Selective Alignment Transfer for Domain Adaptation in Skin Lesion Analysis

Nurjahan Sultana, Wenqi Lu(✉), Xinqi Fan, and Moi Hoon Yap

Department of Computing and Mathematics, Manchester Metropolitan University, Manchester, UK
nurjahan.sultana@stu.mmu.ac.uk,{w.lu, x.fan, m.yap}@mmu.ac.uk

Abstract. Domain adaptation is crucial for deep learning in skin lesion analysis because models trained on dermoscopic images often struggle to generalise to clinical images, which exhibit variations in lighting, resolution, and background conditions. We propose Selective Alignment Transfer for Domain Adaptation (SAT-DA), a fully supervised framework that significantly reduces this domain gap by dynamically assigning feature importance weights based on statistical moments from both source and target domains. SAT-DA emphasises domain-invariant features and suppresses domain-specific noise to preserve crucial diagnostic cues. Our multi-loss strategy combines classification, alignment, and diversity losses to optimise feature selection and prevent feature collapse onto a narrow set. SAT-DA was evaluated on six public datasets comprising dermoscopic and clinical images and consistently outperformed state-of-the-art supervised and unsupervised methods. On Derm7pt-Derm to Derm7pt-Clinic, SAT-DA achieves 82.46% AUROC, surpassing the strongest baseline by over 6%. Notably, SAT-DA also maintains high performance on completely unseen datasets not used as source or target, demonstrating robust cross-domain generalisation. Overall, these results highlight SAT-DA's ability to address practical clinical deployment challenges, offering a reliable, fully supervised solution for cross-domain skin lesion analysis. The complete implementation of the SAT-DA method is available at our GitHub repository.

Keywords: Domain Adaptation · Skin Lesion Analysis · Fully Supervised Learning.

1 Introduction

Deep learning models have shown remarkable results for skin cancer classification, frequently outperforming dermatologists in controlled settings [5]. However, these advancements have not been fully translated to clinical practice, largely due to the phenomenon of domain shift. Specifically, models trained on standardised dermoscopic images often struggle to generalise to clinical images captured

via mobile devices, which exhibit significant variability in lighting conditions, imaging perspectives, and skin tone representation [2,24,20].

Although unsupervised domain adaptation (UDA) frameworks such as Domain-Adversarial Neural Networks (DANN) [7], Adversarial Discriminative Domain Adaptation (ADDA) [23], and Deep CORAL [21] have been proposed to mitigate distribution discrepancies, their applicability in medical imaging remains challenging. These methods rely on unlabeled target domain data and often struggle to preserve diagnostic class-specific features that are crucial for medical tasks, potentially compromising model interpretability and clinical reliability [3,14].

In contrast, fully supervised domain adaptation (SDA) integrates the availability of labeled target domain data, enabling more effective alignment between source and target distributions. SDA takes advantage of labeled data in both domains to guide alignment. It groups features from the same class while filtering out irrelevant noise. This helps preserve diagnostic cues across domains and improves classification in real-world clinical settings. By using label information, SDA ensures class-wise consistency and enhances the structure of the feature space, which reduces the impact of domain-specific artefacts and supports more reliable diagnostic performance [16,19,12]. This approach can outperform UDA, particularly in high-stakes domains such as oncology, where diagnostic accuracy and model reliability are paramount[17,11,1]. Early supervised approaches aligned feature distributions by minimising class-wise distance measures or incorporating few labeled target samples [22,17]. For instance, Tzeng et al. [22] combined a domain confusion loss with soft label matching for partial supervision, and Motiian et al. [17] introduced contrastive semantic alignment to handle few-shot targets. However, these methods typically do not address the significant domain gap between dermoscopic and clinical/mobile skin images. Clinical images have varied lighting and uncontrolled backgrounds that produce domain-specific artifacts, unlike the relatively homogeneous conditions in standard benchmarks such as Office-31 or Office-Home [15,6]. Prior studies on dermoscopic adaptation mostly focus on biological or technical factors unique to dermoscopy [2], rather than on the dermoscopic-to-clinical image gap. Carretero et al. [1] proposed a supervised contrastive learning method to address domain shifts caused by staining variations in multi-centre histopathological datasets. However, it differs from those encountered in broader imaging tasks, such as adapting dermoscopic models to clinical images with variations in resolution and lighting.

