



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

D-MASTER: Mask Annealed Transformer for Unsupervised Domain Adaptation in Breast Cancer Detection from Mammograms

Tajamul Ashraf^{1†}[0000-0002-7372-3782], Krithika Rangarajan^{2*}[0000-0001-5376-6390], Mohit Gambhir^{2*}[0000-0001-7071-0460], Richa Gauba^{2*}[0009-0007-1872-0670], and Chetan Arora¹[0000-0003-0155-0250]

¹ Indian Institute of Technology Delhi, New Delhi 110016, INDIA
tajamul@sit.iitd.ac.in

² AIIMS, Delhi, New Delhi 110029, INDIA

Abstract. We focus on the problem of Unsupervised Domain Adaptation (UDA) for breast cancer detection from mammograms (BCDM) problem. Recent advancements have shown that masked image modeling serves as a robust pretext task for UDA. However, when applied to cross-domain BCDM, these techniques struggle with breast abnormalities such as masses, asymmetries, and micro-calcifications, in part due to the typically much smaller size of region of interest in comparison to natural images. This often results in more false positives per image (FPI) and significant noise in pseudo-labels typically used to bootstrap such techniques. Recognizing these challenges, we introduce a transformer-based Domain-invariant Mask Annealed Student Teacher autoencoder (D-MASTER) framework. D-MASTER adaptively masks and reconstructs multi-scale feature maps, enhancing the model’s ability to capture reliable target domain features. D-MASTER also includes adaptive confidence refinement to filter pseudo-labels, ensuring only high-quality detections are considered. We also provide a bounding box annotated subset of 1000 mammograms from the RSNA Breast Screening Dataset (referred to as RSNA-BSD1K) to support further research in BCDM. We evaluate D-MASTER on multiple BCDM datasets acquired from diverse domains. Experimental results show a significant improvement of 9% and 13% in sensitivity at 0.3 FPI over state-of-the-art UDA techniques on publicly available benchmark INBreast and DDSM datasets respectively. We also report an improvement of 11% and 17% on In-house and RSNA-BSD1K datasets respectively. The source code, pre-trained D-MASTER model, along with RSNA-BSD1K dataset annotations is available at [d-master/webpage](https://github.com/tajamul/d-master).

1 Introduction

Deep neural networks (DNNs) have achieved noteworthy breakthroughs in medical image analysis [9,31,30,1,37,11] and have shown exceptional performance in

[†] Corresponding author

* Radiologists who annotated the RSNA-BSD1K dataset

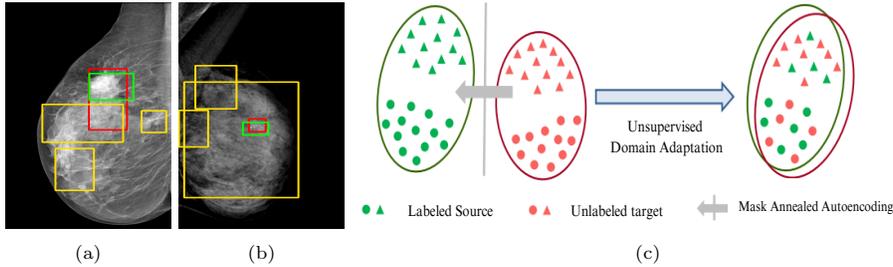


Fig. 1: (a) and (b) depict false positive predictions by current teacher student models in cross-domain BCDM. Red boxes indicate ground truth, yellow boxes show MRT [43] predictions, and green boxes indicate predictions from D-MASTER. As shown in (c), our approach effectively mitigates the domain gap and makes accurate predictions.

specific tasks such as breast cancer detection on mammography. However, they suffer from relatively lower performance when there is a distribution gap between the training data and the deployed environment. This effect is particularly pronounced in medical imaging problems, due to relatively smaller size, and fewer number of annotated datasets, which does not allow a DNN model to capture domain-invariant features. This affects the generalisability of the network across different geographies, different machines, techniques, and protocols of image acquisition. While images in a target population may be available to fine-tune a model, annotations are usually more expensive due to unavailability of a medical expert. Thus, in medical imaging problems, there is a strong need for effective methodologies in unsupervised domain adaptation (UDA).

