Addressing catastrophic forgetting for medical domain expansion

Sharut Gupta

Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston, MA, USA

Praveer Singh

Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston, MA, USA <https://orcid.org/0000-0001-6641-2030>

Ken Chang

Massachusetts Institute of Technology, Cambridge, MA, USA

Liangqiong Qu

Stanford University

Mehak Aggarwal

Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston, MA, USA

Nishanth Arun

Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston, MA, USA

Ashwin Vaswani

Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston, MA, USA

Shruti Raghavan

Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston, MA, USA

Vibha Agarwal

Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston, MA, USA

Mishka Gidwani

Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston, MA, USA

Katharina Hoebel

Massachusetts Institute of Technology, Cambridge, MA, USA

Jay Patel

Massachusetts Institute of Technology, Cambridge, MA, USA

Charles Lu

Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston, MA, USA

Christopher Bridge

Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston, MA, USA

Jimmy Chen

Department of Ophthalmology, Casey Eye Institute, Oregon Health & Science University, Portland, Oregon

R. V. Paul Chan

Department of Ophthalmology and Visual Sciences, Illinois Eye and Ear Infirmary, University of Illinois at Chicago, Chicago, Illinois

Michael Chiang

Department of Ophthalmology, Casey Eye Institute, Oregon Health & Science University, Portland, Oregon

John Campbell

Oregon Health & Science University

Daniel Rubin

Stanford University <https://orcid.org/0000-0001-5057-4369>

Jayashree Kalpathy-Cramer (✉ JKALPATHY-CRAMER@mgh.harvard.edu)

Massachusetts General Hospital <https://orcid.org/0000-0001-8906-9618>

Article

Keywords: Deep Learning, Catastrophic Forgetting, Domain Expansion

Posted Date: November 17th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-1087025/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Addressing catastrophic forgetting for medical domain expansion

Sharut Gupta^{*1,2}, Praveer Singh^{*1}, Ken Chang^{*1,3}, Liangqiong Qu⁴, Mehak Aggarwal¹, Nishanth Arun¹, Ashwin Vaswani¹, Shruti Raghavan¹, Vibha Agarwal^{1,3}, Mishka Gidwani¹, Katharina Hoebel^{1,3}, Jay Patel^{1,3}, Charles Lu¹, Christopher P. Bridge¹, Jimmy S. Chen⁶, R. V. Paul Chan⁵, Michael F. Chiang⁶, J. Peter Campbell⁶, Daniel L. Rubin⁴, and Jayashree Kalpathy-Cramer¹

¹ Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston, MA, USA

jkalpathy-cramer@mgh.harvard.edu

<https://qtim-lab.github.io/>

² Indian Institute of Technology Delhi, New Delhi, India

³ Massachusetts Institute of Technology, Cambridge, MA, USA

⁴ Department of Radiology and Biomedical Data Science, Stanford University, Palo Alto, CA, USA

⁵ Department of Ophthalmology and Visual Sciences, Illinois Eye and Ear Infirmary, University of Illinois at Chicago, Chicago, Illinois

⁶ Department of Ophthalmology, Casey Eye Institute, Oregon Health and Science University, Portland, Oregon

Abstract. Model brittleness is a key concern when deploying deep learning models in real-world medical settings. A model that has high performance at one dataset may suffer a significant decline in performance when tested at on different datasets. While pooling datasets from multiple hospitals and re-training may provide a straightforward solution, it is often infeasible and may compromise patient privacy. An alternative approach is to fine-tune the model on subsequent datasets after training on the original dataset. Notably, this approach degrades model performance at the original datasets, a phenomenon known as *catastrophic forgetting*. In this paper, we develop an approach to address catastrophic forgetting based on elastic weight consolidation combined with modulation of batch normalization statistics under three scenarios: 1) for expanding the domain from one imaging system’s data to another imaging system’s 2) for expanding the domain from a large multi-hospital dataset to another single hospital dataset 3) for expanding the domain from dataset from one geographic region to a dataset from another geographic region. Focusing on the clinical uses cases of mammographic breast density detection and retinopathy of prematurity stage diagnosis, we show that our approach outperforms several other state-of-the-art approaches and provide theoretical justification for the efficacy of batch normalization modulation. The results of this study are generally applicable to the deployment of any clinical deep learning model which requires domain expansion.

Keywords: Deep Learning · Catastrophic Forgetting · Domain Expansion

1 Introduction

Deep learning (DL) models have shown state-of-the-art performance for a wide variety of computer vision [4,33,3,25], biomedical signal processing [11,27], and medical imaging tasks [5,16]. Within the clinical context, there is a need to continually refine these models to achieve high performance on new datasets, such as those from different image acquisition systems or hospitals. Typically, DL models are extended to new datasets via fine-tuning the model weights; a neural network trained on the original dataset is used to initialize a new model, which is then trained on the target domain [23]. However, this can result in *catastrophic forgetting* of the previous dataset, a phenomenon in which models do not preserve previously learned knowledge and consequently result in a degradation of the performance on the original dataset [8]. This poses a major challenge for regulatory agencies, such as the Food and Drug Administration (FDA) in the United States, as DL models that have been fine-tuned after their approval may no longer satisfy the required performance criteria on the original test set. In this work, we explore techniques to mitigate catastrophic forgetting in the setting of fine-tuning on new medical datasets, also known as domain expansion.

Prior work in addressing domain expansion have focused on interventions that target specific layers of a network, *e.g.*, by either fine-tuning trainable parameters of the Batch Normalization (BN) layers [12] or modulating the BN statistics (bias/variance) [17]; both of them attempting to rectify any differences in internal co-variate shift by aligning the distribution of the new dataset with the previous one, resulting in similar model performance across both datasets. However, these techniques necessarily limit the model’s capacity to incorporate new knowledge since all non-BN layers are kept frozen throughout the fine-tuning process.

*These authors contributed equally.

(a)

Dataset	Name	Composition	Acquisition System	Name	Size	Total Size
Institution 1	DMIST	Scanner A	General Electric Medical Systems	SenoGraphe	59411	103890
		-	Fuji Medical	Computed Radiography System	-	
		-	Hologic Mammography System	Digital Mammography System	-	
		-	Hologic Mammography System	Selenia Full Field System	-	
		Scanner B	Fischer Medical	SenoScan	32928	
Institution 2	MGH	-	Hologic mammography system	Lorad Selenia	8603	8603

(b)

Fig. 1: (a) Intensity histograms of 100 randomly selected images from each category showing heterogeneity across digital mammography systems (Scanner A on top-left vs. Scanner B on bottom-left) as well as across medical institutions (Institution 1 on top-right v/s Institution 2 on bottom-right); (b) Summary of dataset composition, acquisition systems, and size.

Some recent works [29,34] have shown that creation of site-specific models can potentially recover performance loss from catastrophic forgetting. However, having individualized models is more complicated to both train and deploy than a single, robust model is. Other approaches specifically handle catastrophic forgetting by constraining gradient updates to model parameters. For example, Elastic Weight Consolidation (EWC) [15] updates model parameters proportionally to the inverse of each parameter’s importance with respect to the original training dataset (*i.e.*, the magnitude of updates for more important parameters is small). Zeng et al. [37] claim EWC to be ineffective in retaining performance on the original task and instead propose Orthogonal Weight Modification (OWM) wherein model parameters are updated in a direction orthogonal to the subspace spanned by the inputs of the model. However, a major limitation of OWM in medical applications is that OWM requires the underlying model to first perform feature extraction from the combined dataset before training a multi-layered perceptron on top of the extracted features. This is usually infeasible in our problem setting where datasets between multiple hospitals cannot be combined at a central location due to infrastructural as well as patient privacy issues. Moreover, both EWC and OWM have thus far been evaluated mainly in the context of *continuous learning*, where the end goal is the sequential learning of separate tasks. This is different from our problem of domain expansion, where the goal is to learn a given task from data of a new domain without forgetting the original domain.

Unlike past works, we focus on the development of techniques to address catastrophic forgetting in a more *realistic setting*. We consider a real-world clinical application of mammographic breast density assessment, which is routinely used to assess breast cancer risk to decrease the chances of breast cancer mortality [32,26]. Specifically, the identification of patients with dense breast tissue warrants additional monitoring, such as supplemental ultrasound or magnetic resonance imaging. The current criteria for mammographic breast density classification is based on the Breast Imaging Reporting and Data System (BI-RADS), which divides breast density into four distinct categories: fatty, scattered, heterogeneously dense, and extremely dense [18]. Given the variability of manual interpretation, there has been interest in developing automated approaches for assessment of mammographic density [31]. A major hurdle for large-scale clinical deployment of a deep learning based breast density assessment tool is the poor generalizability across different hospitals/institutions and scanner types owing to inherent variability in patient demographics, disease prevalence, and imaging acquisition techniques [36,20,1]. Our study addresses heterogeneity in digital mammography systems across different institutions (as depicted in Fig. 1 a) that arises from variability in x-ray tube targets, filters, digital detector technology, and control of automatic exposure [13].