To address this challenge, we propose a fully supervised domain adaptation framework that improves classification on both dermoscopic and clinical images. Our method uses a dynamic feature selection mechanism based on statistical moments from both domains to assign feature importance weights. Unlike purely adversarial [7] or correlation-based methods [21], our approach highlights domain-invariant features that retain essential diagnostic details while suppressing noise from differences in lighting, resolution, viewpoints, and skin tone. We further incorporate a diversity loss to maintain a broad range of features, which is

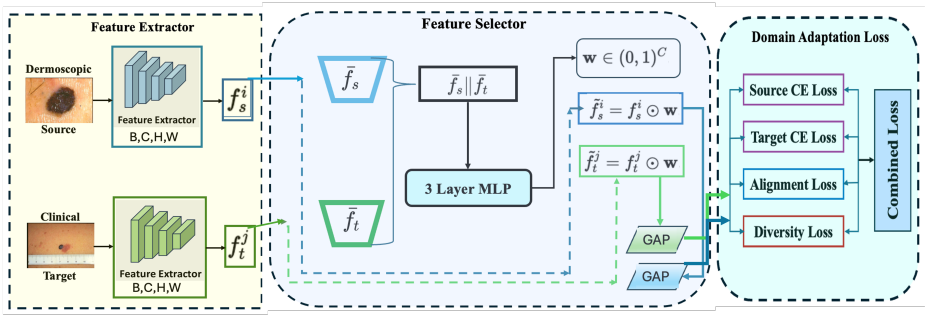

Fig. 1. Architecture of SAT-DA. Feature extraction part processes input images from the source (dermoscopic) and target (clinical) domains using a pretrained model to extract feature maps. Mean feature vectors are pooled for each domain, and a feature selector generates importance weights to emphasise shared features and suppress domain-specific noise. Adjusted feature maps are passed to a classifier for lesion prediction. During training, cross-entropy, feature alignment, and weight diversity losses are used to ensure effective domain adaptation. In inference, stored weights are applied for classification.

crucial in medical imaging where even subtle details like irregular lesion borders can be critical for diagnosis [2]. Our main contributions are as follows:

- **Dynamic feature weighting via statistical moments** We introduce a mechanism to learn domain-specific importance weights to emphasise invariant features and suppress domain-specific noise, substantially enhancing alignment across dermoscopic and mobile images.
- **Multi-loss strategy to preserve diversity** We add a diversity loss alongside classification and alignment losses to prevent reliance on only a few features. This helps capture even small diagnostic cues that are vital in medical imaging.
- **improved cross domain generalisation** Our model ensures reliable cross-domain generalisation, excelling in the target domain and adapting well to unseen datasets.

2 Methodology

An overview of the proposed network for fully supervised domain adaptation in skin lesion analysis is illustrated in Figure 1. The network primarily comprises three parts: 1) in the feature extraction, a pre-trained backbone network is used to map input images into high-level feature representations; 2) in the feature selector, the extracted features are fused and processed to compute channel-wise feature weights; 3) the computed feature weights are applied to both feature representations and a novel classification loss is optimised to ensure domain-invariant class label prediction.

Problem Formulation and Notation: We address the challenge of domain shift between two types of skin images: dermoscopic (source) and clinical (target). We denote the source dataset as $\mathcal{D}_s = \{(x_s^i, y_s^i)\}_{i=1}^{n_s}$ and the target dataset as $\mathcal{D}_t = \{(x_t^j, y_t^j)\}_{j=1}^{n_t}$, where x_s^i, x_t^j are input images, and y_s^i, y_t^j are their corresponding labels. Our goal is to learn a classifier $G(\cdot)$ that performs well on *both* domains despite significant variations, such as lighting conditions, imaging perspectives, and patient skin tones. Unlike adversarial-based methods (*e.g.*, ADDA [23]) which learn domain-invariant representations via a generator-discriminator framework, or approaches like Deep CORAL [21] that align second-order statistics, our method dynamically weighs features while simultaneously encouraging diversity in the channel-wise feature importance. This is achieved by jointly optimising a novel classification loss with a distribution alignment loss, and a diversity-promoting loss.

