

Mitigating annotation shift in cancer classification using single image generative models

Marta Buetas Arcas^a, Richard Osuala^{a, b, c}, Karim Lekadir^{a, d}, and Oliver Díaz^{a, e}

^aDepartament de Matemàtiques i Informàtica, Universitat de Barcelona, Barcelona, Spain

^bHelmholtz Center Munich, Munich, Germany

^cTechnical University of Munich, Munich, Germany

^dInstitució Catalana de Recerca i Estudis Avançats (ICREA), Barcelona, Spain

^eComputer Vision Center, Bellaterra, Spain

1. ABSTRACT

Artificial Intelligence (AI) has emerged as a valuable tool for assisting radiologists in breast cancer detection and diagnosis. However, the success of AI applications in this domain is restricted by the quantity and quality of available data, posing challenges due to limited and costly data annotation procedures that often lead to annotation shifts. This study simulates, analyses and mitigates annotation shifts in cancer classification in the breast mammography domain. First, a high-accuracy cancer risk prediction model is developed, which effectively distinguishes benign from malignant lesions. Next, model performance is used to quantify the impact of annotation shift. We uncover a substantial impact of annotation shift on multiclass classification performance particularly for malignant lesions. We thus propose a training data augmentation approach based on single-image generative models for the affected class, requiring as few as four in-domain annotations to considerably mitigate annotation shift, while also addressing dataset imbalance. Lastly, we further increase performance by proposing and validating an ensemble architecture based on multiple models trained under different data augmentation regimes. Our study offers key insights into annotation shift in deep learning breast cancer classification and explores the potential of single-image generative models to overcome domain shift challenges. All code used for this study is made publicly available at <https://github.com/MartaBuetas/EnhancingBreastCancerDiagnosis>.¹

Keywords: Dataset Shift, Image Synthesis, Mammography, Synthetic Data, GANs, Deep Learning

2. INTRODUCTION

Artificial Intelligence (AI) has emerged as a beneficial tool to assist radiologists in breast cancer detection and diagnosis.^{2–5} Both the quantity and quality of available data have a direct impact on the success of these applications. One major challenge is the scarcity of labeled data,^{4,6} largely attributable to the extensive time, effort, and costs associated with acquiring expert annotations. This often limits the training and evaluation of AI models, causing a lack of generalisation and robustness. This issue is further exacerbated by the varying quality of available expert annotations that commonly display high intra- and inter-observer variability.⁷ This can lead to annotation shift,⁸ where a model’s performance decreases at test time if test annotations differ from their training counterparts, for example, in size, accuracy, delineation, lesion boundary and margin definition, annotation protocol or sourcing modality.

To increase classification model robustness against annotation shift, one approach is to generate additional training images that correspond to the desired annotation characteristics of the target domain⁴ (in-domain). To this end, we propose the selection of a single well-annotated in-domain training image to train a generative AI model, which, in turn, learns to synthesize an arbitrary number of variations. These variations are readily usable as additional classification model training images.

Following this approach, the contributions presented in this study are threefold:

- Design and implementation of a high-accuracy malignancy classification model trained to distinguish cancerous from benign breast lesions.
- Identification and quantification of the impact of annotation shift on multiclass classification performance.
- A novel investigation into single-image generative AI models to mitigate annotation shift by requiring as few as only one additional annotation.

3. METHODS AND MATERIALS

The dataset used in this study is the Breast Cancer Digital Repository (BCDR),⁹ an extensive accessible repository that comprises annotated cases of breast cancer patients from the northern region of Portugal. BCDR provides both normal and annotated patient cases, including digitised screen-film mammography and full-field digital mammography images. Each mammogram is accompanied by corresponding clinical data. The dataset consists on 984 mammograms with breast lesions. Among these, 386 are digitised screen-film and 598 digital images from 425 women. Each image is annotated with a binary mask delineated by radiologists to indicate regions-of-interest. Within this dataset, 699 cases are biopsy-proven benign, and 285 are malignant. Additionally, 200 normal mammograms from 48 women are included. To ensure the classifier’s robustness for both formats, both film and digital images are equally included in all experiments.

