

Achieving Reliable and Fair Skin Lesion Diagnosis via Unsupervised Domain Adaptation

Janet Wang, Yunbei Zhang, Zhengming Ding, Jihun Hamm
Tulane University

{swang47, yzhang111, zding1, jhamm3}@tulane.edu

Abstract

The development of reliable and fair diagnostic systems is often constrained by the scarcity of labeled data. To address this challenge, our work explores the feasibility of unsupervised domain adaptation (UDA) to integrate large external datasets for developing reliable classifiers. The adoption of UDA with multiple sources can simultaneously enrich the training set and bridge the domain gap between different skin lesion datasets, which vary due to distinct acquisition protocols. Particularly, UDA shows practical promise for improving diagnostic reliability when training with a custom skin lesion dataset, where only limited labeled data are available from the target domain. In this study, we investigate three UDA training schemes based on source data utilization: single-source, combined-source, and multi-source UDA. Our findings demonstrate the effectiveness of applying UDA on multiple sources for binary and multi-class classification. A strong correlation between test error and label shift in multi-class tasks has been observed in the experiment. Crucially, our study shows that UDA can effectively mitigate bias against minority groups and enhance fairness in diagnostic systems, while maintaining superior classification performance. This is achieved even without directly implementing fairness-focused techniques. This success is potentially attributed to the increased and well-adapted demographic information obtained from multiple sources.

1. Introduction

AI-assisted diagnostic systems have shown expert-level proficiency in classifying skin cancers, conditions typically identified through visual diagnoses [2, 9, 19]. Such systems can potentially contribute to tele dermatology as diagnostic and decision-support tools, enhancing dermatological access in rural areas where medical resources are scarce [5]. However, the development of reliable diagnostic models is often challenged by the scarcity of labeled images—a

common issue in medical imaging analysis. Efforts to address data scarcity, such as leveraging *Generative Adversarial Networks* (GANs) to augment training sets [1, 13, 25], also demand a large volume of labeled data to achieve ideal performance.

Our motivation stems from a prevalent problem in the medical imaging analysis: the development of classifiers for custom medical datasets often faces a lack of sufficient labeled data, with the target domain distribution remaining unknown. To navigate this challenge, public labeled skin lesion datasets can be utilized to enrich the training set's information and diversity. Each dataset represents a distinct domain, due to varying image acquisition protocols and sources. Therefore, effectively managing the domain shift between source and target domains is essential to ensure that features learned from source domains can be successfully applied to target domains, with limited target labeled data available as reference [21, 32].

Unsupervised domain adaptation (UDA) allows deep neural networks to learn features discriminative between classes and invariant across domains by jointly optimizing underlying features and two components: the label predictor and the domain alignment component [12]. Relevant studies [3, 10] have validated UDA's efficacy for binary melanoma-nevus classification on dermoscopic images. Notably, [3] focuses on the investigation of various UDA methods in a single-source, single-target, homogeneous adaptation context, revealing the prominent effectiveness of adversarial-based UDA in binary classification.

Unlike previous research, our study seeks to explore the potential of UDA in maximizing the utility of diverse, public skin lesion data sources through different training schemes: single-source, combined-source, and multi-source UDA. Following [3], we employ two representative adversarial-based UDA methods, ADDA and DANN, to investigate single-source and combined-source training. Given the absence of prior research on multi-source UDA in skin lesion classification, we have chosen two methods representing moment-matching and adversarial-based UDA approaches for multi-source training, M³SDA and MSDN,