Additionally, we show the utility of our approach for diagnosis of stage in retinopathy of prematurity (ROP) based on retinal fundus photographs. This is of clinical importance because ROP is a leading cause of childhood blindness, especially in low- and middle- income countries. This complication can be prevented with effective screening, either by the clinician or potentially with the assistance of deep-learning based tools, followed by early treatment [2]. Similar to mammographic screening, a hurdle to deployment of algorithms is generalizability across different datasets. Specifically, there can be a drop in performance with different camera systems as well as different patient populations [2]. Our study addresses both of these forms of heterogeneity (as depicted in Fig. 7).

We investigate domain expansion techniques across different digital mammography systems (Senographe as Scanner A and SenoScan as Scanner B), different institutions (DMIST as Institution 1 and MGH as Institution 2)[†], and different countries (North America as Region 1 and Nepal as Region 2). with an objective to mitigate catastrophic forgetting. Dataset composition and size details are enlisted in Fig. 1, 7 b. For simplicity, the original domain (Scanner A, Institution 1, and Region 1), is referred to as *Dataset O* while the target domain (Scanner B, Institution 2, Region 2) is referred as *Dataset T*. The key contributions are as follows:

- We propose a simple yet effective technique to mitigate catastrophic forgetting by utilizing global BN statistics[‡] of Dataset O instead of Dataset T when fine-tuning on T.
- We demonstrate the efficacy of this technique under two different scenarios: first when restricting fine-tuning to only BN layers (motivated by [12]) and second when fine-tuning using all the layers.
- We demonstrate how a commonly-used continuous learning algorithm (EWC [15]) fails for large, real-world datasets and further highlight how augmenting EWC with our technique not only improves domain expansion over T but also mitigates catastrophic forgetting on O.
- Lastly, we provide a theoretical justification for why using global BN statistics of Dataset O instead of Dataset T better mitigates catastrophic forgetting.

2 Results

With the objective of fine-tuning our deep learning model on Dataset T and simultaneously mitigating catastrophic forgetting on Dataset O, we experiment with two broad sets of approaches:

1. Evaluation with global BN statistics of Dataset T (Fig. 2a Top) vs. Dataset O (Fig. 2a Bottom): both when fine-tuning with only BN layers[§] and when fine-tuning with all layers.
2. Evaluation after incorporating EWC (Fig. 2b) using global BN statistics of Dataset O, again when fine-tuning with only BN layers and when fine-tuning with all layers. Please refer to Section 5.2 in Methods for additional details.

[†]DMIST or digital mammographic imaging screening trial comprises of a consortium of institutions acquired from diverse digital mammography systems; MGH or Massachusetts General Hospital is a single institution dataset acquired from a single digital mammography system. See Section 5.4 in Methods for additional details.

[‡]In this work, global BN statistics of a particular dataset refers to the running mean & standard deviation of BN layers computed when training on that dataset. Typically, while batch statistics (batch mean & batch standard deviation) are used during training, the global BN statistics are used only during inference. See Section 5.1 in Methods for further details.

[§]We freeze all non-BN layers, allowing only BN parameters to train in this case.

Fig. 2: (a) The two different BN approaches for fine-tuning on Dataset T: (Top) Using global BN statistics of Dataset T (running mean $\mu_{global,T}$ and running standard deviation $\sigma_{global,T}$ are computed during fine-tuning process); (Bottom) Using global statistics of Dataset O (running mean $\mu_{global,O}$ and running standard deviation $\sigma_{global,O}$ computed when originally training the model on O). (b) A schematic of EWC trajectory with the loss landscape in red corresponding to Dataset O and the one in green to Dataset T. The black line represents the evolution of loss of the fully trained baseline model on Dataset O when fine-tuned on Dataset T using EWC. The loss trajectory culminates in close proximity to both global minimas, resulting in high performance for both datasets.

2.1 Baseline Experiments

As a baseline, we performed experiments with three approaches: 1) training solely on Dataset O, 2) training on Dataset O and fine-tuning on Dataset T, and 3) training on the combined datasets O and T. As shown in Fig. 3a, 3b, 8 (Panel “Baseline”), starting from Dataset O (Scanner A, Institution 1, or Region 1) and moving onto Dataset T (Scanner B, Institution 2, or Region 2), we observe that models exclusively trained on Dataset O do not generalize well on Dataset T (Fig. 3a, 3b, 8 i). However, a model originally trained on O, when fine-tuned on T (Fig. 3a, 3b, 8 ii), abruptly forgets the information it learnt on O ($p < 0.01$). Only when the model is trained collectively on Dataset O and T (Fig. 3a, 3b, 8 iii) does it achieve high performance on both domains (Scanner A κ : 0.67 and Scanner B κ : 0.69, Institution 1 κ : 0.67 and Institution 2 κ : 0.67, Region 1 AUROC: 0.974 and Region 2 AUROC: 0.976).

Fig. 3: Results for domain expansions across (a) digital mammography systems (Scanner A to Scanner B) and (b) institutions (Institution 1 to Institution 2). The different panels for each domain expansion type correspond to experiments for 1) Baseline models, 2) Only BN-layer fine-tuning with and without EWC[15] and 3) All layer Fine-tuning with and without EWC[15]. The parameter λ is used as a trade off between the relative importance of performance over Datasets O and T (higher the λ parameter, more closer will the performance of the fine-tuned model be to the model originally trained on Dataset O). The error bars indicate the 95% confidence interval over 5 independent experimental runs.

2.2 Global BN statistics of target dataset (T) vs. original dataset (O)

Fine-tuning BN layers: With the intent of aligning the data distribution of Dataset T with Dataset O, we run several experiments fine-tuning only the BN layers for Dataset T while freezing all the non-BN layers. We start with the baseline model previously trained on Dataset O. This model is fine-tuned with two distinct approaches. In the first approach (traditional method as show in Fig. 2a Top), this partially frozen model is fine-tuned using the batch statistics of Dataset T (batch mean $\mu_{batch,T}$ and batch standard deviation $\sigma_{batch,T}$) while its global BN statistics (running mean $\mu_{global,T}$ and running standard deviation $\sigma_{global,T}$) are concurrently calculated during the training process. The global BN statistics are then used to evaluate both datasets (Fig. 3a, 3b, 8 iv). In the second approach (as shown in Fig. 2a Bottom), the global BN statistics of Dataset O (running mean $\mu_{global,O}$ and running standard deviation $\sigma_{global,O}$) are used for fine-tuning this model on Dataset T and again used for evaluation on both the datasets (Fig. 3a, 3b, 8 v).

From Fig. 3a (Panel “BN Fine-Tuning + EWC”), we observe that when the global BN statistics of Scanner B (T) are used (Fig. 3a iv), the model undergoes a large reduction ($p < 0.01$) in the performance on Scanner A (O). Moreover, the performance on Scanner B (T) is lower ($p < 0.01$) compared to the performance when fine-tuned with all the layers (Fig. 3a ii). In contrast, when the model is fine-tuned using the global BN statistics of Scanner A (O) (Fig. 3a v), we see a recovery ($p < 0.01$) in the performance on Scanner A (O) as well as an increase in the performance on Scanner B (T) from the baseline model before fine-tuning ($p < 0.01$) (Fig. 3a i). Similar results are obtained for Institution 1 (O) and Institution 2 (T) as well as Region 1 (O) and Region 2 (T). From these results, we can conclude that fine-tuning BN layers using the global BN statistics of Dataset

O confers a performance advantage over fine-tuning using the global BN statistics of Dataset T. Moreover, although the performance on the target domain improves, fine-tuning with BN layers only partially mitigates catastrophic forgetting and therefore does not represent successful domain expansion.