UDA has been extensively studied in case of natural images, utilizing techniques such as adversarial learning [21], pseudo-label training [6,21,40], image-to-image translation [40], graph reasoning [20], and adaptive mean Teacher training [6], improving domain adaptation efficiency of object detectors. Recently multiple works focused on using Mask autoencoders (MAE) methods in large-scale pretraining for vision models, involving masking parts of input and reconstructing them [12,33]. However, these approaches overlook domain shifts. Alternatively, widespread utilization of Teacher-Student models, wherein a Teacher provides pseudo-labels for target domain (unlabeled images) to supervise a Student model, leads to notable adaptation improvements [40,43]. However, these techniques face the challenge of incorrect predictions and excessive false positives per image due to low-quality pseudo-labels, particularly for medical imaging problems. Pseudo-labels are filtered from the Teacher model’s outputs based on the confidence score threshold. Selecting numerous pseudo-labels with low thresholds leads to inclusion of incorrect predictions, and compromising performance. Conversely, higher thresholds yield a limited number of pseudo-boxes, resulting in sub-optimal supervision. Existing Teacher-Student models often produce pseudo-labels riddled with errors and false positives, as illustrated in Fig. (1a) and Fig. (1b). Though [21,40,43] utilize techniques like adversarial alignment,

weak-strong augmentation, and selective retraining of Student model to minimize the false positives in pseudo-labels, these approaches fail on medical images.

Contributions of this work. We note that screening mammography inherently differs from natural images, with breast abnormalities such as masses, asymmetries, and micro-calcifications, typically much smaller in comparison to the salient objects present in natural images, emphasizing the need for approaches specific to this problem. To address these issues, we make following contributions in this work: **(1)** We introduce D-MASTER, a transformer-based Domain-invariant Mask Annealed Student Teacher Autoencoder Framework for cross-domain breast cancer detection from mammograms (BCDM), integrating a novel mask-annealing technique and adaptive confidence refinement module. Unlike pretraining with mask autoencoders (MAEs) [12], leveraging massive datasets for training and then fine-tuning on smaller datasets, we present a novel learnable masking technique for the MAE branch that generates masks of different complexities, which are reconstructed by the DefDETR [44] encoder and decoder. Our approach, as a self-supervised task on target images, enables the encoder to acquire domain-invariant features and learn better target representations as shown in Fig. (1c). **(2)** In Teacher-Student model, since the pseudo-label noise generated by the Teacher affects the Student model severely, we propose an adaptive confidence refinement module that progressively restricts the confidence metric for pseudo-label filtering. During the initial adaptation phase, soft confidence is applied allowing more pseudo-labels to learn better target representations. Subsequently, as confidence gradually increases, the focus shifts towards enhancing detection accuracy by prioritizing more reliable pseudo-labels. **(3)** We release a bounding box annotated subset of 1000 mammograms from the RSNA Breast Screening Dataset (referred to as RSNA-BSD1K) to support further research in BCDM. **(4)** We setup a new state-of-the-art (SOTA) in detection accuracy for UDA settings. We report a sensitivity of 0.74 on INBreast [25] and 0.51 on DDSM [18] at 0.3 FPI, compared to 0.61 and 0.44 using current SOTA respectively. Significant performance gains are also observed on our in-house and RSNA-BSD1K datasets.

2 Proposed Method

Problem Formulation. For a cross domain BCDM problem in a UDA setting, we have an annotated source dataset $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{N_s}$ with N_s samples, where each sample x represents a mammogram, and $y = (b, c)$ denotes annotation for malignancy, including bounding box b and the corresponding malignant class c . If the sample is benign, there is no corresponding bounding box annotation. Moreover, we have an unannotated target dataset $\mathcal{D}_t = \{x_i^t\}_{i=1}^{N_t}$ with N_t samples. The aim is to improve the performance on target data by training a model solely on source dataset \mathcal{D}_s and target images \mathcal{D}_t (target labels are not available).

2.1 Proposed Architecture

Recent advances [21,40,43] in UDA use an Adaptive Teacher-Student model with adversarial alignment and selective retraining to achieve domain adaptation. Our

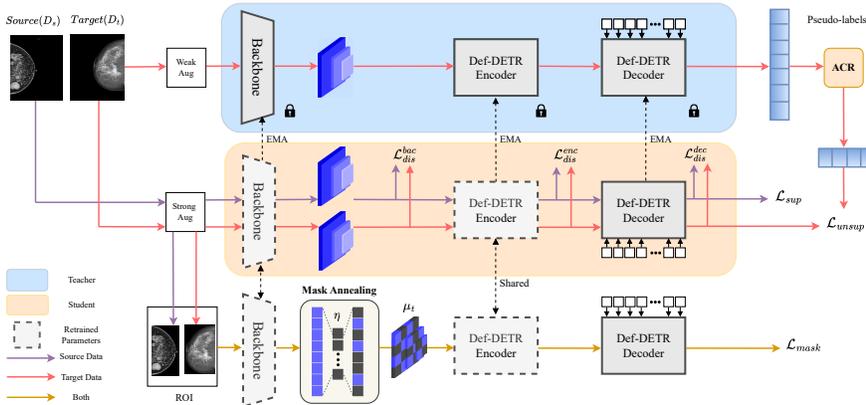


Fig. 2: Architecture of proposed Domain-invariant Mask Annealed Student Teacher Autoencoder (D-MASTER) framework

proposed D-MASTER architecture follows the same style, and is shown in Fig. (2). We utilize *Deformable DETR* [44] (DefDETR) detector, pre-trained on the source domain, \mathcal{D}_s , as the backbone in our architecture.