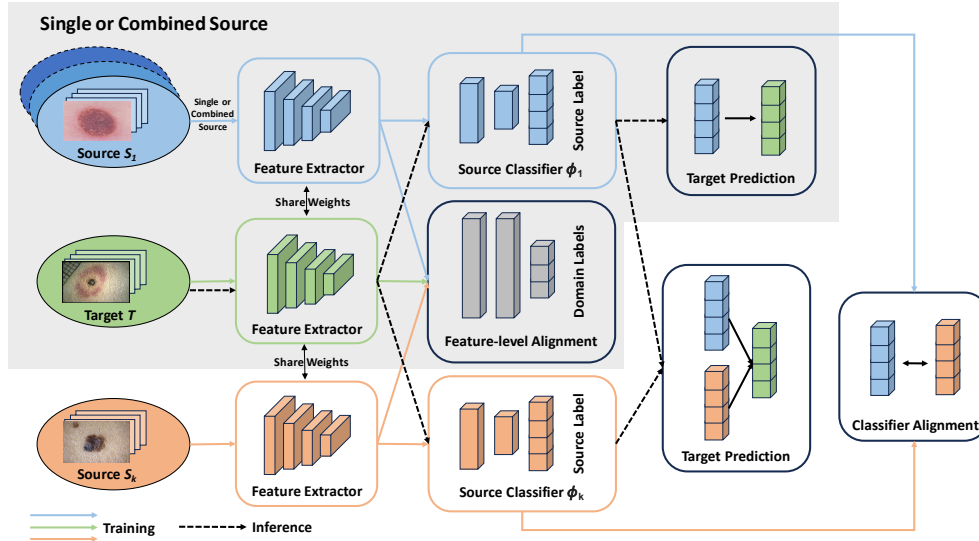


Figure 1. Illustration of single-, combined-, and multi-source UDA training for skin lesion classification. The shaded area demonstrates how single- and combined-source UDA operate. Specifically, combined-source UDA involves an additional step of aggregating multiple datasets into a single source, after which it follows the same training procedure as single-source UDA. The entire figure illustrates the multi-source training scheme, where all source domains (marked in orange and blue) are aligned with the target domain (marked in green) via a domain alignment component (marked in grey). A classifier is trained on each source domain independently, while also aligning with those trained on other source domains. All classifiers will be used to make target predictions during inference.

respectively.

In our experiments, we leverage multiple skin lesion datasets that comprise dermoscopic or clinical images. For each target dataset, we either use a single source or leverage all sources to train a classifier with UDA, under the challenging assumption that labeled data from target domain is inaccessible. If a single-source domain is considered, a corresponding UDA method is applied to adapt it to the target domain. Alternatively, if multiple source domains are considered, we may either combine them for single-source UDA or employ multi-source UDA methods designed for aligning multiple sources. We explore all training schemes in binary and multi-class classification tasks, across the six target sets. Given the richer and well-aligned demographic feature achievable through UDA with multiple training sources, we anticipate a reduced performance gap between groups differentiated by demographic attributes, such as age and skin type. Consequently, we assess the impact of these UDA training schemes on fairness using various metrics. We believe our work holds significant practical value, offering a promising approach to building a reliable diagnostic system on custom datasets, which often suffer from limitation in size and diversity. The key contributions of our work are three-fold:

1. To the best of our knowledge, this is the first systematic evaluation of single- and multi-source UDA methods applied to skin lesion classification across six datasets, for both binary and multi-class tasks.

2. We provide a detailed, quantitative analysis of UDA’s performance in skin lesion classification and find that applying UDA with multiple sources consistently outperforms its single-source counterpart across target domains, for both binary and multi-class classification.

3. To the best of our knowledge, we are the first to assess the effect of UDA on fairness in skin lesion classification. Our analysis reveals that leveraging multiple sources with UDA can effectively ameliorate bias against underrepresented groups.

2. Related Works

2.1. Unsupervised Domain Adaptation

UDA methods aim to mitigate the domain shifts between source and related, yet unlabeled target domains. Existing UDA methods can be categorized based on their knowledge transfer strategies into two primary groups: image alignment and feature alignment. Image alignment uses GANs to translate source images into those that are visually similar to the target domain. In contrast, feature alignment approaches distribution shifts at the feature level. Domain Adversarial Neural Network (DANN) and Adversarial Discriminative Domain Adaptation (ADDA) exemplify adversarial training methods that learn invariant features across domains [11, 31]. There’s a growing interest in adapting multiple resources to an unlabeled domain. For instance, Multi-source Domain Adversarial Networks (MDAN) en-

hance domain adaptation by optimizing task-adaptive generalization bounds [33]. Moreover, the Moment Matching for Multi-Source Domain Adaptation (M³SDA) method transfers knowledge learned from multiple labeled sources to an unlabeled target domain by dynamically aligning moments of their feature distributions [23].

2.2. Domain Adaptation in Skin Lesion Analysis

Prior works have validated the feasibility of DA to enhance performance in skin lesion classification. [15], for example, adopts CycleGAN as a domain adaptation technique to perform invariant attribute translation between skin disease datasets at image level. In their work, source and target domains are aligned to include the same disease classes, under the assumption that target labels are known. On the other hand, [10] focuses on applying single-source UDA to dermoscopic images for melanoma-nevus classification, relaxing the assumption about the target domain. Extending this work, [3] investigates various single-source UDA and finds that adversarial-based UDA training is effective for melanoma-nevus classification.

2.3. Fairness in Skin Lesion Analysis

Recent studies have revealed bias in AI-based diagnosis systems against underrepresented sub-populations [18, 27]. This bias also manifests in skin lesion analysis. Studies have shown that training data with limited skin tone diversity can lead to under-diagnosis of skin cancers in people with darker skin tones, potentially delaying critical treatment and hindering the reliable deployment of AI in tele-dermatology [6, 14]. Furthermore, research suggests that these systems can exhibit performance gaps across different age and gender subgroups, raising concerns about under- and over-diagnosis across groups [10]. While collecting and annotating more demographic-diverse data is ideal, it’s not always practical due to resource limitations. Nevertheless, addressing fairness in AI-based diagnosis still remains crucial in today’s skin lesion analysis.

3. Datasets and Methods

3.1. Dataset

We evaluate the three UDA training schemes using six public and well-annotated skin lesion datasets, consisting of either dermoscopic or clinical images, as shown in Table 1. We specifically separate the dermoscopic and clinical pairs in the Derm7pt dataset into two independent datasets—Derm7pt-derm and Derm7pt-clinic—due to their differing image types. Example images for each dataset can be found in Appendix Sec. 6. We focus on the most common eight skin cancer conditions, whose diagnosis is critical for early screening. Class distributions are detailed in Table 2.