Fine-tuning all layers: Starting with the same baseline model specification as above, we again experiment with fine-tuning all layers of the model including the BN layers to utilize the full capacity of the model for learning new knowledge from Dataset T. Similar to the above setup, this model is fine-tuned for all layers over Dataset T and evaluated first using the global BN statistics calculated while training on Dataset T (Fig. 3a, 3b, 8 ii) and second using the global BN statistics of Dataset O (Fig. 3a, 3b, 8 vii). When a model trained on Dataset O is fine-tuned for all layers on Dataset T using the global BN statistics of Dataset T (Fig. 3a, 3b, 8 ii), we again observe that performance on Dataset O degrades ($p < 0.01$ for Scanner A, Institution 1 and Region 1). Evaluating with global BN statistics of Dataset O (Fig. 3a, 3b, 8 vii) attenuates the performance loss ($p < 0.05$). With this approach, the model is able to perform well on the target domain (Scanner B, Institution 2, and Region 2) irrespective of the choice of global statistics used.

2.3 Incorporation of Elastic Weight Consolidation

Fine-tuning BN Layers with EWC: We explore the influence of EWC on the performance of the original and the target datasets when only the BN layers are fine-tuned using EWC (Fig. 3a, 3b, 8 vi). Performance on both datasets is evaluated after fine-tuning on Dataset T while varying the importance parameter (λ). After experimenting with a wide range of λ values, we observe that when using the global BN statistics of Dataset O we achieve maximum performance on both Dataset O and Dataset T at $\lambda = 0.05$ and $\lambda = 0.005$ for the mammography and ROP datasets, respectively (Fig. 4 (a), 4 (b), 9 i).

Fig. 4: Model performance while varying λ when fine-tuning only BN layers (top row: (i),(ii)) vs. all layers (bottom row: (iii),(iv)) with EWC[15] on (a) Scanner A and Scanner B using global BN statistics of Scanner A (left) vs. Scanner B (right) ; (b) Institution 1 and Institution 2 using global BN statistics of Institution 1 (left) vs. Institution 2 (right). Scanner A and Institution 1 being the original dataset O while Scanner B and Institution 2 being the target dataset T.

Fine-tuning all Layers with EWC: To further mitigate catastrophic forgetting on Dataset O and achieve peak performance on Dataset T, we allow fine-tuning of all layers while also incorporating EWC in the loss function. Performance on both datasets is evaluated after fine-tuning on Dataset T with varying importance parameters (λ). During fine-tuning with the global BN statistics of Dataset O, with increasing λ , performance on Dataset O consistently improves (red or blue curve in Fig. 4(a), 4(b), 9 iii). In contrast, for the target dataset T, we see a degradation in performance with increasing λ (orange or green curve in Fig. 4(a), 4(b), 9 iii). When fine-tuning with the global BN statistics of Dataset O for very high λ values ($\lambda = 1e+5$), catastrophic forgetting for Dataset O is mitigated, with a performance of $\kappa: 0.67$, $\kappa: 0.67$, and AUROC: 0.974 on Scanner A, Institution 1, and Region 1 respectively (Fig. 3a, 3b, 8 ix). However, very high λ values prevent the model from learning the features of the target domain, and hence perform poorly on Scanner B ($\kappa: 0.53$), Institution 2 ($\kappa: 0.23$), and Region 2 (AUROC 0.839) as seen in Fig. 3a, 3b, 9 ix.

Overall, for domain expansion from Scanner A (O) to Scanner B (T) using the global BN statistics of Scanner A (O), the best performance on both Scanner A ($\kappa: 0.60$) and Scanner B ($\kappa: 0.67$) is obtained at $\lambda = 0.005$

(Fig. 3a viii). For domain expansion from Institution 1 (O) to Institution 2 (T) using the global BN statistics of Institution 1 (O), the best performance on both Institution 1 (κ : 0.66) and Institution 2 (κ : 0.62) occurs at $\lambda = 1$ (Fig. 3b viii). For domain expansion from Region 1 (O) to Region 2 (T) using the global BN statistics of Region 1 (O), the best performance on both Region 1 (AUROC: 0.969) and Region 2 (AUROC: 0.951) occurs at $\lambda = .5$ (Fig. 8 viii). Interestingly, at the λ which gives highest performance on both Datasets, fine-tuning only BN layers rather than all layers (red and blue curves in Fig. 4(a), 4(b), 9 when comparing i vs. iii) gives higher performance on Dataset O for all 3 scenarios, although significant for only 2 of the 3 ($p < 0.01$ for domain expansion across Scanner A to Scanner B and Institution 1 to Institution 2; $p = 0.25$ across Region 1 and 2). Conversely, fine-tuning all layers rather than just BN layers (orange and green curves in Fig. 4(a), 4(b), 9 when comparing i vs. iii) results in higher performance on Dataset T ($p < 0.01$ for all datasets). For both BN-only and all Layers fine-tuning, using Elastic Weight Consolidation (EWC) in conjunction with global BN statistics of the target dataset T, is ineffective for mitigating catastrophic forgetting as compared to when using EWC with the global BN statistics of Dataset O (Fig 4(a), 4(b), 9) comparing red or blue curves for left vs. right plots) further substantiating the importance of using global BN statistics of original dataset O ($p < .05$). Thus, in summary, domain expansion is optimally performed by using both, the global BN statistic of Dataset O (to effectively mitigate catastrophic forgetting) and all Layers fine-tuning together with EWC (to attain peak performance on Dataset T).

3 Discussion

Continuous learning and multi-task learning have been an active area of research in the machine learning community. Although some techniques have been proposed [15,38,6,37], they work well on simpler benchmark computer vision datasets (*e.g.*, Permuted MNIST [8]) and network architectures with few layers (2-3 densely connected layers), they often fail (see Section 9.2 in the supplementary materials) when using deeper architectures with more complex layers and functions such as Dropout, ReLU activation, or BN. In this paper, we demonstrate the utility of our approach for domain expansion for real-world use cases in radiology and ophthalmology.

Normalization layers have become a ubiquitous technique for training deep neural networks. Batch Normalization (BN) specifically, normalizing the output from a layer using the mean and standard deviation of the input batch, effectively allows training with higher learning rates and usually improves convergence to a better local optimum [10]. While most existing work focuses on the influence of BN when training models on a *single institution dataset*, few [12] have actually tried to adapt the BN layers for training these models sequentially over multiple datasets with varied scanner/camera types and patient populations. Karani et. al. [12] froze the convolutional layers and trained only the batch norm layers independently for each dataset. Although this approach allows the model to retain performance on the previously trained distribution and task, it restricts model capacity to learn newer tasks or domains. Moreover, every institution would need to have separate BN parameters, which defeats the purpose of *one model for both the original and target datasets*. In our approach, we allow training of all layers on the target dataset using EWC[15], while simultaneously fixing BN statistics to the global BN statistics from the original dataset O, thus preventing catastrophic forgetting on O. Contrary to Karani et. al. [12], our approach provides one universal model for both original and target distributions involved during domain expansion.

Plasticity is the capacity of a network to adapt to new environments while stability ensures retention of previously learned knowledge. Setting the appropriate trade-off between stability and plasticity of a neural network is critical, not only to avoid forgetting but also to learn new tasks quickly. Catastrophic forgetting occurs when a network is overly plastic and not sufficiently stable; *i.e.*, the network is able to quickly acquire new tasks or modalities but does not retain previously learned tasks and modalities. [22] studied this behaviour in the context of dropout layers and showed how dropout can overcome catastrophic forgetting on a previous task when the same model is trained on a new task. Specifically, they showed that multi-layered perceptrons trained with dropout regularization have higher stability and lower variance in the output after ReLU activation function. In other words, activation which are purely active or inactive (values 0 or 1) remain untouched (ensuring stability on previously learned tasks) and only the semi-active neurons are turned off or on, thus promoting plasticity for newer tasks. EWC [15] works by a similar principle, *i.e.*, ensuring stability by constraining those weights which are most important for previous learned tasks, while still retaining enough plasticity by allowing other weights to be trained for newer tasks.

In order to better understand the influence of global BN statistics on overall stability of a network fine-tuned over Dataset T, we plot the distribution of activation values from different BN layer outputs in the network using samples from Dataset O (Fig. 5 for Institution 1). Using global BN statistics of O (Fig. 5 left), we empirically see that deeper layers of the network (Fig. 5 top vs. bottom), have a lower activation variance than when using global BN statistics of T (Fig 5 right). Notably, while the mean variance across different channels for the model fine-tuned using global BN statistics of O increases from 0.01 to 0.63, the model fine-tuned using the global BN statistics of T shows a much greater increase from 0.03 to 0.97. For a detailed, theoretical proof corresponding to this observation, we refer the reader to Section 5.3 in Method.