Teacher-Student model. The model comprises of two main branches: a target-specific Teacher and a cross-domain Student. The Teacher processes weakly augmented images exclusively from target domain (D_t), while the Student handles strongly augmented images from both domains (D_s and D_t). Throughout the training process, the Teacher model generates pseudo-labels for D_t , which are then utilized to train the Student model. Additionally, the Student updates its acquired knowledge to the Teacher through exponential moving average (EMA) after each iteration, $\theta_T \leftarrow \alpha\theta_T + (1 - \alpha)\theta_S$. For source domain, supervised loss \mathcal{L}_{sup} is calculated on x_s using ground truth annotations same as [44], whereas for target domain x_t , we use unsupervised loss, \mathcal{L}_{unsup} , which is the cross-entropy loss with Teacher-generated pseudo-labels.

Discriminators for adversarial alignment. As annotations are solely present for the source images (\mathcal{D}_s), both the Teacher and the Student can be biased towards \mathcal{D}_s during the learning process. To mitigate this [21,43] introduce adversarial learning in the Teacher-Student framework. To achieve adversarial learning, Domain discriminators (D) are used after certain components to predict the domain label of the features, updated by binary cross-entropy loss, \mathcal{L}_{dis} . In D-MASTER, the discriminators are placed after the backbone, DefDETR encoder, and decoder, as shown in Fig. (2). We use standard adversarial loss, $\mathcal{L}_{adv} = \max_{\mathcal{D}_s} \min_{\mathcal{D}_t} \mathcal{L}_{dis}$. Gradient Reverse Layers are used for min-max optimization. The overall objective of the Student (\mathcal{L}_{teach}) is the sum of supervised (\mathcal{L}_{sup}), unsupervised (\mathcal{L}_{unsup}), and adversarial (\mathcal{L}_{adv}) losses.

Selective retraining. Transformer-based models tend to over-fit on target data during such cross-domain training, especially when noisy annotations are included. As the Teacher model undergoes continuous updates via EMA from the

Algorithm 1 Mask Annealing updating process

- 1: **Input:** weight parameters (θ_T), masking ratio (μ_t), step (η), warm restart epoch (T_i), max epochs (T_{\max}), epochs since last restart (T_c), ranges for the steps ($\eta_{\min}^i, \eta_{\max}^i$)
- 2: **Data:** $b_{st} = \{(x_i^s)\}_{i=1}^{N_s} \cup \{x_j^t\}_{j=1}^{N_t}$ \triangleright Batch
- 3: Initialize hyperparameters: $\mu = 0.3$, step $\eta = 0.01$, batch size $b = 16$, max epochs $T_{\max} = 40$, initializing parameters of network θ_T using source trained parameters.
- 4: **for** i in 1 to T_{\max} **do**
- 5: $\eta_t = \eta_{\min}^i + \frac{1}{2}(\eta_{\max}^i - \eta_{\min}^i) \left(1 + \cos\left(\frac{T_c}{T_i} \pi\right)\right)$ *cosine annealing*
- 6: $L = \mathcal{L}_{\text{mask}}(b_{st}, \theta_T)$ \triangleright Computing loss for batch b_{st}
- 7: $\nabla_{\theta} L = \nabla_{\theta} f_t(b_{st}, \theta_T)$ \triangleright Computing gradients
- 8: $\theta'_T = \theta_T - \eta_t \cdot \nabla_{\theta} L$ \triangleright Updating parameters
- 9: $\mu = \mu + \eta_t$ if $L < \bar{L}$, otherwise $\mu = \mu - \eta_t$ \triangleright Updating μ
- 10: $i = i + 1$ \triangleright Incrementing epoch
- 11: **Return** parameters θ'_T, μ' .

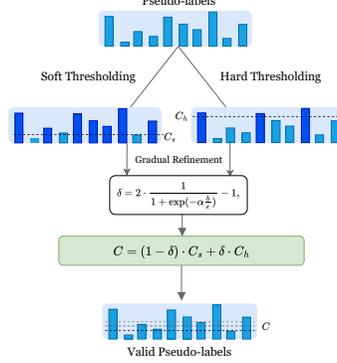


Fig. 3: Mask Annealing Algorithm (**left**) and Adaptive Confidence Refinement (**right**) flowchart depicts the gradual transition of confidence threshold from soft to hard.

Student, it could also be affected and produce incorrect and limited pseudo-labels. To address this, we follow [43], and adopt selective retraining mechanism to help Student jump out of local optimums biased to wrong pseudo-labels. In D-MASTER, the backbone and encoder of the Student are re-initialized with source-trained weights θ_s after certain epochs.