Dataset	Abbrev.	Samples	Type	Diagnostic Classes
ISIC 2018 [4, 30]	isic2018	9688	Dermoscopic	Melanocytic nevus, melanoma, basal cell carcinoma, benign keratosis, dermatofibroma, vascular lesion
ISIC 2020 [26]	isic2020	6000	Dermoscopic	Melanocytic nevus, melanoma, benign keratosis
PAD-UFES-20 [22]	pad	2298	Clinical	Melanocytic nevus, melanoma, basal cell carcinoma, benign keratosis, actinic keratosis, squamous cell carcinoma
Fitzpatrick17k [14]	fitz-roi	1496	Clinical	Melanocytic nevus, melanoma, basal cell carcinoma, benign keratosis, actinic keratosis, squamous cell carcinoma, dermatofibroma
Derm7pt-Derm [17]	d7pt-d	803	Dermoscopic	Melanocytic nevus, melanoma, basal cell carcinoma, benign keratosis, dermatofibroma, vascular lesion
Derm7pt-Clinic [17]	d7pt-c	803	Clinical	Melanocytic nevus, melanoma, basal cell carcinoma, benign keratosis, dermatofibroma, vascular lesion

Table 1. Datasets description.

	isic2018	isic2020	pad	fitz-roi	d7pt-d/c
NEV	6705 (69)	5193 (86)	244 (11)	140 (9)	475 (59)
MEL	1113 (12)	584 (10)	52 (2)	299 (20)	191 (24)
BCC	514 (5)	-	845 (37)	416 (28)	35 (4)
BKL	1099 (11)	223 (4)	235(10)	52(4)	67(8)
AK	-	-	730 (32)	104 (7)	-
SCC	-	-	192 (8)	408 (27)	-
DF	115 (1)	-	-	77 (5)	13 (2)
VL	142 (2)	-	-	-	22 (3)

Table 2. Class distribution for each dataset. The numbers enclosed in parentheses represent percentages.

We adhere to the official data split for Derm7pt, and adopt a 0.2 test split ratio with class stratification for the remaining datasets. For the Fitzpatrick17k subset considered, we remove duplicate or non-lesion data and apply region-of-interest (ROI) preprocessing to the images, using ROI-cropped images in subsequent experiments. Details about the ROI-cropped images can be found in Appendix Sec. 7. Since our objective is to assess the feasibility of UDA in skin lesion analysis, we resize each image to a 64 × 64 pixel format to facilitate the experiments. Throughout the remainder of the paper, we will refer to these datasets by their abbreviated names as listed in Table 1.

3.2. UDA Methods

In this study, we investigate single-source, combined-source, and multi-source UDA training strategies, which are illustrated in Fig. 1. Prior research using adversarial training-based UDA achieved promising results in classifying melanoma and nevus on single-source dermoscopic images [3]. Inspired by this, we selected DANN and ADDA, two representative adversarial-based UDA methods from that study, to investigate combined-source adaptation. Since DANN and ADDN have demonstrated comparable performance in single-source adaptation, we only test single-source UDA with DANN.

