

BiasPruner: Debaised Continual Learning for Medical Image Classification

Nourhan Bayasi¹[0000-0003-4653-6081], Jamil Fayyad²[0000-0003-1553-8754],
Alceu Bissoto³[0000-0003-2293-6160], Ghassan Hamarneh⁴[0000-0001-5040-7448],
and Rafeef Garbi¹[0000-0001-6224-0876]

¹ University of British Columbia, Vancouver, BC, Canada

² University of Victoria, Victoria, BC, Canada

³ University of Campinas, Campinas, Brazil

⁴ Simon Fraser University, Burnaby, BC, Canada
nourhanb@ece.ubc.ca

Abstract. Continual Learning (CL) is crucial for enabling networks to dynamically adapt as they learn new tasks sequentially, accommodating new data and classes without catastrophic forgetting. Diverging from conventional perspectives on CL, our paper introduces a new perspective wherein forgetting could actually benefit the sequential learning paradigm. Specifically, we present **BiasPruner**, a CL framework that intentionally forgets spurious correlations in the training data that could lead to shortcut learning. Utilizing a new bias score that measures the contribution of each unit in the network to learning spurious features, **BiasPruner** prunes those units with the highest bias scores to form a debaised subnetwork preserved for a given task. As **BiasPruner** learns a new task, it constructs a new debaised subnetwork, potentially incorporating units from previous subnetworks, which improves adaptation and performance on the new task. During inference, **BiasPruner** employs a simple task-agnostic approach to select the best debaised subnetwork for predictions. We conduct experiments on three medical datasets for skin lesion classification and chest X-Ray classification and demonstrate that **BiasPruner** consistently outperforms SOTA CL methods in terms of classification performance and fairness. Our code is available here.

Keywords: Continual Learning · Debias · Pruning · Shortcut Learning.

1 Introduction

Humans inherently learn in a continual manner, acquiring new concepts over time without forgetting previous ones. In contrast, deep learning models encounter the challenge of catastrophic forgetting [17], wherein learning new data can override previously acquired knowledge. This issue becomes especially pronounced within the medical domain, given the ever-evolving nature of medical data, the variations in acquisition protocols, the utilization of diverse devices for obtaining medical images, and other factors that contribute to shifts in data distributions or the introduction of new disease classes over time. As a result,

continual learning (CL) [25,27] has emerged as a promising solution, allowing a network to learn continually over a sequence of presented data while forgetting as little as possible about previous knowledge. Several CL methods have emerged within the medical field to address the challenge of forgetting. Replay-based methods [22,13] store a subset of data samples and replay them to retain old information, regularization-based methods [15] impose restrictions on the network parameter updates to preserve prior knowledge while learning new tasks, and architecture-based methods, assign specialized architectural components for each task within the network [1,2] or expand them to accommodate new tasks [9].

While previous CL methods achieved success, they have yet to consider a more realistic setting in which dataset bias exists. In medical imaging, bias could manifest through an imbalanced distribution of sensitive attributes (e.g., gender, age, ethnicity) [4]. Even slight imbalances induce spurious correlations between attributes and the classification target (diagnosis) [3], creating an illusion of predictive power that models can exploit. Leveraging such information compromises the network’s generalization ability, amplifying societal biases over misrepresented populations in data (e.g., detecting melanoma in individuals with dark skin tones). In CL, learning spurious correlations poses a significant challenge due to bias transfer, where biases learned by a model can be transferred to a downstream task even if it has unbiased data [23]. Since CL involves learning a sequence of tasks, the bias transfer can potentially be amplified. Moreover, recent work [5] mathematically proved that handling bias becomes substantially harder when tasks are presented sequentially compared to joint training.