Feature extraction: Given pairs of source and target images, representative features $\boldsymbol{h}_s^i, \boldsymbol{h}_t^j \in \mathbb{R}^{B\times C\times H\times W}$ can be extracted using a pretrained feature extractor $F(\cdot)$ where B is the batch size, C is the number of channels and H and W are the size of each feature channel. The objective of the feature extraction part is to obtain high-level feature representations associated with essential patterns, textures, and structures. It is important to note that our proposed framework is generic to any modern CNN and not restricted to any particular type of features. In this paper, we instantiate $F(\cdot)$ as EfficientNet-B2, yielding $C = 1408$ feature channels after the final convolutional layer.

Dynamic Feature Selection and Weighting: We apply global average pooling (GAP) to the convolutional feature maps $\boldsymbol{h}_s^i, \boldsymbol{h}_t^j$ and generate spatially pooled feature vectors $\boldsymbol{f}_s^i, \boldsymbol{f}_t^i \in \mathbb{R}^{B\times C}$. Within each mini-batch, we then compute the *batch mean* for each domain so that each image is represented by a C-dimensional vector:

$$\bar{\boldsymbol{f}}_s \;=\; \frac{1}{B_s}\sum_{i=1}^{B_s} \boldsymbol{f}_s^i, \quad \bar{\boldsymbol{f}}_t \;=\; \frac{1}{B_t}\sum_{j=1}^{B_t} \boldsymbol{f}_t^j, \tag{1}$$

yielding two vectors $\bar{\boldsymbol{f}}_s, \bar{\boldsymbol{f}}_t \in \mathbb{R}^C$. To capture the interaction between these two domains, we concatenate $\bar{\boldsymbol{f}}_s$ and $\bar{\boldsymbol{f}}_t$ to form a $2C$-dimensional vector, which is then passed through a three-layer feed-forward network with ReLU activations. The final sigmoid layer $\sigma(\cdot)$ produces a channel-wise feature weight vector $\boldsymbol{w} \in (0,1)^C$:

$$\boldsymbol{w} = \sigma\Big(W_3\,\mathrm{ReLU}\big(W_2\,\mathrm{ReLU}(W_1\,[\,\bar{\boldsymbol{f}}_s, \bar{\boldsymbol{f}}_t\,] + b_1) + b_2\big) + b_3\Big). \tag{2}$$

Here, W_1, W_2, W_3 and b_1, b_2, b_3 are learnable parameters. We then refine both source and target features by element-wise multiplication with $\boldsymbol{w}$:

$$\tilde{\boldsymbol{f}}_s^i = \boldsymbol{f}_s^i \odot \boldsymbol{w}, \quad \tilde{\boldsymbol{f}}_t^j = \boldsymbol{f}_t^j \odot \boldsymbol{w}. \tag{3}$$

Although the weight vector $\boldsymbol{w}$ is derived from the *batch-averaged* features, we apply it *element-wise* to each individual feature map in the batch, enabling

domain-invariant weighting at the channel level. After weighting, we apply another GAP and feed these refined features into a classifier $G(\cdot)$.

Loss Functions: The output of SAT-DA is the predicted class probabilities. To optimise the model, we propose a novel loss function comprising following three components:

i. Classification Loss: Because we have labels for both source and target domains, we compute cross-entropy losses on both domains:

$$L_{\text{cls}} \;=\; \frac{1}{n_s}\sum_{i=1}^{n_s} \ell\big(G(\tilde{\boldsymbol{f}}_s^{\,i}),\, y_s^i\big) \;+\; \frac{1}{n_t}\sum_{j=1}^{n_t} \ell\big(G(\tilde{\boldsymbol{f}}_t^{\,j}),\, y_t^j\big), \tag{4}$$

where $\ell(\cdot,\cdot)$ denotes the cross-entropy function.

ii. Alignment Loss: To reduce domain discrepancies, we align the weighted feature representations via

$$L_{\text{align}} = \Big\| \mathbb{E}[\tilde{\boldsymbol{f}}_s] \;-\; \mathbb{E}[\tilde{\boldsymbol{f}}_t] \Big\|_2^2, \tag{5}$$

which penalises the squared ℓ_2 distance between domain means. We facilitate domain-invariant learning by ensuring $\tilde{\boldsymbol{f}}_s$ and $\tilde{\boldsymbol{f}}_t$ have similar distributions. Unlike Deep CORAL [21] which aligns second-order statistics (*e.g.*, covariance), we focus on mean alignment for simplicity and computational efficiency.

iii. Diversity Loss: A key novelty in this paper is to diversify the weighting vector $\boldsymbol{w}$ to avoid trivial solutions that focus on a narrow subset of features. The diversity loss prevents the feature selector from relying on limited features and helps capture important diagnostic cues like irregular borders or heterogeneous textures. We encourage weight variability through

$$L_{\text{diverse}} = -\,\text{std}(\boldsymbol{w}), \tag{6}$$

where $\text{std}(\cdot)$ is the standard deviation of all elements in $\boldsymbol{w}$. This negative sign ensures we maximise the standard deviation.