The classification task in this study is to obtain a pre-biopsy cancer risk prediction of breast lesions, which can assist healthcare professionals in making informed decisions regarding further diagnostic procedures, interventions, and treatment plans. Our multiclass deep-learning classification model distinguishes between (i) healthy tissue, (ii) benign, and (iii) malignant lesions, thereby extending on previous approaches, which focused on general healthy vs non-healthy classification.¹⁰ It works at patch level, classifying regions-of-interest extracted as grayscale patches with pixel dimension of 224x224 stacked across 3 channels. Pixel values were normalised to fall within the range of 0 to 1. As classification model, a ResNet50¹¹ was used, which we initialise with weights pretrained on the ImageNet¹² dataset. To optimise the training process, only the parameters of the last layer were kept trainable. For the multiclass task, there were finally 6147 trainable parameters.

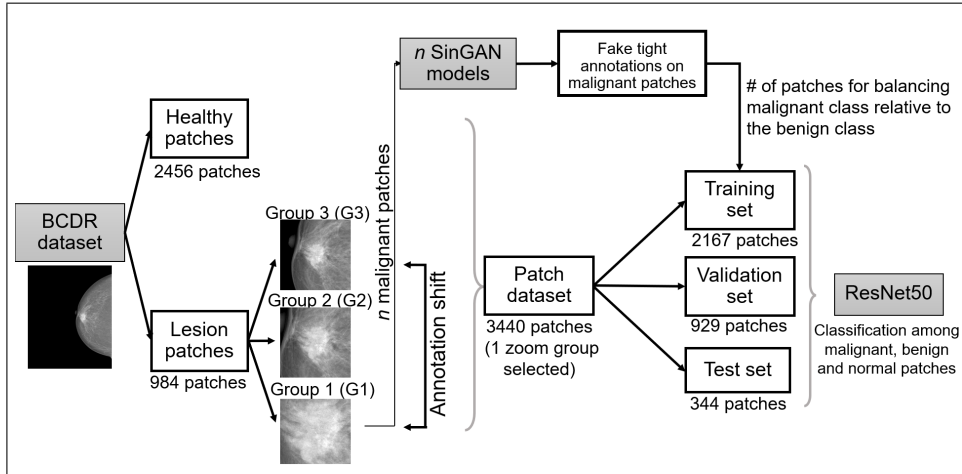


Figure 1. General pipeline of experiments. Patches from healthy and lesion samples are extracted from the BCDR dataset,⁹ with three patches at distinct zoom levels (G1, G2, G3) for each lesion. The dataset is randomly split into training, validation, and test sets using three folds for all experiments, ensuring that images from the same patient are in a single set. To augment the malignant class, patches from G1 are selected to train different SinGAN models individually, each on a distinct selected patch. Once trained, a synthetic dataset is created assembling generated samples from n SinGAN models. This synthetic dataset is incorporated to balance the training dataset.

3.1 Simulating annotation shift

To simulate annotation shift, we extract patches from the lesions from more and less tightly fitting bounding boxes surrounding the lesion, i.e., with different levels of zoom. In practice, an accurate lesion delineation allows to extract a tight lesion bounding box. On the other hand, rectangle lesion annotations (e.g. performed either by human experts or by object detection models^{5,6,13}) contain varying amounts of healthy tissue surrounding the lesions. Therefore, our bounding boxes — extracted based on different zoom levels — simulate varying annotation protocols (annotation shift) and thus allow to measure their influence on classification performance.

Thus, for each lesion, three patches are defined and extracted with different levels of zoom, capturing varying percentages of adjacent healthy tissue. Group 1 (G1) patches correspond to the most accurate bounding box

defined around the original annotated lesion delineation mask. Group 2 (G2) and 3 (G3) capture patches with double (200%) and triple (300%) the height and width of the original bounding box, respectively. Figure 2 exemplifies the aforementioned lesion patches extracted from different zoom levels from a digital mammogram. A total of 3440 such patches are extracted, of which 1406 originated from film-scanned and 2034 from digital mammograms. Among these, 2456 patches are from normal mammograms, 699 from mammograms with benign biopsy results, and 285 from mammograms with malignant biopsy results.