Fig. 5: Probability Density of per-channel activation values computed over a batch from original Dataset O (Institution 1) after fine-tuning on target Dataset T (Institution 2) using global statistics of O (left) and T (right) for an initial layer Layer1.0.bn1 (top) and a deeper layer Layer4.2.bn3 (bottom) of the Resnet-50 architecture. Here layer1.0.bn1 represents the first BN layer in the Bottleneck layer of ResNet50 architecture while layer4.2.bn3 represents the last BN layer before the final classification layer

Inspired by [22], we further hypothesize that as we move deeper into a network fine-tuned over Dataset T, the lower variance (from using global BN statistics of O instead of T) promotes greater stability of the network and would therefore lead to lower plasticity and less catastrophic forgetting. To further investigate the influence of global BN statistics on catastrophic forgetting, we visualize the UMAPs [21] of features computed over Dataset O from the various models in our experiments. Fig. 6(i) depicts UMAPs of sample features for the model trained originally on Dataset O (Institution 1). We observe clear separation for different classes, aligned along a particular direction in low-dimensional space that likely corresponds with high overall performance. In Fig. 6(ii) or 6(iv) (representing models fine-tuned with only BN layers and fine-tuned with all layers respectively by using global BN statistics of T (Institution 2)), we observe that UMAPs of features from Dataset O have classes which are highly entangled with each other, which correlate with lower, overall performance on O. However, for Fig. 6(iii) or 6(v) (representing models fine-tuned with only BN layers and fine-tuned with all layers, respectively, by using global BN statistics of O), we observe how different classes which show high separation results in avoiding any forgetting. This is also observed when fine tuning all layers after incorporating EWC and using global statistics of O (Fig 6 (vi) shows classes are well segregated and aligned in one direction). Overall, this illustrates how global BN statistics of O prevent large variance shifts in the deeper layers, which may avoid any serious entanglement of features in deeper layers, and ultimately results in higher stability of the network to mitigate catastrophic forgetting.

There are several limitations and possible future directions to this work. First, we only evaluate our approach in the context of domain expansion where a model trained on one dataset is fine tuned on a second dataset. Future work can investigate the scalability of our approach when domain expansion is performed across a larger number of datasets. Secondly, both clinical questions investigated in this study were classification tasks. Additional study evaluating the utility of our approach for other clinical tasks such as detection and segmentation is warranted.

Lastly, federated learning has received considerable attention recently, primarily for training models over multi-institutional medical datasets, which are difficult to share due to both privacy and legal issues [29,30]. Federated learning presents an alternative to serial fine-tuning by training on multiple datasets simultaneously. However, data heterogeneity remains an optimization challenge for federated learning. Further comparison of performance of our approach with federated learning under real world settings can be a topic of future investigation.

Fig. 6: UMAPs for original Dataset O (Institution 1) when (i) Model is trained just on Dataset O; When fine-tuning only BN layers with global BN statistics of (ii) T (Institution 2); (iii) O; Fine-tuning all layers with global BN statistics of (iv) T; (v) O; (vi) O after incorporating EWC (with optimal λ value)

4 Conclusion

In this study, we develop a simple yet novel approach to avoid catastrophic forgetting while fine-tuning deep learning models on new datasets, enabling continuous learning and refinement. Catastrophic forgetting has major implications for clinical deployment of deep learning models, where generalizability remains a major hurdle to large-scale deployment. Specifically, we demonstrate that relying on global BN statistics of the original Dataset O when fine-tuning on the target Dataset T alleviates forgetting to a large extent, though the performance achieved on Dataset T isn't optimal. However, when this technique is used in conjunction with EWC, we are not only able to completely mitigate catastrophic forgetting on Dataset O but also achieve peak performance on Dataset T.

5 Methods

5.1 Global BN statistics of target dataset (T) vs. original dataset (O)

Batch normalization (BN) has become a popular technique to regulate shifts in the distribution of network activations during training [10]. BN normalizes each batch of input to the network layers by subtracting the batch mean and dividing by the batch standard deviation. Fig. 2a (top row) illustrates pipeline for traditional fine-tuning on Dataset T starting from a pre-trained model on Dataset O. During fine tuning, output from the previous convolutional layer is passed to the BN layer, which normalizes the input $x_1, x_2 \dots x_m$ (where m is the batch size):

$$y_i = \gamma_T \frac{x_i - \mu_{batch,T}}{\sqrt{\sigma_{batch,T}^2 + \epsilon}} + \beta_T, \quad (1)$$

where y_i is the output after applying BN on input x_i , $\mu_{batch,T}$ and $\sigma_{batch,T}^2$ are the batch mean and batch variance of Dataset T, respectively, ϵ is a stability parameter, and γ_T and β_T are trainable parameters of BN layers when fine-tuned with Dataset T. During inference, BN uses the global BN statistics of Dataset T (running mean $\mu_{global,T}$ and the running standard deviation $\sigma_{global,T}$) which were calculated while training; specifically, at training iteration k , the running mean $\mu_{global,T}[k]$ and running variance $\sigma_{global,T}^2[k]$ for Dataset T are updated:

$$\mu_{global,T}[k] = \alpha \mu_{global,T}[k-1] + (1 - \alpha) * \mu_{batch,T}[k]$$

$$\sigma_{global,T}^2[k] = \alpha \sigma_{global,T}^2[k-1] + (1 - \alpha) * \sigma_{batch,T}^2[k]$$

where the default momentum parameter, α , is set at 0.9. During inference, we replace $\mu_{batch,T}$ and $\sigma_{batch,T}^2$ in equation 1 with the global BN statistics of Dataset T computed during training, resulting in:

$$y_i = \gamma_T \frac{x_i - \mu_{global,T}}{\sqrt{\sigma_{global,T}^2 + \epsilon}} + \beta_T \quad (2)$$

In addition to the traditional approach to compute global BN statistics ($\mu_{global,T}$ and $\sigma_{global,T}$) fine-tuned on Dataset T and evaluated with the same global BN statistics for both datasets, we also conduct experiments to both fine-tune and evaluate using global BN statistics of Dataset O (as shown in bottom row of Fig 2a). Thus, the BN (both for fine-tuning and inference) on an input batch can be represented by:

$$y_i = \gamma_T \frac{x_i - \mu_{global,O}}{\sqrt{\sigma_{global,O}^2 + \epsilon}} + \beta_T \quad (3)$$

where $\mu_{global,O}$ and $\sigma_{global,O}$ are the global BN statistics of Dataset O (running mean and running standard deviation computed when training the model originally on Dataset O); γ_T and β_T denote the trainable parameters of BN layers when fine-tuning with Dataset T but using global BN statistics of O.

We experiment with fine-tuning the model using the above two global BN statistics techniques under the following two scenarios:

- **Fine tuning only BN layers:** As shown in Fig 1, the intensity distributions of Dataset O (Scanner A in red and Institution 1 in blue) differ considerably from the Dataset T (Scanner B in orange and Institution 2 in green), despite all being mammography screening datasets. With the intuition of aligning the data distribution of Dataset T with Dataset O, and retaining the original model performance for Dataset O, we first test run several experiments by fine-tuning only BN layers on Dataset T while freezing all other layers.
- **Fine-tuning all layers (BN and convolutional layers):** With the primary objective of achieving the best performance on Dataset T while retaining the original performance on Dataset O, we further investigate the effects of fine-tuning all layers. Training using all layers aids in leveraging the full capacity of a model for fine-tuning on Dataset T.

5.2 Incorporation of Elastic Weight Consolidation (EWC)

While fine-tuning a model (previously trained on Dataset O) on Dataset T, EWC [15] effectively constrains the updates of those weights which are most important for Dataset O, thus allowing the model to converge to a minimum close to global minima of both datasets T and O (as shown in Fig 2b and ultimately prevents catastrophic forgetting for Dataset O).

We implement this constraint using a quadratic penalty on the change in model parameters. In order to identify the salient parameters for model trained on Dataset O, we first compute the empirical Fisher Information Matrix (FIM). [19] The Fisher information matrix is defined as the covariance of the score function used to calculate the quality of parameter estimation. Due to the complexity of the likelihood function, computing this expectation becomes intractable and hence an approximation called the empirical Fisher (F) is defined as

$$F = \frac{1}{N} \sum_{i=1}^N \nabla_{\theta} \log p(x|\theta) \nabla_{\theta} \log p(x|\theta)' \quad (4)$$

This empirical Fisher matrix when combined with the quadratic penalty adds a constraint to the important parameters of model trained on Dataset O. The combined loss function $L(\theta)$ for elastic weight consolidation is given by

$$L(\theta) = L_T(\theta) + \sum_i \frac{\lambda}{2} F_i (\theta_i - \theta_{O,i}^*)^2 \quad (5)$$

where $L_T(\theta)$ is the loss function for Dataset T, θ_O^* represents the optimal model parameters for Dataset O and i in θ_i iterates over all the current model parameters which are being fine-tuned on Dataset T. The parameter λ is used as a trade off between the relative importance of performance over Datasets O and T (the

higher the λ parameter, the more closer will be the performance of the fine-tuned model to the model originally trained on Dataset O).