Overall Objective: We combine these three losses into one final training objective:

$$L \;=\; L_{\text{cls}} \;+\; \lambda_1\, L_{\text{align}} \;+\; \lambda_2\, L_{\text{diverse}}, \tag{7}$$

where λ_1 and λ_2 control the relative importance of alignment and diversity, respectively.

3 Experimental Results and Discussions

Datasets: We evaluate our method on six publicly available skin lesion datasets, each containing two classes: melanoma (MEL) and nevus (NEV). Three of these datasets are dermoscopic (ISIC-2017 [10], ISIC-2018 [4], Derm7pt-Derm [13]) and three are clinical (Derm7pt-Clinic [13], Fitzpatrick17k [9], PAD-UFES-20 [18]). We split the Derm7pt dataset into Derm (D7D) and Clinic (D7C) subsets. For all datasets, we adopt stratified splits with 20% for validation and 20% for

Table 1. Overview of the six skin lesion datasets and their two classes.

Dataset	Abbrev.	Type	MEL	NEV	Train / Val / Test
ISIC 2017 [10]	ISIC-17	Derm.	347	2045	0.6 / 0.2 / 0.2
ISIC 2018 [4]	ISIC-18	Derm.	1113	6705	0.6 / 0.2 / 0.2
Derm7pt-Derm [13]	D7D	Derm.	191	475	0.6 / 0.2 / 0.2
Derm7pt-Clinic [13]	D7C	Clinic	191	475	0.6 / 0.2 / 0.2
Fitzpatrick17k [9]	Fitz	Clinic	299	140	0.6 / 0.2 / 0.2
PAD-UFES-20 [18]	PAD	Clinic	52	244	0.6 / 0.2 / 0.2

Table 2. AUROC comparison of SDA, UDA, and SAT-DA, where dermoscopic images act as the source domain and clinical images the target domain, with additional evaluation on four unseen datasets. For D7D→D7C, the target domain is D7C and the unseen datasets are ISIC 2017, ISIC 2018, Fitz, and PAD. For ISIC2017→Fitz, the target domain is Fitz and the unseen datasets are ISIC 2018, D7C, D7D and PAD. For ISIC2018→PAD, the target domain is PAD and the unseen datasets are ISIC 2017, D7D, D7C, and Fitz.