2.2 Mask Annealed Autoencoder

DefDETR encoder in our proposed D-MASTER architecture uses multi-scale feature maps $\{X_i \in \mathbb{R}^{C_i \times H_i \times W_i}\}_{i=1}^K$, where K represents the number of feature map layers. We propose a novel mask annealing technique to mask the feature maps $\{m_i \in \{0, 1\}^{H_i \times W_i}\}_{i=1}^K$ with initial masking ratio μ_t , implying that $\mu_t\%$ of the pixels in the feature map are masked (set to zero). The easy-to-hard mask annealing curriculum, as shown in Fig. (3)(left), is devised to adaptively mask patches to make the reconstruction task easier or difficult. During training, the step η_t is optimized using stochastic gradient descent with warm restarts [23] as:

$$\eta_t = \eta_{\min}^i + \frac{1}{2}(\eta_{\max}^i - \eta_{\min}^i) \left(1 + \cos\left(\frac{T_c}{T_i} \pi\right)\right). \quad (1)$$

Here, η_{\min}^i and η_{\max}^i denote the ranges for the steps, T_c reflects the number of iterations has been completed since the last restart and T_i is the no of iterations when a new warm restart of SGD is to be performed,

Reconstruction. Following the masking of feature maps, deformable attention is applied to further encode them by the DefDETR encoder [44]. In the encoder’s output features, we use a shared mask query q_m to fill the masked portion and

Datasets	Model Name	Venue	R@0.05	R@0.1	R@0.3	R@0.5	R@1.0	R@2.0	Accuracy	F1-score
DDSM [18] to In-house	SFA[36]	MM'21	0.05	0.12	0.19	0.24	0.31	0.39	0.336	0.287
	UMT[6]	CVPR'21	0.02	0.07	0.11	0.16	0.25	0.37	0.308	0.235
	D-Adapt[14]	ICLR'22	0.08	0.13	0.19	0.27	0.36	0.44	0.335	0.215
	AT[21]	CVPR'22	0.13	0.18	0.22	0.40	0.52	0.66	0.654	0.314
	H2FA[38]	CVPR'22	0.05	0.09	0.15	0.21	0.28	0.36	0.291	0.225
	AQT[13]	LJCAI'22	0.08	0.11	0.18	0.26	0.33	0.47	0.378	0.306
	HT[7]	CVPR'23	0.12	0.17	0.32	0.42	0.63	0.75	0.316	0.260
	ConFMIX[24]	WACV'23	0.09	0.13	0.19	0.30	0.46	0.53	0.517	0.218
	CLIPGAP[34]	CVPR'23	0.11	0.14	0.31	0.40	0.55	0.67	0.723	0.348
	MRT[43]	ICCV'23	0.13	0.17	0.34	0.46	0.55	0.70	0.816	0.363
	Ours		0.19	0.26	0.43	0.54	0.68	0.78	0.835	0.392
DDSM [18] to INBreast [25]	SFA [36]	MM'21	0.03	0.04	0.13	0.23	0.35	0.53	0.241	0.338
	UMT[6]	CVPR'21	0.01	0.05	0.09	0.15	0.21	0.25	0.204	0.319
	D-Adapt[14]	ICLR'22	0.11	0.23	0.29	0.45	0.59	0.65	0.731	0.414
	AT[21]	CVPR'22	0.19	0.27	0.38	0.65	0.75	0.78	0.721	0.512
	H2FA[38]	CVPR'22	0.13	0.20	0.32	0.45	0.52	0.58	0.575	0.312
	AQT[13]	LJCAI'22	0.11	0.24	0.37	0.44	0.57	0.64	0.680	0.349
	HT[7]	CVPR'23	0.17	0.31	0.49	0.61	0.69	0.73	0.704	0.362
	ConFMIX[24]	WACV'23	0.19	0.29	0.47	0.58	0.73	0.75	0.737	0.409
	CLIPGAP[34]	CVPR'23	0.15	0.30	0.55	0.61	0.75	0.79	0.712	0.445
	MRT[43]	ICCV'23	0.16	0.31	0.54	0.64	0.72	0.77	0.789	0.489
	Ours		0.25	0.3	0.61	0.70	0.82	0.82	0.808	0.524
In-house to INBreast [25]	SFA [36]	MM'21	0.01	0.03	0.09	0.16	0.19	0.22	0.425	0.271
	UMT[6]	CVPR'21	0.05	0.09	0.11	0.15	0.21	0.26	0.362	0.216
	D-Adapt[14]	ICLR'22	0.04	0.06	0.12	0.18	0.29	0.36	0.471	0.263
	AT[21]	CVPR'22	0.16	0.21	0.28	0.35	0.42	0.48	0.712	0.338
	H2FA[38]	CVPR'22	0.03	0.07	0.13	0.18	0.26	0.35	0.486	0.236
	AQT[13]	LJCAI'22	0.02	0.10	0.17	0.25	0.33	0.39	0.569	.318
	HT[7]	CVPR'23	0.07	0.11	0.18	0.26	0.31	0.38	0.648	0.297
	ConFMIX[24]	WACV'23	0.03	0.08	0.15	0.23	0.28	0.32	0.527	0.285
	CLIPGAP[34]	CVPR'23	0.06	0.13	0.19	0.28	0.36	0.41	0.621	0.310
	MRT[43]	ICCV'23	0.32	0.43	0.52	0.69	0.72	0.78	0.779	0.352
	Ours		0.46	0.53	0.71	0.75	0.79	0.82	0.812	0.395
In-house to RSNA-BSD1K	SFA[36]	MM'21	0.02	0.05	0.09	0.12	0.15	0.19	0.501	0.430
	UMT[6]	CVPR'21	0.04	0.06	0.10	0.13	0.19	0.21	0.528	0.287
	D-Adapt[14]	ICLR'22	0.03	0.07	0.09	0.14	0.18	0.27	0.518	0.329
	AT[21]	CVPR'22	0.18	0.24	0.29	0.36	0.39	0.48	0.698	0.601
	H2FA[38]	CVPR'22	0.02	0.07	0.13	0.19	0.23	0.28	0.632	0.325
	AQT[13]	LJCAI'22	0.07	0.12	0.17	0.25	0.31	0.38	0.574	0.415
	HT[7]	CVPR'23	0.09	0.15	0.19	0.28	0.37	0.45	0.621	0.539
	ConFMIX[24]	WACV'23	0.18	0.21	0.28	0.33	0.41	0.48	0.578	0.413
	CLIPGAP[34]	CVPR'23	0.12	0.18	0.23	0.27	0.36	0.42	0.769	0.586
	MRT[43]	ICCV'23	0.28	0.37	0.44	0.53	0.68	0.71	0.812	0.524
	Ours		0.36	0.41	0.58	0.65	0.70	0.79	0.888	0.653