As a result, a constraint on some (not all) model parameters, allows the solution to stay in a low-error region which is optimal for both the datasets [15]. We experiment with fine-tuning the model using EWC together with global BN statistics of O under the following two conditions: fine-tuning only the BN layers, and fine-tuning all layers.

5.3 Comparative analysis of variance shifts with global BN statistics of T and O

From [22], we know that the variance of the penultimate layer's activations of a network directly impacts the stability of a network. In other words, the lower the variance of the activations of the latter layers, the smaller the extent of model's catastrophic forgetting is likely to be. In order to illustrate the influence of using the global BN statistics of T and O on the overall stability of a network, we compute the variance of the outputs of the BN layers of a model fully fine-tuned on Dataset T (without EWC), given in equations 2 and 3. The variance of a layer's activations for the model trained using global BN statistics of T from equation 2 is given by:

$$Var[y_i] = Var \left[\gamma_T \frac{x_i - \mu_{global,T}}{\sqrt{\sigma_{global,T}^2 + \epsilon}} + \beta_T \right] \quad (6)$$

Assuming γ_T , β_T , $\mu_{global,T}$ and $\sigma_{global,T}^2$ to be constant at the time of inference and thereafter for the purpose of understanding catastrophic forgetting, Dataset O is used for inference:

$$Var[\beta_T] = 0 \quad (7)$$

$$Var \left[\gamma_T \frac{x_i - \mu_{global,T}}{\sqrt{\sigma_{global,T}^2 + \epsilon}} \right] = \frac{\gamma_T^2}{\sigma_{global,T}^2 + \epsilon} Var(x_i) \quad (8)$$

Hence, the final variance of output from BN layer of a model trained using global BN statistics of T is given by:

$$Var[y_i] = C_T \times Var(x_i) \quad (9)$$

where C_T is any constant. Next, the variance for the global BN statistics of O from equation 3 is given by:

$$Var[y_i] = Var \left[\gamma_T \frac{x_i - \mu_{global,O}}{\sqrt{\sigma_{global,O}^2 + \epsilon}} + \beta_T \right] \quad (10)$$

Again assuming γ_T , β_T , $\mu_{global,O}$ and $\sigma_{global,O}^2$ to be constant at the time of inference and thereafter for the purpose of understanding catastrophic forgetting, Dataset O is used for inference:

$$Var[\beta_T] = 0 \quad (11)$$

$$Var \left[\frac{x_i - \mu_{global,O}}{\sqrt{\sigma_{global,O}^2 + \epsilon}} \right] \approx 1 \quad (12)$$

Hence, the final variance of output from BN layer of a model trained using global BN statistics of O is given by:

$$Var[y_i] = \gamma_T^2 = C_T \quad (13)$$

Thus while the variance in the case of the model trained using the global BN statistics of O (equation 13) is simply a constant (based on the scaling parameter of the corresponding BN layer), the variance in the case of model trained using the global BN statistics of T (equation 9) is a constant times the variance of the input to the BN layer. In other words, though we see some variance shift in both cases, the shift in case of global BN statistics of T (equation 9) keeps on increasing as we go deeper into the network (owing to its direct proportionality to input variance). This increase in the variance shift over the layers of a network, ultimately results in extreme distortions in the features computed from the penultimate layers (as seen in Fig. 6) of a model trained using global BN statistics of T when tested on the original dataset O, thus leading to poor stability of the network and ultimately resulting in adverse catastrophic forgetting over O.

5.4 Datasets and Preprocessing

Mammographic Breast Density Digital screening mammograms from 33 institutions were retrospectively obtained through the Digital Mammographic Imaging Screening Trial (DMIST), the details of which were previously published [24]. This study was approved by the Institutional Review Board (IRB) of the American College of Radiology Imaging Network (ACRIN), by the IRB and the Cancer Therapy Evaluation Program at the National Cancer Institute. For DMIST, a total of 92 radiologists from the United States and Canada read the exams. Readers in the United States were all qualified interpreters of mammograms under federal law. Canadian readers met equivalent standards. Each site’s lead radiologist received training to read for DMIST and in turn trained the site’s other readers. DMIST images were previously de-identified for this study. For DMIST, 5 digital mammography systems were used: SenoScan (Fischer Medical), the Computed Radiography System for Mammography (Fuji Medical), the Senographe 2000D (General Electric Medical Systems), the Digital Mammography System (Hologic), and the Selenia Full Field Digital Mammography System (Hologic)[24]. The mammograms were saved in DICOM format with 4 different image data formats, corresponding to different digital-mammography systems or different versions of the same system: 12 bit Monochrome 1 (30.3%), 12 bit Monochrome 2 (11.2%), 14 bit Monochrome 1 (58.0%), and 14-bit Monochrome 2 (0.5%). 14-bit Monochrome 2 images were excluded to ensure that each image data format included in our study had adequate representation for training of our deep learning model. The final DMIST patient cohort consisted of 108,230 digital screening images from 21,759 patients.

We also obtained digital screening mammograms from Massachusetts General Hospital (MGH) from 2010 following IRB approval. Patients who had prior surgery or implants were excluded. For MGH, each image was read by a breast imaging radiologist as part of routine clinical practice. All mammograms were acquired using a Lorad Selenia mammography system (Hologic). The final MGH patient cohort consisted of 8,603 digital screening images from 1,856 patients. All images from DMIST and MGH were interpreted by a single radiologist from a pool of radiologists using the ACR BI-RADS breast density lexicon (Category A: fatty, Category B: scattered, Category B: heterogeneously dense, Category D: extremely dense)[18]. The wide variations in intensity histograms between the four datasets is illustrated in Fig. 1a.

The SenoScan (Scanner B, $n = 32928$), Senographe (Scanner A, $n = 59411$), DMIST (Institution 1, $n = 103890$), MGH (Institution 2, $n = 8603$) patient cohorts were split into training, validation, and testing sets in a 7:2:1 ratio, on a patient level. The training set was used to develop the model and the validation set was used to assess model performance during training to prevent overfitting. The test set was unseen until the model training was completed. The intensity of each image was scaled to be between 0 and 1 by dividing by the maximum value of the image format (4095 for 12 bit and 16383 for 14 bit). All Monochrome 1 images were inverted to make them equivalent to Monochrome 2. To ensure proper input size to the pre-trained neural network architectures, the images were resized to $224 \times 224 \times 3$.

Retinopathy of Prematurity Stage Retinal fundus photographs were acquired from two geographic regions: North America and Nepal. The North America dataset came from 9 institutions (Oregon Health and Science University, Columbia University, Cornell University, University of Illinois at Chicago, William Beaumont Hospital, Children’s Hospital Los Angeles, Cedars-Sinai Medical Center, University of Miami, and Asociacion para Evitar la Ceguera en Mexico). The Nepal dataset came from 4 urban hospitals Kathmandu, Nepal (Patan Hospital, Kanti Children’s Hospital, Paropakar Maternity and Women’s Hospital, and the Tilganga Institute of Ophthalmology). The study was approved by the IRB at all institutions [2].

The North American dataset contained nasal, temporal, and posterior pole retinal fundus images acquired using a RetCam camera. The Nepal dataset contained images from all fields of view acquired using a 3nethra camera. All images were classified as either stage (1, 2, or 3) or without stage using both the image and clinical examination findings based on the consensus of 3 experts as previously described [2]. Images were excluded if treatment (such as laser photocoagulation) was visible, if camera artifact obscured >50% of the image, or if retinal detachment was present (stage ≥ 4)[2].

For preprocessing, images were resized into 224×224 as previously described [2]. Each dataset was then split into training, validation, and testing sets on the patient level. The North American dataset contained 4121 training (3019 stage, 1102 no stage), 1114 validation (734 stage, 380 no stage), and 660 testing (449 stage, 211 no stage) images. The Nepal dataset contained 3830 training (3661 stage, 169 no stage), 972 validation (933 stage, 39 no stage), and 247 testing (203 stage, 44 no stage).