Setting	Dataset	ISIC 2018	ISIC 2017	PAD	Fitz	D7D	D7C	Average
ATDOC [16]	D7D→D7C	70.50 ± 0.35	61.73 ± 0.28	54.19 ± 0.14	71.92 ± 0.56	-	75.96 ± 0.18	66.86 ± 0.34
(SDA)	ISIC2017→Fitz	80.25 ± 0.22	-	70.20 ± 0.33	71.48 ± 0.19	74.02 ± 0.24	57.76 ± 0.43	72.84 ± 0.30
	ISIC2018→PAD	-	68.61 ± 0.31	68.94 ± 0.25	69.30 ± 0.29	74.89 ± 0.20	54.56 ± 0.44	67.26 ± 0.31
MCC [12]	D7D→D7C	51.40 ± 0.35	53.30 ± 0.21	47.17 ± 0.39	52.72 ± 0.44	-	51.65 ± 0.18	56.83 ± 0.31
(SDA)	ISIC2017→Fitz	52.08 ± 0.05	-	76.94 ± 0.41	62.19 ± 0.34	57.49 ± 0.18	55.56 ± 0.39	60.85 ± 0.31
	ISIC2018→PAD	-	54.74 ± 0.30	47.62 ± 0.27	53.41 ± 0.16	65.70 ± 0.15	59.62 ± 0.08	56.22 ± 0.09
LIC [19]	D7D→D7C	78.30 ± 0.05	62.25 ± 0.08	59.46 ± 0.12	74.10 ± 0.01	-	72.30 ± 0.05	69.28 ± 0.03
(SDA)	ISIC2017→Fitz	78.08 ± 0.22	-	67.10 ± 0.03	69.27 ± 0.17	75.60 ± 0.19	72.05 ± 0.12	72.42 ± 0.07
	ISIC2018→PAD	-	71.10 ± 0.23	68.60 ± 0.17	75.68 ± 0.11	71.40 ± 0.21	69.38 ± 0.09	71.23 ± 0.08
DANN [7]	D7D→D7C	74.60 ± 0.26	55.10 ± 0.03	61.47 ± 0.22	69.30 ± 0.12	-	72.65 ± 0.31	66.62 ± 0.21
(UDA)	ISIC2017→Fitz	82.40 ± 0.08	-	76.45 ± 0.17	72.60 ± 0.19	82.20 ± 0.23	62.54 ± 0.04	75.24 ± 0.16
	ISIC2018→PAD	-	62.50 ± 0.10	66.30 ± 0.31	69.40 ± 0.42	78.04 ± 0.21	68.50 ± 0.09	68.95 ± 0.26
ADDA [23]	D7D→D7C	59.30 ± 0.18	64.70 ± 0.31	46.04 ± 0.23	52.85 ± 0.09	-	66.21 ± 0.04	59.30 ± 0.18
(UDA)	ISIC2017→Fitz	63.10 ± 0.11	-	65.25 ± 0.06	46.70 ± 0.03	66.20 ± 0.09	49.50 ± 0.02	58.15 ± 0.07
	ISIC2018→PAD	-	68.30 ± 0.25	70.60 ± 0.07	67.50 ± 0.12	75.10 ± 0.09	57.20 ± 0.08	67.74 ± 0.14
Deep CORAL [21]	D7D→D7C	72.60 ± 0.12	64.22± 0.14	58.50 ± 0.33	74.70 ± 0.25	-	67.10 ± 0.02	67.42 ± 0.19
(UDA)	ISIC2017→Fitz	80.40 ± 0.07	-	80.05 ± 0.04	72.15 ± 0.10	80.20 ± 0.24	61.15 ± 0.10	74.79 ± 0.12
	ISIC2018→PAD	-	70.08 ± 0.04	72.40 ± 0.12	70.30 ± 0.09	73.60 ± 0.11	61.15 ± 0.09	69.51 ± 0.09
SAT-DA	D7D→D7C	**84.40 ± 0.41**	**66.53 ± 0.21**	**65.22 ± 0.19**	**79.62 ± 0.45**	-	**82.46 ± 0.56**	**75.65 ± 0.39**
(Ours)	ISIC2017→Fitz	82.05 ± 0.22	-	**77.77 ± 0.33**	72.58 ± 0.19	79.50 ± 0.51	64.50 ± 0.39	**75.28 ± 0.35**
	ISIC2018→PAD	-	69.52 ± 0.27	**76.68 ± 0.57**	**76.50 ± 0.06**	**83.58 ± 0.39**	**71.51 ± 0.19**	**75.56 ± 0.33**

testing. All images are resized to 288×288 pixels. For Fitzpatrick17k, we additionally remove duplicates and non-lesion entries, then apply region-of-interest cropping before resizing. Table 1 summarises the distribution and abbreviations used throughout the paper.

Training and Evaluation: Our method is implemented in PyTorch (v2.4.0 with CUDA 12.1) and deployed on an NVIDIA GeForce RTX 3090 (24 GB) GPU. We use an Adam optimizer with a learning rate of 1×10^{-4} and set $\lambda_1 = 0.1$ and $\lambda_2 = 0.01$. We benchmark against three supervised models: ATDOC [16], MCC [12], and LIC [19], as well as three unsupervised models: DANN [7], ADDA [23], and Deep CORAL [21]. For evaluation, we use AUROC, accuracy and F1 score.

Comparison with other Methods: In Table 2, we present a comparative analysis of our proposed SAT-DA framework against both SDA and UDA methodologies across multiple domain shifts, utilising AUROC as the evaluation metric.