Table 1: Comparison with state of the art UDA techniques.

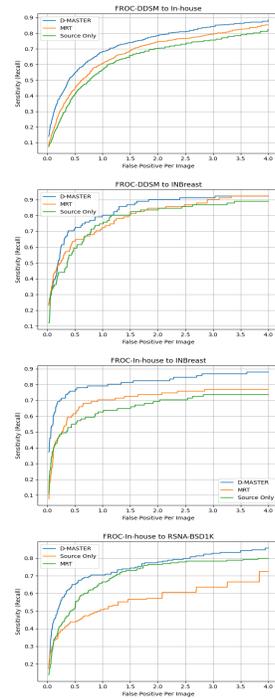


Fig. 4: FROC curves

send it to the MAE decoder D_s to reconstruct the masked portion. Given that the last layer of the feature maps, denoted as X_K , encapsulates all the semantic information [44], we solely reconstruct this layer to expedite faster convergence and lower computational overhead. The decoder's last layer consists of a linear projection, with output channels matching those of X_K . Finally, to supervise the reconstruction loss $\mathcal{L}_{\text{mask}}$, we compute the mean square error between output reconstructed feature maps (\hat{X}_K) and original feature maps (X_K). Hence, the overall objective of the Student model \mathcal{L} is sum of the Teacher based loss ($\mathcal{L}_{\text{teach}}$), and the reconstruction loss ($\mathcal{L}_{\text{mask}}$).

Adaptive confidence refinement (ACR). The correctness of the pseudo-labels is tightly related to confidence used for filtering. In the early stages of learning, confidence tends to be less reliable due to large domain shift. Therefore, we propose a gradual transition of confidence from soft to hard manner (Fig. (3) (right)). To this end, we start with the soft confidence C_s , and as iterations continue, we progressively assign more importance to hard confidence C_h with a shifting weight δ , s.t.: $C = (1 - \delta) \cdot C_s + \delta \cdot C_h$. The weight δ is determined as: $\delta = 2 \cdot \frac{1}{1 + \exp(-\alpha \frac{t}{e})} - 1$, where, t denotes current iteration, e denotes total iterations, and α is a hyperparameter. At each iteration, a pseudo-label from the Teacher is considered valid (used to compute $\mathcal{L}_{\text{unsup}}$) if its confidence exceeds C .