Existing literature hasn’t explored the application of

Source		isic2018	isic2020	pad	fitz-roi	d7pt-d	d7pt-c	average
single	isic2018	-	86.8±1.8	60.0±1.5	81.5±1.7	80.3±1.6	70.8±2.3	75.9±2.0
	isic2020	86.6±1.7	-	78.4±1.6	67.7±1.9	60.4±2.1	58.5±2.2	70.3±1.8
	pad	67.1±2.0	61.2±1.8	-	65.6±1.6	77.5±2.1	56.9±2.3	65.7±1.7
	fitz-roi	71.1±2.2	66.5±2.0	58.8±1.7	-	72.4±1.8	70.8±2.1	67.9±2.0
	d7pt-d	78.3±1.9	72.6±1.7	55.1±2.3	72.7±1.8	-	69.0±2.2	69.5±1.8
	d7pt-c	71.9±2.1	67.3±1.9	74.1±2.2	62.1±2.0	78.5±2.0	-	70.8±2.1
	average	75.0±2.0	70.9±1.8	65.3±1.9	69.9±1.8	73.8±1.9	65.2±2.2	-
single-DANN	isic2018	-	86.8±2.3	72.2±3.4	85.5±2.8	81.8±3.1	73.1±2.7	79.9±3.0
	isic2020	79.9±2.6	-	82.6±3.1	84.2±3.3	74.8±2.9	64.4±2.5	77.2±2.8
	pad	63.5±3.1	64.2±2.8	-	83.2±2.5	71.5±2.7	62.4±3.3	69.0±2.9
	fitz-roi	70.7±3.0	68.6±3.2	72.5±3.3	-	78.0±2.9	73.2±3.1	72.6±3.1
	d7pt-d	71.7±2.9	73.0±3.2	68.6±2.8	84.8±2.7	-	72.5±3.0	74.1±2.8
	d7pt-c	75.4±3.2	64.2±3.1	73.9±2.6	83.7±3.1	81.5±3.3	-	75.7±3.0
	average	72.3±3.1	71.4±3.0	74.0±3.2	84.2±2.9	77.5±2.8	69.1±3.3	-
combined	ERM	85.9±1.7	87.1±2.2	74.0±1.9	76.5±2.3	76.7±2.1	74.0±2.4	79.0±2.0
	rs-ERM	86.3±2.1	87.8±2.0	74.9±2.3	84.0±2.1	80.5±2.4	74.6±2.2	81.4±2.2
	DANN	86.9±2.3	85.4±1.9	70.0±2.1	87.7±2.2	82.7±2.0	74.9±2.3	81.3±2.1
	rs-DANN	86.4±2.1	87.3±2.3	76.9±2.2	92.2±2.3	85.5±2.2	74.3±2.4	83.8±2.2
	ADDA	84.4±1.9	84.2±2.0	73.9±2.2	83.3±1.9	78.0±2.1	72.8±2.3	79.4±2.1
	rs-ADDA	86.1±2.3	85.2±2.1	79.8±2.4	79.2±2.0	82.7±2.2	71.2±2.1	80.7±2.2
multi	MDAN	84.3±3.6	85.7±3.5	70.8±3.2	88.6±3.3	81.8±3.5	74.8±3.1	81.0±3.4
	rs-MDAN	86.4±3.4	85.6±3.9	76.3±3.6	84.4±3.8	82.1±3.8	75.1±3.4	81.7±3.6
	M ³ SDA	85.7±3.3	85.7±3.6	86.4±3.5	77.5±3.4	74.5±3.7	71.6±3.6	80.2±3.6
	rs-M ³ SDA	86.3±3.5	89.4±3.8	84.9±3.4	83.4±3.8	85.1±3.6	78.0±3.7	84.5±3.7

Table 3. AUROC for melanoma-nevus binary classification, with each column showing results for a target domain. Single-source training treats each non-target dataset as a source. Combined-source aggregates non-target datasets, applying single-source UDA thereafter. Multi-source results are derived without intentional dataset aggregation, instead relying on UDA methods designed to adapt multiple sources simultaneously. "rs-" denotes the inclusion of a weighted random sampler.

multi-source UDA to the area of skin lesion analysis. Therefore, we selected two methods representing moment-matching and adversarial-based UDA approaches for multi-source training, M³SDA and MSDN, respectively. It's important to note that a comprehensive comparison of various UDA methods falls outside the scope of this work. Our focus here is on evaluating the effectiveness of different UDA training schemes (i.e., single-source, combined-source, and multi-source UDA) in integrating available public data resources and bridging the gap between their domains, for skin lesion classification.

3.3. Evaluating Fairness

Adapting the methods from [7, 8], fairness is quantified using three metrics: Predictive Quality Disparity (**PQD**), Demographic Disparity Metric (**DPM**), and Equality of Opportunity Metric (**EOM**). Let S be the sensitive group, for example, $S = \{\text{male, female}\}$ and denote the label set by $\{1, \dots, M\}$, we can formulate the three metrics as follow:

1. Predictive Quality Disparity (**PQD**) measures the predic-

tion quality difference between each sensitive groups:

$$\text{PQD} = \frac{\min(\text{acc}_j, j \in S)}{\max(\text{acc}_j, j \in S)}$$

here acc_j is the accuracy (or AUROC) of the data in j -th sensitive group.

2. Demographic Disparity Metric (**DPM**) computes the percentage diversities of positive outcomes for each sensitive group.

$$\text{DPM} = \frac{1}{M} \sum_{i=1}^M \frac{\min[p(\hat{y} = i | s = j), j \in S]}{\max[p(\hat{y} = i | s = j), j \in S]}$$

where \hat{y} is the prediction of the model.

3. Equality of Opportunity Metric (**EOM**) requires that different sensitive group should have similar true positive rates.

$$\text{EOM} = \frac{1}{M} \sum_{i=1}^M \frac{\min[p(\hat{y} = i | y = i, s = j), j \in S]}{\max[p(\hat{y} = i | y = i, s = j), j \in S]}$$

where y is the true label.