Notably, SAT-DA achieves performance that is either competitive with or surpasses existing state-of-the-art methods. For D7D→D7C, SAT-DA obtains

82.46 ± 0.56 on the target domain, signifying a significant improvement over the strongest SDA method, ATDOC (75.96 ± 0.18), and the best UDA approach, DANN (72.65 ± 0.31). In ISIC2018→PAD, SAT-DA reaches 76.68 ± 0.57, again indicating a significant improvement over all baselines. Performance margins are closer in ISIC2017→Fitz, where SAT-DA achieves 72.58 ± 0.19, yet its average performance on unseen datasets remains higher than competing methods. We also conducted statistical tests to determine the significance of the results. Since we evaluated each model multiple times on different unseen datasets, we generated a distribution of results for each method. A one-way ANOVA test conducted on these distributions indicated a statistically significant difference between models ($p < 0.05$). Further pairwise comparisons using Tukey's HSD test confirmed that SAT-DA significantly performed better than the baseline models. This further highlights the limitations of current SDA methods, which suffer from domain dependency, and the instability of UDA methods, which fail to adapt consistently across different datasets. We further evaluated each method's ability to reduce domain disparity by calculating the Domain Shift Cover(DSC), as shown in Table 4. DSC quantifies the extent to which the performance gap is bridged when moving from the base model (trained solely on dermoscopic images) to DA models with clinical images. Our SAT-DA achieves the highest DSC across AUROC, Accuracy and F1 Score, demonstrating its superior ability to facilitate robust feature alignment and effective knowledge transfer across domains. By contrast, certain UDA methods (e.g. ADDA) show negative DSC in specific metrics, suggesting instability in purely label-free adaptation. While existing SDA methods benefit from target labels, they remain inferior to SAT-DA in effectively mitigating domain shifts comprehensively.

Table 3. Ablation study of SAT-DA: AUROC results from D7D (source) to D7C (target) and evaluation on four unseen datasets (ISIC 2018, ISIC 2017, PAD, Fitz). F.S: Feature selection. L_{diverse}: diversity loss. L_{align}: alignment loss.

Method	ISIC 2018	ISIC 2017	PAD	Fitz	D7C	Average
W/O (F.S & L_{diverse} & L_{align})	48.46 ± 0.25	47.50 ± 0.37	52.70 ± 0.23	52.00 ± 0.17	55.59 ± 0.46	50.93 ± 0.31
W/O (F.S & L_{diverse})	54.46 ± 0.39	52.30 ± 0.26	62.20 ± 0.07	55.58 ± 0.14	59.60 ± 0.48	56.83 ± 0.31
SAT-DA (Full Application)	84.40 ± 0.41	66.53 ± 0.21	65.22 ± 0.19	79.62 ± 0.45	82.46 ± 0.56	75.65 ± 0.39

Table 4. Domain Shift Covered across different settings on the D7D → D7C dataset. Base model (EfficientNet B2) trained on dermoscopic images and evaluated on clinical images. The Combined setting blends dermoscopic and clinical images, split into training, validation, and test sets. The Domain Shift Cover (DSC) is computed as $\mathbf{DSC}(\%) = \frac{\mathbf{DA\ Model\ Value - Base\ Value}}{\mathbf{Target\ Value - Base\ Value}} \times 100$; Target Value=100 [8].

Setting	**AUROC**		**Accuracy**		**F1 Score**	
Base	68.26		67.81		07.38	
	Absolute	DSC (%)	Absolute	DSC (%)	Absolute	DSC (%)
Combined	73.50	16.55	68.65	2.62	58.45	55.01
ATDOC [16] (SDA)	70.50	7.04	73.44	17.66	35.38	30.23
MCC [12] (SDA)	51.40	-53.12	68.12	0.96	0.02	-7.95
LIC [19] (SDA)	72.35	12.89	71.88	12.64	24.28	18.25
DANN [7] (UDA)	74.60	19.88	70.63	8.76	54.30	50.70
ADDA [23] (UDA)	59.30	-28.31	60.30	-23.33	47.16	42.95
Deep CORAL [21] (UDA)	67.70	-1.76	63.12	-14.57	40.34	35.59
SAT-DA (Ours)	82.46	**44.85**	76.92	**28.30**	60.52	**57.37**

Ablation Study: Table 3 examines the impact of removing individual components—feature selection, diversity loss, and alignment loss—on the performance of SAT-DA. The evaluation is conducted in the context of domain adaptation from D7D (source) to D7C (target) and is further validated on four previously unseen datasets. Without any component, the model barely reaches 50.93% average AUROC, reflecting poor cross-domain adaptation. Adding alignment loss alone raises this to 56.83%, indicating some benefit from reducing distribution gaps but still insufficient for challenging domain shifts. In contrast, the full SAT-DA framework achieves a substantially higher average AUROC of 75.65%, highlighting the critical role of integrating channel-wise feature selection and diversity regularisation with alignment loss.