To address this gap and tackle bias in CL, we propose **BiasPruner**, a fixed-size network capable of learning sequentially and fairly over time by dedicating a unique debiased subnetwork for each task. **BiasPruner** leverages a newly proposed bias score to measure the contribution of each unit in the network to learning spurious features. Units with high bias scores are pruned to form a task-specific debiased subnetwork, which is kept frozen to avoid forgetting, whereas the remaining pruned units are subsequently offered for learning new tasks. Fig. 1 presents an overview of our method. We evaluate our solution on three medical imaging classification datasets, each with different bias attributes. Our results demonstrate **BiasPruner**’s superior performance in both classification accuracy and fairness. While a few recent methods have addressed fairness in CL [14,16,6], **BiasPruner**, to the best of our knowledge, is the first work in the medical field covering different benchmarks and bias attributes in a class-incremental setup. Crucially, **BiasPruner** does *not* require dataset biases to be explicitly annotated. This is particularly relevant in healthcare, where identifying biases is complex and costly, compounded by patient data privacy concerns [19].

2 Methodology

BiasPruner employs a fixed-size network, f , capable of learning T tasks sequentially, one at a time, where T is not pre-determined, without forgetting any of the previously learned tasks. During training the t -th domain, where




Fig. 1: (Left) **BiasPruner** learns sequentially, allocating a subnetwork for each task. (Right) **BiasPruner** evaluates each network unit’s contribution to learning spurious features from biased training data, assigning bias scores. High-score units are pruned, and the subnetwork is finetuned on both easy and hard samples.

$t \in \{1, 2, \dots, T\}$, the network does not have access to old data, i.e., it exclusively receives biased training data $D_t = (x_i, y_i)$ specific to the current task, where (x_i, y_i) represent the training samples, consisting of a total of N_t images and $y_i \in \mathcal{C}_t$ classes (note the subscript t emphasizing that the set of classes may change, including adding new classes, for new tasks). For clarity, we employ the symbol c to denote any class within the set $\mathcal{C}_t$. **BiasPruner** creates a debiased subnetwork for the t -th task by pruning units in the network that are mostly correlated with unknown bias(es) in D_t . Furthermore, **BiasPruner** transfers knowledge through pruning of the original network, including units of previously created subnetworks, for each new task. At inference, **BiasPruner** identifies the optimal subnetwork for predictions on a given data in a task-agnostic setup; i.e., information about the task origin of a test image is unknown or unavailable.

2.1. Detecting Spurious Features through Bias Scoring. Given a biased dataset D_t , one of the key causes of learning shortcut predictions occurs when the model finds it easier to learn spurious features rather than the intended ones [19]. Consequently, we propose to intentionally encourage the network f to quickly fit on the easier features from the training data of D_t . To achieve this, we adopt the generalized cross entropy (GCE) [29] loss, $\mathcal{L}_{\text{GCE}}$, which was originally proposed to address noisy labels by fitting on the easier clean data and slowly memorizing the hard noisy samples. The GCE loss is formulated as follows:

$$\mathcal{L}_{\text{GCE}}(p(x; \theta), y) = \frac{1 - p_y(x; \theta)^q}{q}$$

where $q \in (0, 1]$ is a hyperparameter controlling the degree of bias amplification, $p(x; \theta)$ and $p_y(x; \theta)$ are the softmax output of the network and its probability assigned to the target label y , respectively. Due to the GCE loss’s gradient, which up-weights samples with a high probability of predicting the correct target, the network quickly becomes biased to easier samples and learns shortcuts [21].