Source		isic2018	isic2020	pad	fitz-roi	d7pt-d	d7pt-c	average
single	isic2018	-	86.5±3.8	23.9±3.1	21.3±2.5	66.9±3.9	54.2±2.7	50.6±3.6
	isic2020	70.4±4.2	-	10.4±3.4	21.0±4.1	47.6±2.8	46.1±4.4	39.1±3.9
	pad	61.1±2.9	63.7±3.6	-	30.0±2.8	45.3±4.3	44.5±3.2	48.9±3.8
	fitz-roi	56.9±4.3	60.2±3.7	38.3±3.5	-	48.6±3.1	46.6±4.2	50.1±4.4
	d7pt-d	64.9±4.0	69.1±2.9	20.4±4.4	21.0±3.6	-	54.2±2.5	45.9±3.9
	d7pt-c	67.9±4.5	78.5±4.1	13.3±3.8	23.3±4.3	60.8±2.9	-	48.8±3.4
	average	64.2±3.2	71.6±4.3	21.3±3.9	23.3±3.6	53.8±4.0	49.1±3.5	47.2±4.1
combined	ERM	69.7±2.7	86.8±2.8	35.2±2.2	33.3±2.6	64.4±2.1	58.0±2.3	57.9±2.9
	rs-ERM	72.6±3.8	88.7±4.0	43.3±3.7	36.7±3.4	69.7±4.0	57.8±4.3	61.4±3.8
	DANN	69.5±3.6	86.1±4.0	33.9±3.1	33.0±3.9	61.3±3.5	56.0±3.2	56.6±3.6
	rs-DANN	72.8±3.5	88.8±3.9	45.9±4.3	37.7±3.8	64.1±3.3	61.8±3.6	61.8±4.0
	ADDA	68.7±4.2	86.6±3.9	35.7±4.3	19.3±3.7	65.7±4.0	58.0±3.9	55.7±4.1
	rs-ADDA	74.0±3.6	86.9±4.3	40.7±3.2	38.7±3.8	67.9±3.5	63.4±3.9	61.9±3.6
multi	MDAN	66.7±3.8	83.6±4.1	38.7±4.2	36.3±3.9	60.8±3.7	53.7±4.3	56.6±3.8
	rs-MDAN	70.2±4.0	86.7±4.3	38.3±3.6	40.7±3.9	65.7±4.1	60.0±3.8	60.3±4.2
	M ³ SDA	67.4±3.6	81.2±4.3	38.0±3.8	38.7±4.0	63.8±4.1	59.1±3.9	58.0±3.7
	rs-M ³ SDA	71.3±3.9	86.6±3.7	40.7±3.5	40.0±4.1	65.9±4.3	62.3±4.0	61.1±3.9

Table 4. Accuracy of 8-label classification. The eight skin conditions are listed in Table 2.

3.4. Experiment Design

In our experimental design, we treat each dataset as a target domain in turn to assess UDA for binary and multi-class classification. For each target domain considered, we employ single-source, combined-source, and multi-source training schemes. In binary melanoma-nevus classification, we first train classifiers using both single and combined sources for each target domain without deploying UDA, aiming to compare the performance of each single source with an aggregated source. Next, we implement UDA for both single-source and combined-source cases, utilizing the adversarial-based DANN and ADDA methods, which have been proven effective in prior studies for binary skin lesion classification [3, 10]. Lastly, we explore the efficacy of multi-source UDA in binary classification by applying two representative methods, M³SDA and MDAN, that are designed for multi-domain adaptation. To address class imbalance, a weighted random sampler is incorporated into both combined-source and multi-source training, assuming a uniform target label distribution, same as [3]. This technique ensures balanced class ratios within each training mini-batch. Performance is evaluated using the AUROC metric for binary classification.

For multi-class classification, we apply the same training strategies but exclude single-source UDA due to the possible absence of target classes from the source domain. Therefore, we focus on combined-source and multi-source UDA training, whose merits shine when single-source training underperforms due to label space mismatch. For this task, we use accuracy as the evaluation metric. This discrep-

ancy in metrics is due to the impracticality of computing AUROC when classes are missing from either the source or target domain (i.e., label space mismatch).

Fairness assessment is conducted on three datasets—Fitzpatrick17k, ISIC2020, and PAD-UFES-20—for binary classification. The training and test splits from these datasets are combined into one dataset for evaluation to ensure sufficient representation of each group. In alignment with [6, 14], we consider skin color in Fitzpatrick17k as a sensitive attribute, categorizing the data into three groups according to their Fitzpatrick Skin Type (FST): light skin tone (FST 1-2), medium skin tone (FST 3-4), and dark skin tone (FST 5-6). For ISIC2020 and PAD-UFES-20, age is regarded as a sensitive attribute and can be used to form two groups (≤ 30 and >30). The distribution of these sensitive attributes is detailed in Appendix Sec. 8. We evaluate fairness using three common metrics (PQD, DPM, and EOM) in our analysis across five training strategies: single-source and combined-source training, both with and without UDA, in addition to multi-source UDA. It’s important to note that in the single-source scenarios, we use the average fairness results from training on the other 5 datasets independently as the results.

3.5. Implementation Details

All of our experiments utilized a pre-trained VGG16-BN[28] in PyTorch as a feature extractor. Our methodology for single- and combined-source UDA experiments is based on a well-established repository¹. The default hyperparameters were employed for both single and combined sources,

¹<https://github.com/thuml/Transfer-Learning-Library>

	Wasserstein Distance in Input Space					Chi-Square Divergence in Label Space					Test Error				
ISIC2018															
ISIC2020	52.79					3.36					0.30				
PAD-UFES-20	68.45	62.80				4.59	6.42				0.39	0.36			
Fitzpatrick	92.96	77.46	68.14			4.76	7.02	1.37			0.43	0.40	0.62		
Derm7pt-D	80.34	88.06	97.22	116.37		0.10	0.32	110.89	81.00		0.35	0.31	0.80	0.79	
Derm7pt-C	77.69	84.95	83.48	107.48	80.33	0.10	0.32	110.89	81.00	0.00	0.32	0.21	0.87	0.77	0.39

Figure 2. Feature distance (left), label distance (middle), and single-source test error for multi-class classification on the target test set (right). Pearson correlation coefficient is 0.31 between feature distance and test error and 0.78 between label distance and test error.