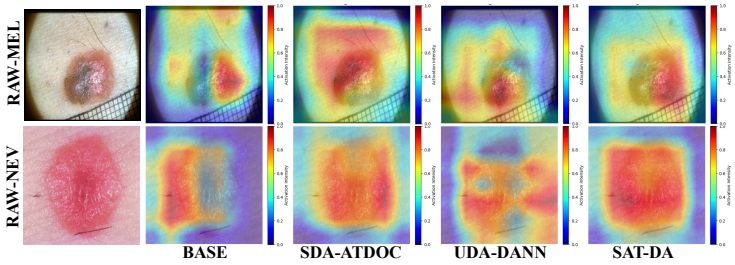

Fig. 2. Visual comparisons on two example images. From left to right: raw clinical lesion image, Grad-CAM heatmaps for four models: Base (no adaptation), SDA (ATDOC), UDA (DANN), and our SAT-DA. The heatmaps highlight each model's attention over the lesion; SAT-DA shows tighter focus on the malignant region.

Qualitative Analysis: We generate Grad-CAM visualisations for the entire D7C test set and observed that, in most cases, our model exhibites superior localisation of lesion areas than other models. Figure 2 presents example images demonstrating that SAT-DA focuses more precisely on the malignant regions, whereas the other models tend to produce broader or misplaced activations.

Limitations: Our model struggles with large domain gaps, particularly with Fitzpatrick data, where variations in image acquisition angles significantly differ from the training distribution, even after ROI preprocessing. This issue is also evident in Grad-CAM visualisations on the D7C test set, where the model sometimes fails to accurately localise the lesion when the image is captured from a significantly different angle.

4 Conclusion

We propose a selective alignment transfer approach to address domain shift in skin lesion analysis. This method focuses on the most discriminative features, thereby preserving crucial clinical information while reducing the influence of confounding factors. Our results demonstrate that strategically guided domain adaptation improves lesion classification performance across diverse datasets. Our approach enhances resilience and accuracy, making it more suitable for real-world clinical applications.

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

References

1. Carretero, I., Meseguer, P., del Amor, R., Naranjo, V.: Enhancing whole slide image classification through supervised contrastive domain adaptation (12 2024), http://arxiv.org/abs/2412.04260
2. Chamarthi, S., Fogelberg, K., Maron, R.C., Brinker, T.J., Niebling, J.: Mitigating the influence of domain shift in skin lesion classification: A benchmark study of unsupervised domain adaptation methods on dermoscopic images (10 2023), http://arxiv.org/abs/2310.03432
3. Chen, C., Dou, Q., Chen, H., Qin, J., Heng, P.A.: Aaai-19) provincial key laboratory of computer vision and virtual reality technology, shenzhen institutes of advanced technology. In: The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19) (2019). https://doi.org/https://doi.org/10.48550/arXiv.1901.08211
4. Codella, N., Rotemberg, V., Tschandl, P., Celebi, M.E., Dusza, S., Gutman, D., Helba, B., Kalloo, A., Liopyris, K., Marchetti, M., 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)
5. Combalia, M., Codella, N., Rotemberg, V., Carrera, C., Dusza, S., Gutman, D., Helba, B., Kittler, H., Kurtansky, N.R., Liopyris, K., Marchetti, M.A., Podlipnik, S., Puig, S., Rinner, C., Tschandl, P., Weber, J., Halpern, A., Malvehy,

J.: Validation of artificial intelligence prediction models for skin cancer diagnosis using dermoscopy images: the 2019 International Skin Imaging Collaboration Grand Challenge. The Lancet Digital Health **4**(5), e330–e339 (5 2022). https://doi.org/10.1016/S2589-7500(22)00021-8