5.5 Training and Prediction

Mammographic Breast Density For the baseline classification model, a Resnet50 architecture with ImageNet-pretrained weights[9] was used. Cross entropy loss function and the Adam optimizer [14] ($\text{lr} = 10^{-6}$, $\beta_1 = 0.9$, $\beta_2 = 0.99$, $\epsilon = 10^{-7}$) were used across models. Batches were randomly sampled using a fixed batch size of 32 images. Early stopping with a patience of 20 epochs is used to prevent overfitting. Checkpoints were saved after

each epoch based on the performance on the validation set and the model with highest validation accuracy was saved as the final model. The training set is augmented in real time by means of random flips and rotations (-45° - 45°).

To combine predictions from all images across all mammography views from a given patient study into a patient-level assessment, the output probabilities for all corresponding images from a patient study are averaged. The averaged probabilities are then used to determine the predicted breast density class. All models are evaluated using Cohen's Kappa scores with linear weighting (κ). For reference, a κ of 0.21-0.40, 0.41-0.60, and 0.61-0.80 represents fair, moderate, and substantial agreement, respectively [28]. A non-parametric two-sided Wilcoxon signed-rank test at a significance level of $p = 0.05$ was used for statistical comparisons of model performance.

Retinopathy of Prematurity Stage For the baseline classification model, a Resnet18 architecture with ImageNet-pretrained weights[9] was used. Binary cross entropy loss function and the Adam optimizer [14] ($\text{lr} = 10^{-4}$, $\beta_1 = 0.9$, $\beta_2 = 0.99$, $\epsilon = 10^{-7}$) were used across models. Batches were randomly sampled using a fixed batch size of 64 images. Checkpoints were saved after each epoch based on the performance on the validation set and the model with lowest validation loss was saved as the final model. The training set is augmented in real time by means of random intensity scaling (0-1), rotation (-15° - 15°), and flips.

Performance was evaluated by means of area under the receiver operating characteristic curve on the image level. A Wilcoxon signed-rank test at a significance level of $p = 0.05$ was used for statistical comparisons of model performance.

6 Data Availability Statement

The main data supporting the results in this study are available within the paper and its Supplementary Information. The raw datasets from DMIST, MGH, North America, and Nepal is not publicly available due to patient privacy restrictions, though potential collaborators are directed to the corresponding author.

7 Code Availability Statement

All code from our method is available at: <https://github.com/QTIM-Lab/MedicalDomainExpansion>

References

1. Chang, K., Beers, A.L., Brink, L., Patel, J.B., Singh, P., Arun, N.T., Hoebel, K.V., Gaw, N., Shah, M., Pisano, E.D., et al.: Multi-institutional assessment and crowdsourcing evaluation of deep learning for automated classification of breast density. *Journal of the American College of Radiology* (2020)
2. Chen, J.S., Coyner, A.S., Ostmo, S., Sonmez, K., Bajimaya, S., Pradhan, E., Valikodath, N., Cole, E.D., Al-Khaled, T., Chan, R.P., et al.: Deep learning for the diagnosis of stage in retinopathy of prematurity: Accuracy and generalizability across populations and cameras. *Ophthalmology Retina* (2021)
3. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 801–818 (2018)
4. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709* (2020)
5. Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S.: Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **542**(7639), 115–118 (2017)
6. Farajtabar, M., Azizan, N., Mott, A., Li, A.: Orthogonal gradient descent for continual learning. In: *International Conference on Artificial Intelligence and Statistics*. pp. 3762–3773. PMLR (2020)
7. Farquhar, S., Gal, Y.: Towards robust evaluations of continual learning (2019)
8. Goodfellow, I.J., Mirza, M., Xiao, D., Courville, A., Bengio, Y.: An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211* (2013)
9. 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)
10. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167* (2015)
11. Jing, J., Sun, H., Kim, J.A., Herlopian, A., Karakis, I., Ng, M., Halford, J.J., Maus, D., Chan, F., Dolatshahi, M., et al.: Development of expert-level automated detection of epileptiform discharges during electroencephalogram interpretation. *JAMA neurology* **77**(1), 103–108 (2020)
12. Karani, N., Chaitanya, K., Baumgartner, C., Konukoglu, E.: A lifelong learning approach to brain mr segmentation across scanners and protocols. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 476–484. Springer (2018)

13. Keavey, E., Phelan, N., O'Connell, A., Flanagan, F., O'Doherty, A., Larke, A., Connors, A.: Comparison of the clinical performance of three digital mammography systems in a breast cancer screening programme. *The British journal of radiology* **85**(1016), 1123–1127 (2012)
14. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
15. 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)
16. Li, M.D., Chang, K., Bearce, B., Chang, C.Y., Huang, A.J., Campbell, J.P., Brown, J.M., Singh, P., Hoebel, K.V., Erdoğan, D., et al.: Siamese neural networks for continuous disease severity evaluation and change detection in medical imaging. *NPJ digital medicine* **3**(1), 1–9 (2020)
17. Li, Y., Wang, N., Shi, J., Liu, J., Hou, X.: Revisiting batch normalization for practical domain adaptation. *arXiv preprint arXiv:1603.04779* (2016)
18. Liberman, L., Abramson, A.F., Squires, F.B., Glassman, J., Morris, E., Dershaw, D.D.: The breast imaging reporting and data system: positive predictive value of mammographic features and final assessment categories. *AJR. American journal of roentgenology* **171**(1), 35–40 (1998)
19. Ly, A., Marsman, M., Verhagen, J., Grasman, R., Wagenmakers, E.J.: A tutorial on fisher information (2017)
20. Mårtensson, G., Ferreira, D., Granberg, T., Cavallin, L., Oppedal, K., Padovani, A., Rektorova, I., Bonanni, L., Pardini, M., Kramberger, M.G., et al.: The reliability of a deep learning model in clinical out-of-distribution mri data: a multicohort study. *Medical Image Analysis* p. 101714 (2020)
21. McInnes, L., Healy, J., Melville, J.: Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426* (2018)
22. Mirzadeh, S.I., Farajtabar, M., Ghasemzadeh, H.: Dropout as an implicit gating mechanism for continual learning. *arXiv preprint arXiv:2004.11545* (2020)
23. Mohamed, A., Berg, W., Peng, H., Luo, Y., Jankowitz, R., Wuand, S.: A deep learning method for classifying mammographic breast density categories. *Medical physics* **45**(1), 314–321 (2017)
24. Pisano, E.D., Gatsonis, C., Hendrick, E., Yaffe, M., Baum, J.K., Acharyya, S., Conant, E.F., Fajardo, L.L., Bassett, L., D'Orsi, C., et al.: Diagnostic performance of digital versus film mammography for breast-cancer screening. *New England Journal of Medicine* **353**(17), 1773–1783 (2005)
25. Rana, A., Singh, P., Valenzise, G., Dufaux, F., Komodakis, N., Smolic, A.: Deep tone mapping operator for high dynamic range images. *IEEE Transactions on Image Processing* **29**, 1285–1298 (2020). <https://doi.org/10.1109/TIP.2019.2936649>
26. Razzaghi, H., Troester, M.A., Gierach, G.L., Olshan, A.F., Yankaskas, B.C., Millikan, R.C.: Mammographic density and breast cancer risk in white and african american women. *Breast cancer research and treatment* **135**(2), 571–580 (2012)
27. Ribeiro, A.H., Ribeiro, M.H., Paixão, G.M., Oliveira, D.M., Gomes, P.R., Canazart, J.A., Ferreira, M.P., Andersson, C.R., Macfarlane, P.W., Wagner Jr, M., et al.: Automatic diagnosis of the 12-lead ecg using a deep neural network. *Nature communications* **11**(1), 1–9 (2020)
28. Richard, L., Gary, K.: The measurement of observer agreement for categorical data. *arxiv e-prints*, page. *Biometrics* **33**, 159–174 (1977)
29. Roth, H.R., Chang, K., Singh, P., Neumark, N., Li, W., Gupta, V., Gupta, S., Qu, L., Ihsani, A., Bizzo, B.C., et al.: Federated learning for breast density classification: A real-world implementation. *Lecture Notes in Computer Science* p. 181–191 (2020). https://doi.org/10.1007/978-3-030-60548-3_18, http://dx.doi.org/10.1007/978-3-030-60548-3_18
30. Sheller, M.J., Edwards, B., Reina, G.A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R.R., et al.: Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific reports* **10**(1), 1–12 (2020)
31. Sprague, B.L., Conant, E.F., Onega, T., Garcia, M.P., Beaber, E.F., Herschorn, S.D., Lehman, C.D., Tosteson, A.N., Lacson, R., Schnall, M.D., et al.: Variation in mammographic breast density assessments among radiologists in clinical practice: a multicenter observational study. *Annals of internal medicine* **165**(7), 457–464 (2016)
32. Tabár, L., Vitak, B., Chen, H.H.T., Yen, M.F., Duffy, S.W., Smith, R.A.: Beyond randomized controlled trials: organized mammographic screening substantially reduces breast carcinoma mortality. *Cancer: Interdisciplinary International Journal of the American Cancer Society* **91**(9), 1724–1731 (2001)
33. Tan, M., Pang, R., Le, Q.V.: Efficientdet: Scalable and efficient object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10781–10790 (2020)
34. Yu, T., Bagdasaryan, E., Shmatikov, V.: Salvaging federated learning by local adaptation. *arXiv preprint arXiv:2002.04758* (2020)
35. Zagoruyko, S., Komodakis, N.: Wide residual networks. *arXiv preprint arXiv:1605.07146* (2016)
36. Zech, J.R., Badgeley, M.A., Liu, M., Costa, A.B., Titano, J.J., Oermann, E.K.: Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. *PLoS medicine* **15**(11) (2018)
37. Zeng, G., Chen, Y., Cui, B., Yu, S.: Continual learning of context-dependent processing in neural networks. *Nature Machine Intelligence* **1**(8), 364–372 (2019)
38. Zenke, F., Poole, B., Ganguli, S.: Continual learning through synaptic intelligence. *Proceedings of machine learning research* **70**, 3987 (2017)