Once the network is biased, it becomes logical to identify the units that have contributed the most to learning the shortcut in each class. To achieve this, we

partition the training data $\{x, y\}$ into two groups per groundtruth class c : The biased sample set, $\mathcal{E}_c^t$, consists of (x_i, y_i) pairs that are correctly classified by the biased network with a probability $p_{y,i} \geq \tau$; i.e., samples that are $\mathcal{E}$ asier for the network to learn. Similarly, the unbiased sample set, $\mathcal{H}_c^t$, comprises (x_i, y_i) pairs that are misclassified by the biased network; i.e., samples that are $\mathcal{H}$ arder to learn, as follows (Refer to supplementary material, Fig. 4, for visualizations):

$$\mathcal{E}_c^t = \{i \mid y_i = c_i \ \& \ p_{y,i} \geq \tau\} \quad , \quad \mathcal{H}_c^t = \{i \mid y_i \neq c_i \ \& \ p_{y,i} \geq \tau\}.$$

Next, we define a bias score, $\mathcal{S}_{c,n}^t$, for each unit n in the biased network relative to a given class c by analyzing each unit’s ReLU activation, a_i^n , as follows:

$$\mathcal{S}_{c,n}^t = \frac{1}{|\mathcal{E}_c^t|} \sum_{i \in \mathcal{E}_c^t} \text{Var}(a_i^n) - \frac{1}{|\mathcal{H}_c^t|} \sum_{i \in \mathcal{H}_c^t} \text{Var}(a_i^n).$$

$\text{Var}(a_i^n)$ represents the variance of the feature map a_i^n over its spatial dimensions (w, h) . The final unit-based bias score $\bar{\mathcal{S}}_n^t$ is calculated by averaging the results over all class-specific scores. Units that respond more strongly to biased samples (c.f. $\mathcal{E}_c^t$) than to unbiased samples (c.f. $\mathcal{H}_c^t$) are assigned higher bias scores, designating them as the main contributors to learning shortcuts in the network.

2.2 Forming Subnetworks by Bias-aware Pruning and Finetuning. To ensure fairness in CL, we form a task-specific, debiased subnetwork, f_t , by selectively removing the units responsible for learning the bias in D_t . The pruning involves removing the top $\gamma\%$ of units, which includes the output feature maps with the highest bias scores and their corresponding filters, leaving $(1 - \gamma\%)$ for each f_t . To counteract potential performance drop post-pruning while prioritizing improved performance on harder-to-learn samples, we propose a new weighted cross entropy loss, $\mathcal{L}_{\text{WCE}}$, for fine-tuning f_t on D_t over a few epochs:

$$\mathcal{L}_{\text{WCE}}(x) = \mathcal{W}(x) \cdot \mathcal{L}_{\text{CE}}(f(x), y), \quad \text{where} \quad \mathcal{W}(x) = \exp(\alpha \cdot \mathcal{L}_{\text{GCE}}(x)).$$

$\alpha \in (0, 1)$ is a trainable parameter, and $\mathcal{L}_{\text{GCE}}(x)$ is the sample’s GCE loss value determined as discussed in Sec. 2.1. With this weighted function, the influence of training samples in the finetuning process varies according to their bias alignment; i.e., easy samples (c.f., $\mathcal{L}_{\text{GCE}}(x)$ is small) are down-weighted, whereas hard samples (c.f., $\mathcal{L}_{\text{GCE}}(x)$ is large) are up-weighted, exponentially.

2.3. Debiased Knowledge Transfer for Enhanced Task Adaptation. When learning a new task, **BiasPruner** facilitates knowledge transfer (KT), which is achieved by pruning the entire original network f to create the new task-specific subnetwork, including both free units and pre-assigned debiased subnetworks of previous tasks. To avoid forgetting the previously acquired knowledge, the subnetworks associated with prior tasks are kept frozen and only the free units are updated to learn the new task.

Table 1: Details on the multi-class disease datasets used in our experiments.