Source	AT [21]	MA	ACR	R@0.025	R@0.05	R@0.1	R@0.3	R@0.5	R@0.9
✓				0.19	0.34	0.43	0.57	0.65	0.69
✓		✓		0.39	0.46	0.53	0.71	0.75	0.78
✓	✓			0.28	0.35	0.42	0.61	0.64	0.69
✓	✓		✓	0.30	0.42	0.51	0.61	0.70	0.71
✓	✓	✓		0.37	0.46	0.59	0.72	0.75	0.75
✓	✓	✓	✓	0.39	0.51	0.53	0.71	0.75	0.79

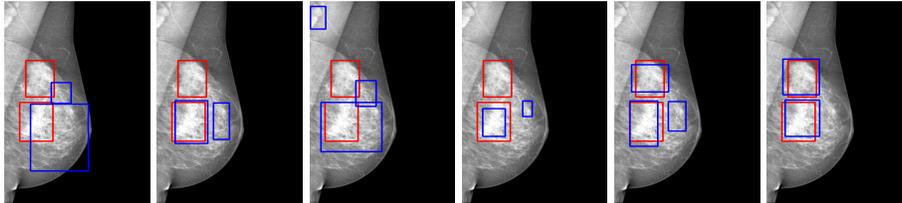


Table 2: Ablation study to understand impact of each proposed module for In-house to INBreast adaptation. “Source” denotes the source-only trained model, “Baseline” the basic teacher-student architecture, “MA” the proposed mask annealing technique, and “ACR” denotes adaptive confidence refinement module. The figures from left to right correspond to qualitative results from row 1 to row 6 respectively. Red boxes denote the ground truth, and blue boxes show the predicted regions.

3 Experiments and Results

Dataset and evaluation methodology. We evaluate our proposed D-MASTER on INBreast [25], DDSM [18], RSNA [2], and in-house datasets. Our in-house diagnostic dataset contains 3,501 mammograms, including both CC and MLO views, with 615 malignancies. The original RSNA dataset [2] consists of 54,706 screening mammograms, containing 1,000 malignancies from 8,000 patients. We curated a subset named RSNA-BSD1K, comprising 1,000 mammograms with 200 malignant cases, annotated at the bounding box level by 2 expert radiologists. Note that unlike single domain detection techniques which use a particular subset of the dataset for training and remaining for testing, our technique does not require any labels from the target dataset. Hence, for our problem, it seems logical to use the whole dataset during training and testing, and not just the respective train or test split. Hence, when reporting results for “Dataset A to Dataset B”, we imply that the model is trained on $\mathcal{D}_s = A$ (whole dataset for the training), and adapted for $\mathcal{D}_t = B$ (whole dataset for UDA and testing).

Evaluation metric. We use Free-Response Receiver Operating Characteristic (FROC) curves [8] for reporting our results. The curves provide a graphical representation of sensitivity/recall values at different false positives per image (FPI). We follow related works in this area [27] and consider a prediction as true positive if the center of the predicted bounding box lies within the ground-truth box.

Implementation details. We employ a 2-layer asymmetric decoder [12] in the MAE with an initial mask ratio of 0.2. Network optimization uses the Adam optimizer [17] with an initial learning rate of 2×10^{-4} , and a batch size of 16.

Table 3: **Left:** Demonstrating superiority of our approach for natural image setting with Sim10K [16] as the source, and Cityscapes [5] as the target dataset. **Right:** Comparing object detectors proposed for natural images with our method.

Detector	Venue	R@0.1	R@0.3	R@0.5	Acc	F1	Detector	Venue	AP(car)
FRCNN [28]	NIPS'15	0.25	0.40	0.45	0.828	0.297	FRCNN [28]	NIPS'15	39.4
YOLO-v5 [10]	ICCVW'21	0.26	0.34	0.36	0.828	0.311	SSAL [26]	NIPS'21	42.5
Cond.DETR [10]	ICCV'21	0.30	0.47	0.51	0.833	0.430	DA-Faster [4]	CVPR'18	41.9
Dab-DETR [22]	ICLR'22	0.43	0.57	0.63	0.843	0.552	MeGACDA [35]	CVPR'21	44.8
Dab DEF [22]	ICLR'22	0.40	0.46	0.50	0.830	0.546	ViSGA [29]	ICCV'21	49.3
DN-DETR [19]	CVPR'22	0.38	0.47	0.58	0.828	0.393	AQT [13]	IJCAI 2022	53.4
DN DEF. [19]	CVPR'22	0.39	0.44	0.47	0.825	0.526	DA-Detr [42]	CVPR'23	54.7
YOLO-V8 [15]	UT'23	0.06	0.08	0.10	0.258	0.346	KTNet [32]	ICCV'21	50.7
FocalNET [39]	NIPS'23	0.56	0.69	0.74	0.865	0.641	ConfMix [24]	CVPR'23	56.3
DINO [41]	ICLR'23	0.35	0.54	0.65	0.820	0.275	PT [3]	PMLR'22	55.1
Ours	-	0.53	0.71	0.75	0.879	0.773	Ours	-	61.9

Data augmentation techniques include random horizontal flips, high-resolution mammograms (4K) for weak augmentation, Gaussian blurring, low-resolution mammograms (1K), and random contrast change for strong augmentations. Detailed information is in supplementary Section 4.