Figure 3. t-SNE figures of binary classification. Here target domain is fitz-roi and source domains are the other 5. Diagrams 1-2: combined-source without DA. Diagram 3-4: combined-source DANN. Diagrams are colored by domains (1, 3) and by labels (2, 4). After DANN, domains are well aligned and classes remain separable.

with the exception of the learning rates for single-source training, which were slightly decreased to 1×10^{-4} . As for the multi-source UDA, we utilized the official code with the default hyperparameters^{2,3}. Each experiment was repeated for three times. The data loading process was specifically adjusted to ensure the compatibility of these codes with our datasets. We deploy models on an RTX A6000 and an RTX 3090.

4. Results

4.1. Single Source vs. Multiple Sources

Table 3 presents the results for binary classification, with each column representing results for a dataset as the target domain. An immediate observation is that leveraging multiple datasets yields the best results across all target domains. As highlighted in the table, the best results are either derived from combined-source or multi-source training schemes, even when data from the target domain isn't available at all, validating the effectiveness of using mul-

²<https://github.com/hanzhaoml/MDAN>

³<https://github.com/VisionLearningGroup/M3SDA>

iple public datasets. The most prominent improvements are observed in cases where non-ISIC competition datasets are the target domain. Both ISIC2018 and ISIC2020 contain dermoscopic images and share similar class distribution that is highly skewed towards nevus (NEV) in binary classification. It is not surprising to observe that when each of them serves as the other's source domain in single-source training, the results are sufficiently good due to their well aligned feature and label spaces. In contrast, the other four non-ISIC competition datasets benefit more from multiple sources, since they have distinct data distributions that can not be accommodated by a single dataset for training a reliable classifier. Considering single-source training alone, the classifier trained on the ISIC2018 dataset performs best on average across target domains, whereas using the PAD-UFES-20 dataset as a source yields the lowest average results. The performance gap may be attributed to the differences in dataset information: ISIC2018 covers 6 of 8 classes with the most examples, whereas PAD-UFES-20 has the fewest example. A similar pattern emerges in the more challenging multi-class task, as shown in Table 4. Since ISIC2018 and ISIC2020 share a similar distribution




Figure 4. Fairness results on the Fitzpatrick17k, ISIC2020, and PAD-UFES-20 datasets for binary classification. Skin type is considered a sensitive attribute for Fitzpatrick17k, while age for ISIC2020 and PAD-UFES-20. All UDA-based methods in this experiment incorporate a weighted random sampler. Results can be found in Appendix Tab. 5

that skewed towards nevus (NEV) and melanoma (MEL), they can sufficiently serve as each other’s training source and gain less advantage from multiple sources compared to other datasets.

To better understand these results, we use Wasserstein Distance [20, 24, 29] to measure the pixel-level difference and Chi-square divergence to measure label shift, for each single source and target pair (Fig.2), in multi-class classifi-

cation. Calculating their Pearson Correlations with the test error, we find that label shift is highly correlated with test error, yielding a Pearson Correlation coefficient as high as 0.78. The significant role of label shift is intuitive and anticipated. For instance, a classifier trained on ISIC 2020 performs uniformly poorly on other target domains, with the exception of ISIC 2018, which has a small feature and label distance from ISIC 2020. This result further confirms that the utilization of multiple sources can effectively mitigate label mismatch, especially when the distribution of the target domain is unknown.

4.2. With vs. Without Domain Adaptation

In binary classification, the single-source UDA method consistently outperforms its non-DA counterpart across all target domains in average AUROC. The only exception is when ISIC2018 serves as the target domain, where we observe a slight drop in average performance (see the last rows of the single and single-DANN sections in Table 3). Performance improves further when UDA is applied to multiple sources, as highlighted in the previous section. Upon examining the results for combined-source and multi-source training more closely, we find that both strategies do not significantly outperform the ERM baseline (non-DA) on ISIC2018 and ISIC2020. As previously analyzed, this outcome is intuitive because these two ISIC datasets are each other's sufficient training sets due to their similar feature and label distributions, with or without domain adaptation. Thus, including other smaller, non-ISIC datasets in the training set adds less informative value to the dominant ISIC source. In contrast, applying UDA on multiple sources consistently brings sizable improvements over the ERM baselines when non-ISIC datasets are the target domains. This indicates that useful information from public datasets has been effectively transferred to the target domain as a result of UDA. Taking Fitzpatrick17k as an example target domain, Fig. 3 demonstrates that applying DANN on a combined source forces the domain representations to integrate while preserving the separability of classes. For eight-class classification, despite the significantly increased task difficulty compared to binary classification, training on multiple sources with UDA generally outperforms non-DA training (Table 4).