Dataset	Number of images	Classes	Tasks	Classes per task	Dataset bias
FITZ	16,012	114	6	[19, 19, 19, 19, 19, 19]	Skin tone (I, II, III, IV, V, VI)
HAM	8,678	7	3	[2, 2, 3]	Age (age $\geq$ 60, age $<$ 60)
NIH	19,993	14	3	[4, 5, 5]	Gender (male, female)

2.4. Task-agnostic Inference. BiasPruner addresses a practical scenario where the task identity of a test image is unknown during inference. In other words, the specific task to which an image belongs to is not explicitly provided. Given a test batch of size s as $\mathbf{X}^{\text{test}}$, we employ a ‘maxoutput’ strategy for task prediction [7], which involves identifying the task with the maximum output response: $t^* = \arg \max_{t=1,2,\dots,T} \sum_{i=1}^s \max \varphi_t(\theta_t(\mathbf{x}_i^{\text{test}}))$, where φ_t is the fully connected layer of the t -th subnetwork. Subsequently, we use the selected t^* task to make the final prediction $\hat{\mathbf{y}}$ based on the corresponding subnetwork; $\hat{\mathbf{y}} = f_{t^*}(\mathbf{X}^{\text{test}})$.

3 Experiments and Results

Datasets. We selected/constructed datasets based on three primary considerations: (a) the presence of a dataset bias that is spuriously correlated with the disease classes; (b) The need for a variety of classes to facilitate the CL setup; and (c) publicly available to ensure reproducibility. Hence, we include Fitzpatrick17K (FITZ) [10], HAM10000 (HAM) [24] and NIH ChestX-Ray14 (NIH) [26]. Each dataset has 114, 7 and 14 distinct classes, respectively, that are split into 6, 3 and 3 tasks, respectively, with non-overlapping classes, as shown in Table 1 (Refer to supplementary material for dataset (Table 5) & bias (Fig. 5) details).

Evaluation Metrics. We assess the performance of BiasPruner using both the accuracy and fairness metrics. We use the commonly used F1-score (F) and balanced accuracy (ACC) metrics. We report the accuracy per sensitive attribute (e.g., male, female) as well as overall class performance (Overall). For fairness, we use the demographic parity ratio (DPR) and equal opportunity difference (EOD) metrics. Similar to other CL methods, we report all metrics at the end of learning (i.e., after training the model on all T tasks), averaged across all tasks.

Implementation Details. We use ResNet-50 [11] as the backbone for feature extraction and a unified classifier for all tasks during inference. We use the Adam optimizer with a batch of 32 images for 200 epochs to train BiasPruner with $\mathcal{L}_{\text{GCE}}$, having early stopping in case of overfitting. We set q in $\mathcal{L}_{\text{GCE}}$ to 0.7 (default) and the confidence threshold to $\tau = 0.70$. We set the pruning ratio to $\gamma = 0.6$ for all tasks. For the finetuning with $\mathcal{L}_{\text{WCE}}$, we train the debiased subnetwork for 20 epochs, and we saved the weights with the highest ACC and EOD on the validation set. In all experiments, we report averaged results across three random task orders, aiming to neutralize any potential impact of the order in which tasks are processed during network training.

I. Quantitative Results on Skin-tone-biased Dataset (FITZ) are reported in Table 2. First, we compare BiasPruner (Exp $\mathcal{D}$) against three common CL

Table 2: Classification performance and fairness on FITZ. Best results marked in **bold** (except upper-bound). Higher is better for all metrics except EOD.