Comparison with SOTA techniques. Table 1 shows the comparative results with other domain adaptation techniques, including those proposed for natural images. Fig. (4) depicts corresponding FROC curves comparison with the nearest competitors only (to avoid clutter). Table S3 in the supplementary show quantitative results for more datasets.

Ablation Study. Our ablation study include three experiments: (1) Table 2 provides the quantitative and qualitative analysis to understand the impact of various proposed components in the proposed D-MASTER architecture. (2) Our core contributions are not specific to medical imaging, and are expected to be useful for natural images as well. Hence, we compare our model with existing domain adaptive object detection (DOAD) techniques when trained on Sim10K [16] data and unsupervised domain adaptation on Cityscapes [5] dataset. Since our model is designed for a single label (presence or absence of breast cancer), hence we perform this experiment only for single label, “car” on the two datasets. Table 3(left) shows the result. (3) To understand the benefit of overall solution with respect to recent SOTA object detection techniques, in Table 3(right), we show the result when these techniques are trained on in-house dataset and directly tested on INBreast without any domain adaptation.

4 Conclusion

Lack of generalisability of Deep Neural Networks for breast cancer detection, on the data obtained across different geographies, with different distributions, acquired on different machines and imaging protocols, is a big barrier towards their clinical adoption. In this paper, we address this problem by proposing a new unsupervised domain adaptation framework, which can help such models adapt to newer sample distributions. Experimental results, carried out with four datasets

from different domains, demonstrated the superior performance over other competing methods. We hope that our work will improve the clinical adoption of automated breast cancer detection models with improved generalization using our technique.

Acknowledgement. We acknowledge and thank the funding support from Ministry of Education, Government of India, Central Project Management Unit, IIT Jammu with sanction number IITJMU/CPMU-AI/2024/0002.

Disclosure of Interests. The authors have no competing interests to declare relevant to this article’s content.

References

1. Cai, M., Li, M., Xiong, Z., Zhao, P., Li, E., Tang, J.: An advanced deep learning framework for video-based diagnosis of asd. In: MICCAI (2022)
2. Carr, C., Kitamura, F., K-Cramer, J., Mongan, J., Andriole, K., V, M., Riopel, M., Ball, R., Dane, S.: Rsnr screening mammography breast cancer detection. 2022
3. Chen, M., Chen, W., Yang, S., Song, J., Wang, X., Zhang, L., Yan, Y., Qi, D., Zhuang, Y., Xie, D., et al.: Learning domain adaptive object detection with probabilistic teacher. arXiv:2206.06293 (2022)
4. Chen, Y., Li, W., Sakaridis, C., Dai, D., Van Gool, L.: Domain adaptive faster r-cnn for object detection in the wild. In: CVPR (2018)
5. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: CVPR (2016)
6. Deng, J., Li, W., Chen, Y., Duan, L.: Unbiased mean teacher for cross-domain object detection. In: CVPR (2021)
7. Deng, J., Xu, D., Li, W., Duan, L.: Harmonious teacher for cross-domain object detection. In: CVPR (2023)
8. Egan, J.P., Greenberg, G.Z., Schulman, A.I.: Operating characteristics, signal detectability, and the method of free response. JASA (1961)
9. Fang, X., Kim, D., Xu, X., Kuang, T., Deng, H.H., Barber, J.C., Lampen, N., Gateno, J., Liebschner, M.A., Xia, J.J., et al.: Deep learning-based facial appearance simulation driven by surgically planned craniomaxillofacial bony movement. In: MICCAI (2022)
10. Ge, Z., Liu, S., Wang, F., Li, Z., Sun, J.: Yolox: Exceeding yolo series in 2021. arXiv:2107.08430 (2021)
11. Gonzalez-Jimenez, A., Lionetti, S., Gottfrois, P., Gröger, F., Pouly, M., Navarini, A.A.: Robust t-loss for medical image segmentation. In: MICCAI (2023)
12. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: CVPR (2022)
13. Huang, W.J., Lu, Y.L., Lin, S.Y., Xie, Y., Lin, Y.Y.: Aqt: Adversarial query transformers for domain adaptive object detection. In: IJCAI (2022)
14. Jiang, J., Chen, B., Wang, J., Long, M.: Decoupled adaptation for cross-domain object detection. In: ICLR (2022)
15. Jocher, G., Chaurasia, A., Qiu, J.: YOLO by Ultralytics (2023), <https://github.com/ultralytics/ultralytics>