4.3. Impact on Fairness

Fairness is tested on three datasets - Fitzpatrick17k, ISIC 2020, and PAD-UFES-20 - for the binary classification task, using three metrics: PQD, DPM, and EOM. According to Table 3, DANN and M³SDA with imbalance technique deliver the best average results in combined-source and multi-source training, respectively. Therefore, these two methods are selected to examine the impact of combined-source and multi-source training on fairness. As illustrated in Fig. 4,

applying UDA on multiple sources consistently achieves superior classification performance while maintaining a minimal performance gap between groups. Despite the absence of fairness-focused techniques specifically aimed at minority groups, the results indicate that UDA effectively adapts lesion-related features and translates rich demographic information to the target domain, nudging the model towards a more equitable diagnostic system.

5. Conclusion

In this study, we have explored three UDA training strategies: single-source, combined-source, and multi-source, when labeled target data is unavailable. The results validate the effectiveness of using multiple sources over a single source and establish the superiority of UDA over non-DA training in binary and multi-class classifications. Our results show that leveraging UDA with multiple sources enriches the class information of the training set and ensures integration of diverse domains. The merits of applying UDA on multiple sources particularly stand out when not a single dataset can cover the feature or label spaces of the target domain. Additionally, this study highlights the potential of utilizing UDA with multiple sources to promote fairness and mitigate bias against minority groups. We believe the practicality of our work is evident in the medical field, where custom datasets are usually too small to train a reliable diagnosis tool, and public datasets present a new, cost-free source of information. This work also opens up other interesting questions, such as the root cause of the improved fairness after adopting UDA on multiple sources, which will be left as a new direction for future work.

Acknowledgement

This work was partly supported by the NSF EPSCoR-Louisiana Materials Design Alliance (LAMDA) program #OIA-1946231 and partly by the Harold L. and Heather E. Jurist Center of Excellence for Artificial Intelligence at Tulane University.

References

- [1] Alceu Bissoto, Eduardo Valle, and Sandra Avila. Gan-based data augmentation and anonymization for skin-lesion analysis: A critical review. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1847–1856, 2021. 1
- [2] Titus Brinker, Achim Hekler, Alexander Enk, Joachim Klode, Axel Hauschild, Carola Berking, Bastian Schilling, Sebastian Haferkamp, Jochen Utikal, Christof Kalle, Collaborators, Stefan Fröhling, and Michael Weichenthal. A convolutional neural network trained with dermoscopic images performed on par with 145 dermatologists in a clinical melanoma image classification task. *European Journal of Cancer*, 111:148–154, 2019. 1