Exp	Method	F	ACC							DPR	EOD
			Type-I	Type-II	Type-III	Type-IV	Type-V	Type-VI	Overall		
Comparison against Baselines											
$\mathcal{A}$	JOINT	0.256	0.269	0.304	0.335	0.309	0.365	0.245	0.324	0.137	0.298
	SINGLE	0.435	0.410	0.469	0.465	0.495	0.492	0.430	0.472	0.185	0.251
	SeqFT	0.188	0.187	0.261	0.299	0.254	0.214	0.192	0.221	0.051	0.721
Comparison against CL Methods											
$\mathcal{B}$	EWC	0.325	0.254	0.356	0.355	0.401	0.412	0.244	0.324	0.212	0.342
	PackNet	0.433	0.366	0.402	0.445	0.447	0.479	0.319	0.414	0.154	0.425
	SupSup	0.451	0.254	0.298	0.441	0.452	0.436	0.410	0.425	0.162	0.431
Comparison against CL with Bias Mitigation Methods											
$\mathcal{C}$	EWC+S	0.308	0.264	0.357	0.324	0.411	0.417	0.385	0.341	0.228	0.311
	PackNet+S	0.495	0.434	0.485	0.494	0.565	0.562	0.584	0.501	0.184	0.248
	SupSup+S	0.466	0.418	0.467	0.432	0.554	0.561	0.534	0.492	0.182	0.221
	EWC+W	0.321	0.251	0.356	0.334	0.392	0.401	0.398	0.346	0.216	0.298
	PackNet+W	0.527	0.405	0.477	0.480	0.529	0.546	0.524	0.472	0.144	0.246
	SupSup+W	0.457	0.425	0.451	0.448	0.530	0.561	0.544	0.508	0.178	0.254
Our Proposed Fair Continual Learning Method											
$\mathcal{D}$	BiasPruner	0.540	0.457	0.502	0.435	0.551	0.563	0.584	0.512	0.331	0.202
[Upper-bound] Comparison against a Bias Mitigation Method											
$\mathcal{E}$	FairDisCo	0.542	0.479	0.523	0.468	0.571	0.574	0.615	0.548	0.474	0.192

baselines (Exp **A**): JOINT, which consolidates data from all tasks for joint model training; SINGLE, which trains separate models for each task and deploys task-specific models during inference; and SeqFT, which finetunes a single model on the current task without addressing forgetting. We observe that SINGLE outperforms JOINT as each task is learned independently, leading to improved classification and fairness results, and that SeqFT exhibits a significant performance drop due to catastrophic forgetting. Notably, **BiasPruner** outperforms baselines in terms of overall accuracy and fairness, attributing this to its ability to reduce the training data bias and transfer knowledge across the tasks.

Secondly, we compare **BiasPruner** against three CL methods (Exp **B**): EWC [12], a regularization-based method, and PackNet [20] and SupSup [28], both subnetwork-based like ours. We notice that these CL methods demonstrate lower fairness compared to **BiasPruner**, which is expected as they overlook dataset bias. Specifically, PackNet and SupSup exhibit higher accuracy but lower fairness compared to EWC. This is mainly due to their subnetwork-based nature, which can inadvertently worsen accuracy disparities, particularly among specific subgroups, during the removal of unimportant parameters [18].

Thirdly, we enhance the competing CL methods by augmenting each of them with pre-processing bias mitigation algorithms (Exp **C**). Specifically, we apply the Resampling Algorithm (S), which balances the dataset by oversampling minorities and undersampling majorities within each pair of skin label and tone. Additionally, we explore the Reweighting Algorithm (W) [8], which assigns lower weights to images that have been disadvantaged or favored to prevent the model from learning discriminatory features. While showing improved accuracy and fairness compared to Exp **B**, they fall short of our **BiasPruner**'s performance.

Table 3: Classification performance and fairness on HAM and NIH. Best results marked in **bold** (except upper-bound).