6. Cui, S., Wang, S., Zhuo, J., Su, C., Huang, Q., Tian, Q.: Gradually Vanishing Bridge for Adversarial Domain Adaptation. In: Computer Vision and Pattern Recognition. IEEE, Seattle, WA, USA (2020). https://doi.org/https://doi.org/10.48550/arXiv.2003.13183
7. Ganin, Y., Lempitsky, V., Ru, L.: Unsupervised Domain Adaptation by Backpropagation. In: ICML'15: Proceedings of the 32nd International Conference on International Conference on Machine Learning. pp. 1180–1189 (2015). https://doi.org/https://doi.org/10.48550/arXiv.1409.7495
8. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., March, M., Lempitsky, V.: Domain-adversarial training of neural networks. Journal of machine learning research **17**(59), 1–35 (2016)
9. Groh, M., Harris, C., Soenksen, L., Lau, F., Han, R., Kim, A., Koochek, A., Badri, O.: Evaluating deep neural networks trained on clinical images in dermatology with the fitzpatrick 17k dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1820–1828 (2021)
10. Gutman, D., Codella, N.C., Celebi, E., Helba, B., Marchetti, M., Mishra, N., Halpern, A.: Skin lesion analysis toward melanoma detection: A challenge at the international symposium on biomedical imaging (isbi) 2016, hosted by the international skin imaging collaboration (isic). arXiv preprint arXiv:1605.01397 (2016)
11. Hedegaard, L., Sheikh-Omar, O.A., Iosifidis, A.: Supervised domain adaptation: A graph embedding perspective and a rectified experimental protocol (4 2020). https://doi.org/10.1109/TIP.2021.3118978, http://arxiv.org/abs/2004.11262http://dx.doi.org/10.1109/TIP.2021.3118978
12. Jin, Y., Wang, X., Long, M., Wang, J.: Minimum class confusion for versatile domain adaptation. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16. pp. 464–480. Springer (2020)
13. Kawahara, J., Daneshvar, S., Argenziano, G., Hamarneh, G., Checklist, S.P.: Skin lesion classification using multitask multimodal neural nets. IEEE Journal of Biomedical and Health Informatics **23**, 538–546 (2019)
14. Kumari, S., Singh, P.: Deep learning for unsupervised domain adaptation in medical imaging: Recent advancements and future perspectives. Computers in Biology and Medicine **170** (3 2024). https://doi.org/10.1016/j.compbiomed.2023.107912
15. Li, S., Liu, C.H., Lin, Q., Xie, B., Ding, Z., Huang, G., Tang, J.: Domain conditioned adaptation network. In: The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20) (2020), www.aaai.org
16. Liang, J., Hu, D., Feng, J.: Domain adaptation with auxiliary target domain-oriented classifier. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16632–16642 (2021)
17. Motiian, S., Piccirilli, M., Adjeroh, D.A., Doretto, G.: Unified deep supervised domain adaptation and generalization (9 2017), http://arxiv.org/abs/1709.10190
18. Pacheco, A.G.C., Lima, G.R., Salomão, A.S., Krohling, B.A., Biral, I.P., De Angelo, G.G., Alves, F.C.R., Esgario, J.G.M., Simora, A.C., Castro, P.B.C., Rodrigues, F.B., Frasson, P.H.L., Krohling, R.A., Knidel, H., Santos, M.C.S., Santo, R.B.D.E., Macedo, T.L.S.G., Canuto, T.R.P., De Barros, L.F.S.: PAD-UFES-20:

A skin lesion dataset composed of patient data and clinical images collected from smartphones. Data in brief **32**, 106221 (10 2020). https://doi.org/10.1016/j.dib.2020.106221
19. Presta, A., Spadaro, G., Tartaglione, E., Fiandrotti, A., Grangetto, M.: Domain adaptation for learned image compression with supervised adapters. In: 2024 Data Compression Conference (DCC). pp. 33–42. IEEE (2024)
20. Sultana, N., Lu, W., Fan, X., Yap, M.H.: Domain adaptation for skin lesion: Evaluating real-world generalisation. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 3433–3443 (2025)
21. Sun, B., Saenko, K.: Deep CORAL: Correlation Alignment for Deep Domain Adaptation. In: Computer Vision – ECCV 2016 Workshops. pp. 443–450. Amsterdam, The Netherlands (7 2016). https://doi.org/10.48550/arXiv.1607.01719
22. Tzeng, E., Hoffman, J., Darrell, T., Saenko, K.: Simultaneous deep transfer across domains and tasks. In: Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV). pp. 4068 – 4076 (2015). https://doi.org/https://doi.org/10.1109/ICCV.2015.463
23. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017. vol. 2017-January, pp. 2962–2971. Institute of Electrical and Electronics Engineers Inc. (11 2017). https://doi.org/10.1109/CVPR.2017.316
24. Wang, J., Zhang, Y., Ding, Z., Hamm, J.: Achieving reliable and fair skin lesion diagnosis via unsupervised domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 5157–5166 (June 2024)