8 Acknowledgements

Research reported in this publication was supported by a training grant from the National Institute of Biomedical Imaging and Bioengineering (NIBIB) of the National Institutes of Health under award number 5T32EB1680 to J. B. Patel. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

This publication was supported from the Martinos Scholars fund to K. Hoebel. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the Martinos Scholars fund.

This study was supported by National Institutes of Health (NIH) grants U01CA154601, U24CA180927, and U24CA180918 to J. Kalpathy-Cramer, U01CA242879 to D. Rubin and J. Kalpathy-Cramer, and National Science Foundation (NSF) grant NSF1622542 to J. Kalpathy-Cramer.

This research was carried out in whole or in part at the Athinoula A. Martinos Center for Biomedical Imaging at the Massachusetts General Hospital, using resources provided by the Center for Functional Neuroimaging Technologies, P41EB015896, a P41 Biotechnology Resource Grant supported by the National Institute of Biomedical Imaging and Bioengineering (NIBIB), National Institutes of Health.

Data Collection as well as Cleaning for the ROP datasets was supported by grants R01 EY19474, R01 EY031331, R21 EY031883, and P30 EY10572 from the National Institutes of Health (Bethesda, MD), and by unrestricted departmental funding and a Career Development Award (JPC) from Research to Prevent Blindness (New York, NY). The Nepal screening program was supported by the United States Agency for International Development and Helen Keller International.

We would additionally like to acknowledge the tireless effort of Susan Ostmo, Kemal Sonmez, Sanyam Bajimaya, Eli Pradhan, Nita Valikodath, Emily D Cole and Tala Al-Khaled for participating in annotation of the Nepal and ICROP datasets for ROP stage diagnosis.

Lastly, we would also like to thank ECOG-ACRIN for providing access to the DMIST dataset and to the ACR-DSI for valuable discussions. Original data collection for ACRIN 6652 (DMIST) was supported by NCI Cancer Imaging Program grants.

9 Supplementary Material

9.1 Domain expansion across different regions for diagnosis of stage in retinopathy of prematurity

(a)

Dataset	Name	Camera System	Total Size
Region 1	North America	RetCam	5895
Region 2	Nepal	3nethra	5049

(b)

Fig. 7: (a) Intensity histograms across RGB channels of 100 randomly selected images from each category showing heterogeneity across retinal fundus photographs ((i) represents Region 1 and (ii) represents Region 2) (b) Summary of camera systems and size.

Fig. 8: Results for domain expansions across regions (Region 1 to Region 2). The different panels for each domain expansion type correspond to experiments for 1) Baseline models, 2) Only BN-layer fine-tuning with and without EWC[15] and 3) All layer Fine-tuning with and without EWC[15]. The parameter λ is used as a trade off between the relative importance of performance over Datasets O and T (higher the λ parameter, more closer will the performance of the fine-tuned model be to the model originally trained on Dataset O)

Fig. 9: Model performance while varying λ when fine-tuning only BN layers (top row: (i),(ii)) vs. all layers (bottom row: (iii),(iv)) with EWC[15] on Region 1 and Region 2 using global BN statistics of Region 1 (left) vs. Region 2 (right). Region 1 being the original Region 2 being the target dataset T.

9.2 Other State Of The Art Comparisons

In the results section, We have compared the results of our approach with two existing state-of-the-art techniques: BN fine-tuning [12] and baseline EWC [15]. For additional comparison, we trained models on Institution 1 and

Institution 2 with Zenke et. al.’s code** [38], resulting in a $\kappa = -0.054$ on Institution 1 and $\kappa = 0.573$ on Institution 2 ($p < 0.01$), which shows a lack of effectiveness in mitigating catastrophic forgetting for our real world dataset. Despite rigorous experimentation with the *strength* parameter (denoted by c in the paper [38]), which controls the trade-off between memory retention and learning rate, we failed to observe any improvement in performance on Institution 1.

We did not compare our results with more recent continuous learning techniques, such as OWM [37], which are limited to extracting features from a model trained on the “combined dataset”. This was out of scope for our problem because data between different medical institutions is not allowed to be shared due to patient privacy regulations.

We did, however, compare our results with recent work by [22], in which they show how dropout implicitly produces a gating mechanism in the network that can address catastrophic forgetting. Since their results were only demonstrated for a multi-layer perceptron (MLP), we attempted to adapt their idea for more modern, deeper network architectures, such as ResNet50, in order to accommodate larger and more heterogeneous datasets. We trained a ResNet50 model on two tasks of permuted MNIST described in their paper. The permuted MNIST dataset is generated by shuffling pixels such that the permutation is the same between images of the same task but is different across the tasks [8]. When a standard Resnet50 architecture is trained on task 1 and fine-tuned on task 2, performance on task 1 drops sharply (as shown in Fig. 10 i). Although we were able to replicate their results (perform domain expansion with minimal catastrophic forgetting) using a two layered MLP (with dropout), we were unable to achieve high performance using a Resnet50 architecture. Indeed, [7] previously criticized Permuted MNIST for being an “unrealistic best case scenario for continuous learning”, since in a real world, new dataset would have some concurrence with the old dataset. In order to adapt [22] to our Resnet50 architecture, we inserted dropout with probability of 0.5 before the last layer. This led to significant levels of catastrophic forgetting with accuracy of 0.12 and 0.91 on validation datasets from task 1 and 2 respectively after training on the second task as shown in Fig. 10 ii. This shows that while their dropout technique works well with a shallow network like MLP, the technique fails for complex model architectures^{††}. However, when a ResNet50 architecture was fine-tuned on task 2 using global BN statistics of task 1, catastrophic forgetting was significantly mitigated. Performance on task 1 and task 2 was 0.95 and 0.93 respectively as shown in Fig. 10 iii.

Moreover, in order to test dropouts in the intermediate convolution layers, we additionally experimented with Wide ResNet architecture from [35], given that they could successfully incorporate dropouts inside each residual blocks (after every convolutional layer). We used the same tasks 1 and 2 from Permuted MNIST as used for Resnet50 architecture. We first trained a vanilla Wide ResNet architecture (without dropout) on task 1 and subsequently fine-tuning on task 2 that led to drastic decrease in performance on task 1 as shown in Fig. 10 iv. The model attained accuracies of 0.18 and 0.94 on task 1 and task 2, respectively. Adding dropouts to the wide residual network, similar to [35] with widening factor of 2 and 50 convolution layers (WRN-50-2), resulted in considerable levels of catastrophic forgetting as shown in Fig. 10 v. For a dropout probability of 0.10, accuracies for task 1 and 2 were 0.36 and 0.86 respectively. At higher dropout probabilities, the model was unable to converge for task 2. As shown in Fig. 10 vi, when the model was fine-tuned using global BN statistics of task 1, performance on task 1 and task 2 was 0.88 and 0.95, respectively, hence alleviating catastrophic forgetting. In conclusion, our results show that all the above state-of-the-art techniques are unable to alleviate catastrophic forgetting for modern CNN architectures on commonly used Permuted MNIST dataset even with extensive hyper-parameter tuning. However using the global BN statistics technique, we can easily mitigate this catastrophic forgetting to a large extent.

**<https://github.com/ganguli-lab/pathint>

^{††}One of the possible reasons can be that MLP, due to its fully connected layers, is somewhat blind to the permutations and hence doesn’t forget much from task 1 when trained on task 2.