Exp	Method	HAM						NIH					
		F	ACC		Overall	DPR	EOD	F	ACC		Overall	DPR	EOD
			<60	≥60					M	F			
Comparison against Baselines													
$\mathcal{F}$	JOINT	0.755	0.781	0.665	0.738	0.239	0.320	0.282	0.306	0.259	0.285	0.706	0.325
	SINGLE	0.836	0.819	0.834	0.841	0.609	0.131	0.434	0.428	0.403	0.417	0.728	0.311
	SeqFT	0.431	0.372	0.404	0.416	0.201	0.558	0.219	0.251	0.217	0.231	0.246	0.544
Comparison against CL Methods													
$\mathcal{G}$	EWC	0.788	0.773	0.804	0.772	0.561	0.360	0.398	0.428	0.405	0.417	0.562	0.264
	PackNet	0.824	0.807	0.799	0.808	0.620	0.302	0.434	0.47	0.444	0.458	0.588	0.284
	SupSup	0.831	0.788	0.845	0.822	0.625	0.296	0.448	0.451	0.441	0.445	0.571	0.293
Comparison against CL with Bias Mitigation Methods													
$\mathcal{H}$	EWC+S	0.834	0.821	0.832	0.827	0.575	0.172	0.412	0.434	0.416	0.421	0.567	0.259
	PackNet+S	0.839	0.849	0.817	0.829	0.613	0.181	0.419	0.44	0.425	0.434	0.640	0.211
	SupSup+S	0.849	0.802	0.811	0.817	0.639	0.204	0.432	0.456	0.448	0.451	0.662	0.204
	EWC+W	0.791	0.778	0.784	0.781	0.544	0.168	0.418	0.441	0.423	0.432	0.569	0.251
	PackNet+W	0.814	0.877	0.819	0.842	0.549	0.189	0.443	0.462	0.456	0.459	0.704	0.192
	SupSup+W	0.846	0.797	0.809	0.803	0.536	0.213	0.458	0.481	0.463	0.474	0.731	0.184
Our Proposed Fair Continual Learning Method													
$\mathcal{I}$	BiasPruner	0.860	0.851	0.852	0.858	0.642	0.127	0.488	0.525	0.484	0.507	0.821	0.188
[Upper-bound] Comparison against a Bias Mitigation Method													
$\mathcal{J}$	FairDisCo	0.873	0.876	0.904	0.893	0.682	0.113	0.486	0.545	0.512	0.538	0.855	0.150

Finally, we compare **BiasPruner** to FairDisCo [8], a (non-CL) bias mitigation technique for medical applications, which uses bias annotations in training. Therefore, it can set an upper bound on the performance. For a fair comparison, we allow FairDisCo to learn each task independently and report the average performance over all tasks (Exp $\mathcal{E}$). Despite not using bias annotations, **BiasPruner** exhibits slightly lower but comparable performance to FairDisCo.

II. Quantitative Results on Age- and Gender-biased Dataset (HAM & NIH, respectively) are given in Table 3. **BiasPruner** (Exp $\mathcal{I}$) outperforms other baselines (Exp $\mathcal{F}$), CL methods (Exp $\mathcal{G}$), and CL methods with debiasing (Exp $\mathcal{H}$) in both overall task classification and fairness.


III. Ablation Studies analyze the impact of individual components in **BiasPruner** (Table 4). In Exp $\mathcal{K}$, we train the model using CE loss instead of GCE. In Exp $\mathcal{L}$, we randomly prune the network instead of using our bias-based pruning. In Exp $\mathcal{M}$, we finetune the debiased subnetworks in **BiasPruner** with CE loss without weighting it. In Exp $\mathcal{N}$, we simulate the absence of knowledge transfer (KT) by prohibiting any overlapping between the parameters θ_t and $\theta_{t'}$ for any two tasks t and t' . We observe that 1) the impact of $\mathcal{L}_{\text{WCE}}$ (Exp $\mathcal{M}$) is predominant, as fine-tuning with CE leads to the poorest performance in accuracy and fairness, attributed to the risk of subnetworks potentially relearning bias.

IV. A Sequential Analysis, as illustrated in Fig. 2, showcases the consistently superior performance of **BiasPruner** over other methods in terms of overall accuracy and DPR after each step in the continual learning sequence across FITZ, HAM, and NIH datasets.


V. For Analysis of Model Biases (e.g., skin tone), we trained classifiers for sensitive attribute detection on top of frozen feature extractors from three net-

Table 4: Classification (Overall) & fairness (DPR) results of **BiasPruner** from ablation studies. Best results marked in **bold**.

