E., Marques, J.S., 2021. Explainable skin lesion diagnosis using taxonomies. *Pattern Recognition* 110, 107413.
- Bi, L., Feng, D.D., Fulham, M., Kim, J., 2020. Multi-label classification of multi-modality skin lesion via hyper-connected convolutional neural network. *Pattern Recognition* 107, 107502.
- Brinker, T.J., Hekler, A., Utikal, J.S., Grabe, N., Schadendorf, D., Klode, J., Berking, C., Steeb, T., Enk, A.H., Von Kalle, C., 2018. Skin cancer classification using convolutional neural networks: systematic review. *Journal of medical Internet research* 20, e11936.
- Celebi, M.E., Codella, N., Halpern, A., 2019. Dermoscopy image analysis: overview and future directions. *IEEE journal of biomedical and health informatics* 23, 474–478.
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G., 2020. A simple framework for contrastive learning of visual representations, in: *International conference on machine learning*, PMLR. pp. 1597–1607.
- Chen, W., Gao, L., Li, X., Shen, W., 2022. Lightweight convolutional neural network with knowledge distillation for cervical cells classification. *Biomedical Signal Processing and Control* 71, 103177.
- Codella, N.C., Gutman, D., Celebi, M.E., Helba, B., Marchetti, M.A., Dusza, S.W., Kallou, A., Liopyris, K., Mishra, N., Kittler, H., et al., 2018. Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic), in: *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*, IEEE. pp. 168–172.
- Combalia, M., Codella, N.C., Rotemberg, V., Helba, B., Vilaplana, V., Reiter, O., Carrera, C., Barreiro, A., Halpern, A.C., Puig, S., et al., 2019. Bcn20000: Dermoscopic lesions in the wild. *arXiv preprint arXiv:1908.02288*.
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database, in: *2009 IEEE conference on computer vision and pattern recognition*, Ieee. pp. 248–255.
- Ding, L., Wang, Y., Yuan, K., Jiang, M., Wang, P., Huang, H., Wang, Z.J., 2021a. Towards universal physical attacks on single object tracking, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 1236–1245.
- Ding, X., Wang, Y., Wang, Z.J., Welch, W.J., 2021b. Efficient subsampling for generating high-quality images from conditional generative adversarial networks. *arXiv preprint arXiv:2103.11166*.
- Ding, X., Wang, Y., Xu, Z., Wang, Z.J., Welch, W.J., 2021c. Distilling and transferring knowledge via cgan-generated samples for image classification and regression. *arXiv preprint arXiv:2104.03164*.
- Ding, X., Wang, Y., Xu, Z., Welch, W.J., Wang, Z.J., 2021d. Ccgan: Continuous conditional generative adversarial networks for image generation, in: *International Conference on Learning Representations*, pp. 1–10.
- Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S., 2017. Dermatologist-level classification of skin cancer with deep neural networks. *nature* 542, 115–118.
- Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W., 2018. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231*.
- Gessert, N., Nielsen, M., Shaikh, M., Werner, R., Schlaefel, A., 2020. Skin lesion classification using ensembles of multi-resolution efficientnets with meta data. *MethodsX* 7, 100864.
- Gou, J., Yu, B., Maybank, S.J., Tao, D., 2021. Knowledge distillation: A survey. *International Journal of Computer Vision* 129, 1789–1819.
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R., 2020. Momentum contrast for unsupervised visual representation learning, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9729–9738.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.
- Hinton, G., Vinyals, O., Dean, J., et al., 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* 2.
- Huber, P.J., 1992. Robust estimation of a location parameter, in: *Breakthroughs in statistics*. Springer, pp. 492–518.
- Jing, L., Tian, Y., 2020. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence* 43, 4037–4058.
- Johnson, J., Alahi, A., Fei-Fei, L., 2016. Perceptual losses for real-time style transfer and super-resolution, in: *European conference on computer vision*, Springer. pp. 694–711.
- Kawahara, J., BenTaieb, A., Hamarneh, G., 2016. Deep features to classify skin lesions, in: *2016 IEEE 13th international symposium on biomedical imaging (ISBI)*, IEEE. pp. 1397–1400.
- Kawahara, J., Daneshvar, S., Argenziano, G., Hamarneh, G., 2018. Seven-point checklist and skin lesion classification using multitask multimodal neural nets. *IEEE journal of biomedical and health informatics* 23, 538–546.
- Lin, M., Chen, Q., Yan, S., 2014. Network in network, in: *International Conference on Learning Representations*, pp. 1–10.
- Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., Van Der Laak, J.A., Van Ginneken, B., Sánchez, C.I., 2017. A survey on deep learning in medical image analysis. *Medical image analysis* 42, 60–88.
- Liu, H., HaoChen, J.Z., Gaidon, A., Ma, T., 2022. Self-supervised learning is more robust to dataset imbalance, in: *International Conference on Learning Representations*, pp. 1–8.
- Liu, Y., Cao, J., Li, B., Yuan, C., Hu, W., Li, Y., Duan, Y., 2019. Knowledge distillation via instance relationship graph, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7096–7104.
- Nachbar, F., Stolz, W., Merkle, T., Cognetta, A.B., Vogt, T., Landthaler, M., Bilek, P., Braun-Falco, O., Plewig, G., 1994. The abcd rule of dermoscopy: high prospective value in the diagnosis of doubtful melanocytic skin lesions. *Journal of the American Academy of Dermatology* 30, 551–559.
- Van den Oord, A., Li, Y., Vinyals, O., 2018. Representation learning with contrastive predictive coding. *arXiv e-prints*, arXiv:1807.
- Pacheco, A.G., Krohling, R.A., 2021. An attention-based mechanism to combine images and metadata in deep learning models applied to skin cancer classification. *IEEE journal of biomedical and health informatics* 25, 3554–3563.
- Park, W., Kim, D., Lu, Y., Cho, M., 2019. Relational knowledge distillation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3967–3976.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al., 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32.
- Peng, B., Jin, X., Liu, J., Li, D., Wu, Y., Liu, Y., Zhou, S., Zhang, Z., 2019. Correlation congruence for knowledge distillation, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5007–5016.

- Qin, D., Bu, J.J., Liu, Z., Shen, X., Zhou, S., Gu, J.J., Wang, Z.H., Wu, L., Dai, H.F., 2021. Efficient medical image segmentation based on knowledge distillation. *IEEE Transactions on Medical Imaging* 40, 3820–3831.
- Raghu, M., Zhang, C., Kleinberg, J., Bengio, S., 2019. Transfusion: Understanding transfer learning for medical imaging. *Advances in neural information processing systems* 32.
- Romero, A., Ballas, N., Kahou, S.E., Chassang, A., Gatta, C., Bengio, Y., 2015. Fitnets: Hints for thin deep nets. *International Conference on Learning Representations*.
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C., 2018. Mobilenetv2: Inverted residuals and linear bottlenecks, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4510–4520.
- Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization, in: *Proceedings of the IEEE international conference on computer vision*, pp. 618–626.
- Srinivasu, P.N., SivaSai, J.G., Ijaz, M.F., Bhoi, A.K., Kim, W., Kang, J.J., 2021. Classification of skin disease using deep learning neural networks with mobilenet v2 and lstm. *Sensors* 21, 2852.
- Tang, P., Yan, X., Nan, Y., Xiang, S., Krammer, S., Lasser, T., 2022. Fusionm4net: A multi-stage multi-modal learning algorithm for multi-label skin lesion classification. *Medical Image Analysis* 76, 102307.
- Toğaçar, M., Cömert, Z., Ergen, B., 2021. Intelligent skin cancer detection applying autoencoder, mobilenetv2 and spiking neural networks. *Chaos, Solitons & Fractals* 144, 110714.
- Tschandl, P., Rosendahl, C., Kittler, H., 2018. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data* 5, 1–9.
- Tung, F., Mori, G., 2019. Similarity-preserving knowledge distillation, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1365–1374.
- Wang, L., Yoon, K.J., 2021. Knowledge distillation and student-teacher learning for visual intelligence: A review and new outlooks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wang, Y., Cai, J., Louie, D.C., Wang, Z.J., Lee, T.K., 2021. Incorporating clinical knowledge with constrained classifier chain into a multimodal deep network for melanoma detection. *Computers in Biology and Medicine* 137, 104812.
- Xie, X., Niu, J., Liu, X., Chen, Z., Tang, S., Yu, S., 2021. A survey on incorporating domain knowledge into deep learning for medical image analysis. *Medical Image Analysis* 69, 101985.
- Xu, G., Liu, Z., Li, X., Loy, C.C., 2020. Knowledge distillation meets self-supervision, in: *European Conference on Computer Vision*, Springer. pp. 588–604.
- Yim, J., Joo, D., Bae, J., Kim, J., 2017. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4133–4141.
- Zhang, J., Xie, Y., Wu, Q., Xia, Y., 2019. Medical image classification using synergic deep learning. *Medical image analysis* 54, 10–19.