16. Johnson-Roberson, M., Barto, C., Mehta, R., Sridhar, S.N., Rosaen, K., Vasudevan, R.: Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks? arXiv:1610.01983 (2016)
17. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
18. Lee, R.S., Gimenez, F., Hoogi, A., Miyake, K.K., Gorovoy, M., Rubin, D.L.: A curated mammography data set for use in computer-aided detection and diagnosis research. *Scientific data* (2017)
19. Li, F., Zhang, H., Liu, S., Guo, J., Ni, L.M., Zhang, L.: Dn-detr: Accelerate detr training by introducing query denoising. In: CVPR (2022)
20. Li, W., Liu, X., Yuan, Y.: Sigma: Semantic-complete graph matching for domain adaptive object detection. In: cvpr (2022)
21. Li, Y.J., Dai, X., Ma, C.Y., Liu, Y.C., Chen, K., Wu, B., He, Z., Kitani, K., Vajda, P.: Cross-domain adaptive teacher for object detection. In: CVPR (2022)
22. Liu, S., Li, F., Zhang, H., Yang, X., Qi, X., Su, H., Zhu, J., Zhang, L.: Dab-detr: Dynamic anchor boxes are better queries for detr. arXiv:2201.12329 (2022)
23. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv:1608.03983 (2016)
24. Mattolin, G., Zanella, L., Ricci, E., Wang, Y.: Confmix: Unsupervised domain adaptation for object detection via confidence-based mixing. In: WACV (2023)
25. Moreira, I.C., Amaral, I. D. I., Cardoso, A., Cardoso, M.J., Cardoso, J.S.: Inbreast: toward a full-field digital mammographic database. *Academic radiology* (2012)
26. Munir, M.A., Khan, M.H., S, M., Ali, M.: Ssal: Synergizing between self-training and adversarial learning for domain adaptive object detection. *NuerIPS* (2021)
27. Rangarajan, K., Aggarwal, P., Gupta, D.K., Dhanakshirur, R., Baby, A., Pal, C., Gupta, A.K., Hari, S., Banerjee, S., Arora, C.: Deep learning for detection of iso-dense, obscure masses in mammographically dense breasts. *European radiology* (2023)
28. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *NuerIPS* (2015)
29. Rezaeianaran, F., Shetty, R., Aljundi, R., Reino, D.O., Zhang, S., Schiele, B.: Seeking similarities over differences: Similarity-based domain alignment for adaptive object detection. In: ICCV (2021)
30. Shu, X., Zhang, L., Wang, Z., Lv, Q., Yi, Z.: Deep neural networks with region-based pooling structures for mammographic image classification. *IEEE TMI* (2020)
31. Sun, Z., Jiang, H., Ma, L., Yu, Z., Xu, H.: Transformer based multi-view network for mammographic image classification. In: MICCAI (2022)
32. Tian, K., Zhang, C., Wang, Y., Xiang, S., Pan, C.: Knowledge mining and transferring for domain adaptive object detection. In: ICCV (2021)
33. Tong, Z., Song, Y., Wang, J., Wang, L.: Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *NuerIPS* (2022)
34. Vidit, V., Engilberge, M., Salzmann, M.: Clip the gap: A single domain generalization approach for object detection. In: CVPR (2023)
35. Vs, V., G, V., Oza, P., S, V.A., P, V.M.: Mega-cda: Memory guided attention for category-aware unsupervised domain adaptive object detection. In: CVPR (2021)
36. Wang, W., Cao, Y., Zhang, J., He, F., Zha, Z.J., Wen, Y., Tao, D.: Exploring sequence feature alignment for domain adaptive detection transformers. In: *ACM Multimedia*. pp. 1730–1738 (2021)
37. Xie, Y., Yin, Y., Li, Q., Wang, Y.: Deep mutual distillation for semi-supervised medical image segmentation. In: MICCAI (2023)

38. Xu, Y., Sun, Y., Y, Z., Miao, J., Y, Y.: H2fa r-cnn: Holistic and hierarchical feature alignment for cross-domain weakly supervised object detection. In: CVPR (2022)
39. Yang, J., Li, C., Dai, X., Gao, J.: Focal modulation networks. NeurIPS (2022)
40. Yu, J., Liu, J., Wei, X., Zhou, H., Nakata, Y., Gudovskiy, D., Okuno, T., Li, J., Keutzer, K., Zhang, S.: Cross-domain object detection with mean-teacher transformer. arXiv:2205.01643 (2022)
41. Zhang, H., Li, F., Liu, S., Zhang, L., Su, H., Zhu, J., Ni, L., Shum, H.Y.: DINO: DETR with improved denoising anchor boxes for end-to-end object detection. ICLR (2023)
42. Zhang, J., Huang, J., Luo, Z., Zhang, G., Zhang, X., Lu, S.: Da-detr: Domain adaptive detection transformer with information fusion. In: CVPR (2023)
43. Zhao, Z., Wei, S., Chen, Q., Li, D., Yang, Y., Peng, Y., Liu, Y.: Masked retraining teacher-student framework for domain adaptive object detection. In: ICCV (2023)
44. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv:2010.04159 (2020)