Exp	$\mathcal{L}_{\text{GCE}}$	Bias-aware Pruning	$\mathcal{L}_{\text{WCE}}$	KT	FITZ		HAM		NIH	
					Overall $\uparrow$	DPR $\uparrow$	Overall $\uparrow$	DPR $\uparrow$	Overall $\uparrow$	DPR $\uparrow$
$\mathcal{D}, \mathcal{I}$	✓	✓	✓	✓	0.512	0.331	0.858	0.642	0.507	0.821
$\mathcal{K}$	×	✓	✓	✓	0.498	0.254	0.834	0.579	0.501	0.779
$\mathcal{L}$	✓	×	✓	✓	0.508	0.328	0.842	0.637	0.498	0.814
$\mathcal{M}$	✓	✓	×	✓	0.481	0.247	0.792	0.576	0.468	0.754
$\mathcal{N}$	✓	✓	✓	×	0.504	0.324	0.851	0.630	0.496	0.803

Fig. 2: The overall (dashed) and DPR (dotted) performance of **BiasPruner** and other methods over *all the seen tasks* after each training step in the continual learning sequence, where T_i refers to the i th task.

works: CE-based, GCE-based, and our **BiasPruner**, all pre-trained to diagnose (Fig. 3). The better-than-chance ($\in [0.63, 0.828]$) detection accuracy of sensitive attributes, in CE- and GCE-based training, reveals that sensitive data is embedded in the originally learned features, i.e., the presence of bias. GCE, due to its loss function promoting shortcut learning, showed the most bias. The high accuracy achieved by CE shows that even naïvely trained models are susceptible to bias. In contrast, **BiasPruner** shows minimal bias, reflected by its near-random AUC values ($\in [0.49, 0.67]$) when detecting sensitive information.

Fig. 3: Sensitive attribute detection from frozen models pre-trained to diagnose. **BiasPruner** low AUCs indicate that bias is not encoded in its resulting features.

4 Conclusion

In this paper, we presented **BiasPruner**, a new continual learning (CL) framework that leverages intentional forgetting to improve fairness and mitigate catas-

trophic forgetting. By quantifying the influence of each network unit on learning spurious features, **BiasPruner** identifies and prunes highly biased units to construct a debiased subnetwork for each task. Experimental evaluations on three classification datasets demonstrate that **BiasPruner** consistently outperforms baselines and CL methods in classification performance and fairness. Our results highlight the need to address dataset bias in future CL designs and evaluations, due to the frequent fairness shortcomings of current CL methods.

Acknowledgments. We thank NVIDIA for their hardware grant and the Natural Sciences and Engineering Research Council (NSERC) of Canada for the Vanier PhD Fellowship. A. Bissoto is funded by FAPESP (2019/19619-7, 2022/09606-8).

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

References

1. Bayasi, N., Du, S., Hamarneh, G., Garbi, R.: Continual-GEN: Continual group ensembling for domain-agnostic skin lesion classification. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 3–13. Springer (2023)
2. Bayasi, N., Hamarneh, G., Garbi, R.: Culprit-Prune-Net: Efficient continual sequential multi-domain learning with application to skin lesion classification. In: Medical Image Computing and Computer Assisted Intervention (MICCAI). pp. 165–175. Springer (2021)
3. Bissoto, A., Barata, C., Valle, E., Avila, S.: Even small correlation and diversity shifts pose dataset-bias issues. *Pattern Recognition Letters* (2024)
4. Brown, A., Tomasev, N., Freyberg, J., Liu, Y., Karthikesalingam, A., Schrouff, J.: Detecting shortcut learning for fair medical ai using shortcut testing. *Nature Communications* **14**(1), 4314 (2023)
5. Busch, F.P., Kamath, R., Mitchell, R., Stammer, W., Kersting, K., Mundt, M.: Where is the truth? the risk of getting confounded in a continual world. arXiv preprint arXiv:2402.06434 (2024)
6. Chowdhury, S.B.R., Chaturvedi, S.: Sustaining fairness via incremental learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 6797–6805 (2023)
7. Dekhovich, A., Tax, D.M., Sluiter, M.H., Bessa, M.A.: Continual prune-and-select: class-incremental learning with specialized subnetworks. *Applied Intelligence* **53**(14), 17849–17864 (2023)
8. Du, S., Hers, B., Bayasi, N., Hamarneh, G., Garbi, R.: FairDisCo: Fairer ai in dermatology via disentanglement contrastive learning. In: European Conference on Computer Vision. pp. 185–202 (2022)
9. González, C., Ranem, A., Othman, A., Mukhopadhyay, A.: Task-agnostic continual hippocampus segmentation for smooth population shifts. In: Domain Adaptation and Representation Transfer MICCAI Workshop. pp. 108–118 (2022)