- [3] Sireesha Chamarthi, Katharina Fogelberg, Titus J. Brinker, and Julia Niebling. Mitigating the influence of domain shift in skin lesion classification: A benchmark study of unsupervised domain adaptation methods. *Informatics in Medicine Unlocked*, 44:101430, 2024. 1, 3, 5
- [4] Noel CF Codella, David Gutman, M Emre Celebi, Brian Helba, Michael A Marchetti, Stephen W Dusza, Aadi Kalloo, Konstantinos Liopyris, Nabin Mishra, Harald Kittler, et al. 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)*, pages 168–172. IEEE, 2018. 3
- [5] Alberto Coustasse, Raghav Sarkar, Bukola Abodunde, Brandon J. Metzger, and Chelsea M. Slater. Use of teledermatology to improve dermatological access in rural areas. *Telemedicine journal and e-health : the official journal of the American Telemedicine Association*, 2019. 1
- [6] Roxana Daneshjou, Kailas Vodrahalli, Weixin Liang, Roberto A Novoa, Melissa Jenkins, Veronica Rotemberg, Justin Ko, Susan M Swetter, Elizabeth E Bailey, Olivier Gevaert, Pritam Mukherjee, Michelle Phung, Kiana Yekrang, Bradley Fong, Rachna Sahasrabudhe, James Zou, and Albert Chiou. Disparities in dermatology ai performance on a diverse, curated clinical image set. *Science Advances*, 2022. 3, 5
- [7] Mengnan Du, Fan Yang, Na Zou, and Xia Hu. Fairness in deep learning: A computational perspective, 2020. 4
- [8] Siyi Du, Ben Hers, Nourhan Bayasi, Ghassan Hamarneh, and Rafeef Garbi. Fairdisco: Fairer ai in dermatology via disentanglement contrastive learning, 2022. 4
- [9] Andre Esteva, Brett Kuperl, Roberto A. Novoa, Justin M. Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542:115–118, 2017. 1
- [10] Katharina Fogelberg, Sireesha Chamarthi, Roman C Maron, Julia Niebling, and Titus J Brinker. Domain shifts in dermoscopic skin cancer datasets: Evaluation of essential limitations for clinical translation. *New Biotechnology*, 76:106–117, 2023. 1, 3, 5
- [11] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. 2
- [12] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 1180–1189, Lille, France, 2015. PMLR. 1
- [13] Amirata Ghorbani, Vivek Natarajan, David Coz, and Yuan Liu. DermGAN: Synthetic Generation of Clinical Skin Images with Pathology. In *Proceedings of the Machine Learning for Health NeurIPS Workshop*, pages 155–170. PMLR, 2020. 1
- [14] Matthew Groh, Caleb Harris, Luis Soenksen, Felix Lau, Rachel Han, Aerin Kim, Arash Koochek, and Omar Badri. 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*, pages 1820–1828, 2021. 3, 5
- [15] Yanyang Gu, Zongyuan Ge, C Paul Bonnington, and Jun Zhou. Progressive transfer learning and adversarial domain adaptation for cross-domain skin disease classification. *IEEE journal of biomedical and health informatics*, 24(5):1379–1393, 2019. 3
- [16] Glenn Jocher, Ayush Chaurasia, and Jing Qiu. YOLO by Ultralytics, 2023. 1
- [17] Jeremy Kawahara, Sara Daneshvar, Giuseppe Argenziano, and Ghassan Hamarneh. Seven-point checklist and skin lesion classification using multitask multimodal neural nets. *IEEE journal of biomedical and health informatics*, 23(2):538–546, 2018. 3
- [18] Agostina J. Larrazabal, Nicolás Nieto, Victoria Peterson, Diego H. Milone, and Enzo Ferrante. Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. *Proceedings of the National Academy of Sciences of the United States of America*, 117:12592 – 12594, 2020. 3
- [19] Yuan Liu, Ayush Jain, Clara Eng, David H Way, Kang Lee, Peggy Bui, Kimberly Kanada, Guilherme de Oliveira Marinho, Jessica Gallegos, Sara Gabriele, Vishakha Gupta, Nalini Singh, Vivek Natarajan, Rainer Hofmann-Wellenhof, Greg S Corrado, Lily H Peng, Dale R Webster, Dennis Ai, Susan J Huang, Yun Liu, R Carter Dunn, and David Coz. A deep learning system for differential diagnosis of skin diseases. *Nature Medicine*, 26(6):900–908, 2020. 1
- [20] Akshay Mehra, Yunbei Zhang, and Jihun Hamm. Analysis of task transferability in large pre-trained classifiers. *arXiv preprint arXiv:2307.00823*, 2023. 7
- [21] Yaniv Ovadia, Emily Fertig, Jie Jessie Ren, Zachary Nado, D. Sculley, Sebastian Nowozin, Joshua V. Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. In *Neural Information Processing Systems*, 2019. 1
- [22] Andre G C Pacheco, Gustavo R Lima, Amanda S Salomão, Breno Krohling, Igor P Biral, Gabriel G de Angelo, Fábio C R Alves, Jr, José G M Esgario, Alana C Simora, Pedro B C Castro, Felipe B Rodrigues, Patricia H L Frasson, Renato A Krohling, Helder Knidel, Maria C S Santos, Rachel B do Espírito Santo, Telma L S G Macedo, Tania R P Canuto, and Luíz F S de Barros. PAD-UFES-20: A skin lesion dataset composed of patient data and clinical images collected from smartphones. *Data in Brief*, 32:106221, 2020. 3
- [23] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1406–1415, 2019. 3
- [24] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019. 7
- [25] Zhihang Ren, Yunhui Guo, X Yu Stella, and David Whitney. Improve image-based skin cancer diagnosis with generative

- self-supervised learning. In *2021 IEEE/ACM Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE)*, pages 23–34. IEEE, 2021. 1
- [26] Veronica Rotemberg, Nicholas Kurtansky, Brigid Betz-Stablein, Liam Caffery, Emmanouil Chousakos, Noel Codella, Marc Combalia, Stephen Dusza, Pascale Guitera, David Gutman, Allan Halpern, Brian Helba, Harald Kittler, Kivanc Kose, Steve Langer, Konstantinos Lioprys, Josep Malvehy, Shenara Musthaq, Jabpani Nanda, Ofer Reiter, George Shih, Alexander Stratigos, Philipp Tschandl, Jochen Weber, and H Peter Soyer. A patient-centric dataset of images and metadata for identifying melanomas using clinical context. *Scientific Data*, 8(1):34, 2021. 3
- [27] Laleh Seyyed-Kalantari, Haoran Zhang, Matthew BA McDermott, Irene Y Chen, and Marzyeh Ghassemi. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nature medicine*, 27(12):2176–2182, 2021. 3
- [28] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 5
- [29] Yang Tan, Yang Li, and Shao-Lun Huang. Otce: A transferability metric for cross-domain cross-task representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15779–15788, 2021. 7
- [30] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 5(1):1–9, 2018. 3
- [31] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7167–7176, 2017. 2
- [32] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? *ArXiv*, abs/1411.1792, 2014. 1
- [33] Han Zhao, Shanghang Zhang, Guanhang Wu, José MF Moura, Joao P Costeira, and Geoffrey J Gordon. Adversarial multiple source domain adaptation. *Advances in neural information processing systems*, 31, 2018. 3