Fig. 10: Accuracy of a ResNet50 architecture(top) (i) standard training (ii) using dropout probability of 0.5 before the last layer (iii) using batch statistics of task 1 while fine-tuning; a Wide-ResNet architecture(bottom) (iv) standard training (v) using dropout probability of 0.1 (vi) using batch statistics of task 1 while fine-tuning

9.3 Evaluation of Validation Loss

Fig. 11: Evolution of validation loss during All-layer fine-tuning using EWC for different λ values for (a) Institution 1 (left) and Institution 2 (right); (b) Scanner A (left) and Scanner B (right).

We visualize the evolution of the validation loss (Fig. 11) on both the original and target datasets with varying λ values over the course of fine-tuning. As observed, as the λ decreases, the loss for the original dataset O (Scanner A or Institution 1) increases (Fig. 11a, 11b left plot), validating the decrease in performance on O that was observed with decreasing λ values (red or blue curve in Fig 4a, 4b iii). On the other hand, for the target dataset T (Scanner B or Institution 2), we see a smooth decreasing loss curve with decreasing λ , hence validating the increase in performance with decreasing λ (orange or green curve in Fig 4a, 4b iii).

9.4 Reverse Experiments

In order to further substantiate our hypothesis, we additionally repeat all experiments for the mammography dataset in reverse order, *i.e.*, we consider Scanner B or Institution 2 as the original dataset O and Scanner A or Institution 1 as the target dataset T. Visualizing figures (Fig. 12, 13 and 14) we observe similar conclusions for the reverse experiments as observed in the original order.

Fig. 12: Results for domain expansions across digital mammography systems (Scanner B to Scanner A). The different panels correspond to experiments for 1) Baseline models, 2) Only BN-layer fine-tuning with and without EWC[15] and 3) All layer fine-tuning with and without EWC[15].

As for domain expansion from Scanner B (O) to Scanner A (T) is concerned, using global BN statistics of O instead of T, improves performance for both scanner types under the two scenarios: BN fine-tuning (Fig. 12 iv, v) and all layer fine-tuning (Fig. 12 ii, vii). Remarkable, high performance is maintained for Scanner B (O) even without the usage of EWC both for BN Fine-tuning (Fig. 12 v, vi) as well as all layer fine-tuning (Fig. 12 vii, viii). This can also be seen in Fig. 14a left column, where performance stays consistent for Scanner B (O) with increase in λ parameter even though performance of Scanner A (T) drops. A possible reason could be that Scanner A (T), with a large dataset size ($n = 59411$ for Scanner A vs. $n = 32928$ for Scanner B) and some overlap in intensity histogram with Scanner B (O) as shown in Fig. 1a, potentially helps in fine-tuning a model that is robust enough to scans from original scanner type, Scanner B (O) even when $\lambda = 0$. However, note that this is only applicable when we use global BN statistics of O. With global BN statistics of T, we observe a drastic dip in Scanner B's (O) performance even for high λ values (Fig. 14a right column).

Fig. 13: Results for domain expansions across institutions (Institution 2 to Institution 1). The different panels correspond to experiments for 1) Baseline models, 2) Only BN-layer fine-tuning with and without EWC[15] and 3) All layer Fine-tuning with and without EWC[15].

On the other hand, in case of domain expanding from Institution 2 (O) to Institution 1 (T) (Fig. 13 ii), catastrophic forgetting is lower than when domain expanding from Institution 1 (O) to Institution 2 (T) (Fig. 3b ii). Furthermore, at very high values of λ ($\lambda=1e+5$), performance on Institution 1 (T) becomes significantly lower (Fig. 13 ix) as compared to the performance on target domain when domain expanding from Institution 1 (O) to Institution 2 (T) (Fig. 3b ix). Similar to domain expansion from Scanner B to Scanner A, a possible explanation for this finding may be that Institution 2 (O) is smaller and less heterogeneous than Institution 1 (T). Hence while in the first case (Fig. 13 ii), when the model is fine tuning on Institution 1 (T), it possibly learns a more robust model due to dataset size and heterogeneity, that improves the model’s ability to generalize well on Institution 2 (O). In the second case (Fig. 13 ix), having too high a λ essentially allows it to learn only from the Institution 2 dataset, which is smaller and more homogeneous, resulting in poor performance on Institution 1. However, comparing Fig. 14b left and right column, we can clearly ascertain that global BN statistics of O together with EWC is necessary for the model to avoid catastrophic forgetting on Institution 2 (O). Additionally, fine-tuning all layers is needed to achieve high performance on Institution 1 (T).

Fig. 14: Model performance while varying λ when fine-tuning only BN layers (top) vs. all layers (bottom) with EWC[15] on (a) Scanner B and Scanner A using global statistics of Scanner B (left) vs. Scanner A (right) ; (b) Institution 2 and Institution 1 using global statistics of Institution 2 (left) vs. Institution 1 (right). Here, Scanner B and Institution 2 are used as the original dataset O while Scanner A and Institution 1 are used as the target dataset T.

9.5 Statistical Comparisons

Alternate Hypothesis	p-value				
	Scanner A to Scanner B ^{‡‡}	Scanner B to Scanner A ^{§§}	Institution 1 to Institution 2 ^{¶¶}	Institution 2 to Institution 1 ^{***}	Region 1 to Region 2 ^{†††}
Model trained on O, when fine-tuned on T does not forget the information it learned on O	0.0090	0.0090	0.0090	0.0090	0.0090
When fine-tuning only BN layer using the global BN statistics of T, the model does not undergoes change in the performance on O	0.0090	0.0090	0.0090	0.0090	0.0090
The performance on T when fine-tuning only the BN layers is not different than the performance when fine-tuning with all the layers using the global BN statistics of T	0.0090	0.0090	0.0090	0.0090	0.0090
The performance on O when fine-tuning only the BN layers is not different when using the global BN statistics of O compared to using global BN statistics of T	0.0090	0.0090	0.0090	0.0090	0.0283
The performance on T when fine-tuning only the BN layers is not different when using the global BN statistics of O compared to a model trained solely on O	0.0090	0.0090	0.0090	0.0090	0.0283
The performance on O when fine-tuning all layers using the global BN statistics of O is not different than using global BN statistics of T	0.0090	0.0090	0.0090	0.0090	0.9168
The performance on O when fine-tuning only BN layers using EWC is not different than fine-tuning all layer using global BN stats of O	0.0090	0.3472	0.0090	0.0758	0.2505
The performance on T when fine-tuning all layers using EWC is not different than fine-tuning only BN layer using global BN stats of O	0.0090	0.0090	0.0090	0.0090	0.0090
The performance on O when fine-tuning all layers using EWC is not different when using global BN stats of O compared to using global BN stats of T	0.0090	0.0090	0.0090	0.0090	0.0283
The performance on O when fine-tuning only BN layers using EWC is not different when using global BN stats of O compared to using global BN stats of T	0.0090	0.0090	0.0090	0.0090	0.0090

^{‡‡}Dataset O denotes Scanner A and Dataset T indicates Scanner B

^{§§}Dataset O denotes Scanner B and Dataset T indicates Scanner A

^{¶¶}Dataset O denotes Institution 1 and Dataset T indicates Institution 2

^{***}Dataset O denotes Institution 2 and Dataset T indicates Institution 1

^{†††}Dataset O denotes Region 1 and Dataset T indicates Region 2

9.6 Demographics Details

DMIST Dataset		
Components	Composition	Size
Age (median years, IQR)	-	46 (8)
Gender (%)	Female	100
Race (%)	White	80.3%
	Black or African American	13.9%
	Hispanic or Latino	3.7%
	Asian	1.4%
	American Indian or Alaska	0.1%
	Other or Unknown	0.6%
Breast Density (%)	Fatty	11.5%
	Scattered	44.3%
	Heterogeneously dense	38.2%
	Extremely dense	6.2%

MGH Dataset		
Components	Composition	Size
Breast Density (%)	Fatty	7.3%
	Scattered	51.6%
	Heterogeneously dense	36.8%
	Extremely dense	4.3%

IROP Dataset		
Components	Composition	Size
Gestational Age (median weeks, IQR)	-	26.1 (3.0)
Birth Weight (median grams, IQR)	-	817.0 (395)
Presence of Stage	No Stage	71.3%
	Stage	28.7%

Nepal Dataset		
Components	Composition	Size
Gestational Age (median weeks, IQR)	-	33.0 (5.0)
Birth Weight (median grams, IQR)	-	1900 (686)
Presence of Stage	No Stage	95.0%
	Stage	5.0%

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [RSf.pdf](#)