10. 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)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
12. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al.: Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* **114**(13), 3521–3526 (2017)
13. Kiyasseh, D., Zhu, T., Clifton, D.: A clinical deep learning framework for continually learning from cardiac signals across diseases, time, modalities, and institutions. *Nature Communications* **12**(1), 4221 (2021)
14. Lee, D., Jung, S., Moon, T.: Continual learning in the presence of spurious correlation. *arXiv preprint arXiv:2303.11863* (2023)
15. Lenga, M., Schulz, H., Saalbach, A.: Continual learning for domain adaptation in chest x-ray classification. In: *Medical Imaging with Deep Learning*. pp. 413–423 (2020)
16. Lesort, T.: Spurious features in continual learning. In: *AAAI Bridge Program on Continual Causality*. pp. 59–62 (2023)
17. Lewandowsky, S., Li, S.C.: Catastrophic interference in neural networks: Causes, solutions, and data. In: *Interference and inhibition in cognition*, pp. 329–361 (1995)
18. Lin, X., Kim, S., Joo, J.: FairGrape: Fairness-aware gradient pruning method for face attribute classification. In: *European Conference on Computer Vision*. pp. 414–432 (2022)
19. Luo, L., Xu, D., Chen, H., Wong, T.T., Heng, P.A.: Pseudo bias-balanced learning for debiased chest x-ray classification. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 621–631 (2022)
20. Mallya, A., Lazebnik, S.: Packnet: Adding multiple tasks to a single network by iterative pruning. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 7765–7773 (2018)
21. Nam, J., Cha, H., Ahn, S., Lee, J., Shin, J.: Learning from failure: De-biasing classifier from biased classifier. *Advances in Neural Information Processing Systems* **33**, 20673–20684 (2020)
22. Perkonig, M., Hofmanninger, J., Herold, C.J., Brink, J.A., Panykh, O., Prosch, H., Langs, G.: Dynamic memory to alleviate catastrophic forgetting in continual learning with medical imaging. *Nature communications* **12**(1), 5678 (2021)
23. Salman, H., Jain, S., Ilyas, A., Engstrom, L., Wong, E., Madry, A.: When does bias transfer in transfer learning? *arXiv preprint arXiv:2207.02842* (2022)
24. Tschandl, P., Rosendahl, C., Kittler, H.: The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data* **5**(1), 1–9 (2018)
25. Wang, L., Zhang, X., Su, H., Zhu, J.: A comprehensive survey of continual learning: Theory, method and application. *arXiv preprint arXiv:2302.00487* (2023)
26. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2097–2106 (2017)

- 27. Wang, Z., Yang, E., Shen, L., Huang, H.: A comprehensive survey of forgetting in deep learning beyond continual learning. arXiv preprint arXiv:2307.09218 (2023)
- 28. Wortsman, M., Ramanujan, V., Liu, R., Kembhavi, A., Rastegari, M., Yosinski, J., Farhadi, A.: Supermasks in superposition. Advances in Neural Information Processing Systems **33**, 15173–15184 (2020)
- 29. Zhang, Z., Sabuncu, M.: Generalized cross entropy loss for training deep neural networks with noisy labels. Advances in neural information processing systems **31** (2018)