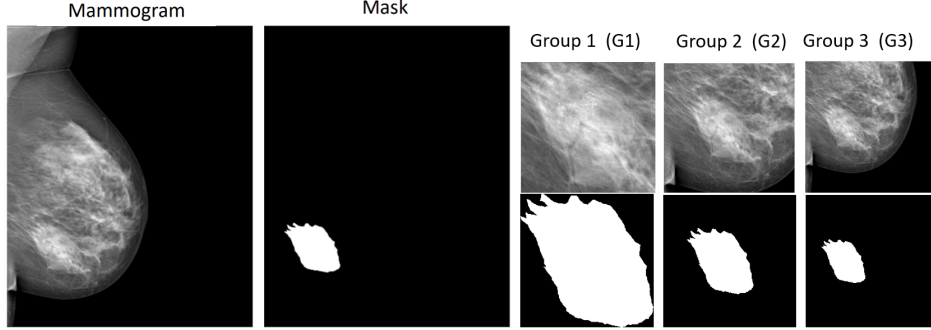


Figure 2. Digital mammogram with a biopsy-proven malignant lesion and its corresponding lesion annotation mask. From the third to the fifth column, the extracted patches are depicted ranging from region-of-interest zoom level group 1 (G1) to group 2 (G2) and 3 (G3) with increasing extend of non-lesion tissue visible on the patch.

3.2 Single image generative model

Given its potential and its successes in medical image synthesis^{4,14–17} and, particularly, lesion region-of-interest generation,¹⁶ we adopt the SinGAN architecture¹⁸ to generate multiple synthetic images from a single training image. SinGAN is a multi-scale generative adversarial network (GAN)¹⁹ containing a generator-discriminator pair in each scale. The input in each such scale s_n consists of a vector z_n drawn from a noise distribution P_z which, if applicable, is added to the image x_{n+1} generated in the previous scale s_{n+1} to generate image x_n . Through such iterative GAN-based upsampling, the generator in the last scale s_0 concludes by outputting a synthetic image x_0 of the size of the original training image.

In summary, SinGAN is a stacked ensemble of \mathcal{N} GANs that operate on different scales of the same image, so each scale n has its own generator (G_n) and discriminator (D_n). Despite the requirement of only a single image as training input, the training process involves sequentially training \mathcal{N} GANs, which is the trade-off for leveraging the capabilities of SinGAN. The training loss for each GAN in each scale incorporates two terms (1): the adversarial loss (\mathcal{L}_{adv}) and the reconstruction loss (\mathcal{L}_{rec}).

$$\min_{G_n} \max_{D_n} \mathcal{L}_{adv}(G_n, D_n) + \alpha \mathcal{L}_{rec}(G_n) \quad (1)$$

The adversarial loss penalises the difference between the distribution of patches in the real image x_n and the distribution of patches in the generated samples \tilde{x}_n . This loss term helps the generator to produce samples that resemble the real data distribution. The Wasserstein GAN with Gradient Penalty (WGAN-GP) loss function,²⁰ is used for training the generator and discriminator. It was observed that this loss function, compared to other GAN variants, enhances training stability.¹⁸ The expression of the WGAN-GP loss function is:

$$\mathcal{L} = \mathbb{E}_{\tilde{\mathbf{x}} \sim \mathbb{P}_g}[f(\tilde{\mathbf{x}})] - \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_r}[f(\mathbf{x})] + \lambda \mathbb{E}_{\tilde{\mathbf{x}} \sim \mathbb{P}_{\hat{\mathbf{x}}}}[(\|\Delta_{\hat{\mathbf{x}}} f(\hat{\mathbf{x}})\|_2 - 1)^2] \quad (2)$$

In Equation 3.2, the first two terms are the original WGAN loss and the right term is the gradient penalty with λ the penalty coefficient.

We implement SinGAN based on the official code repository¹⁸ with adaptations for hyperparameter setup. A scale factor of 0.8 was chosen, which means that the resolution of the image when passing it to the next scale

is reduced by 20%. For each scale, the final model after 30 training epochs was selected. The default values were used for the rest of the hyperparameters. The framework of the SinGAN architecture is depicted in Figure 3, illustrating this hierarchical structure of the model.

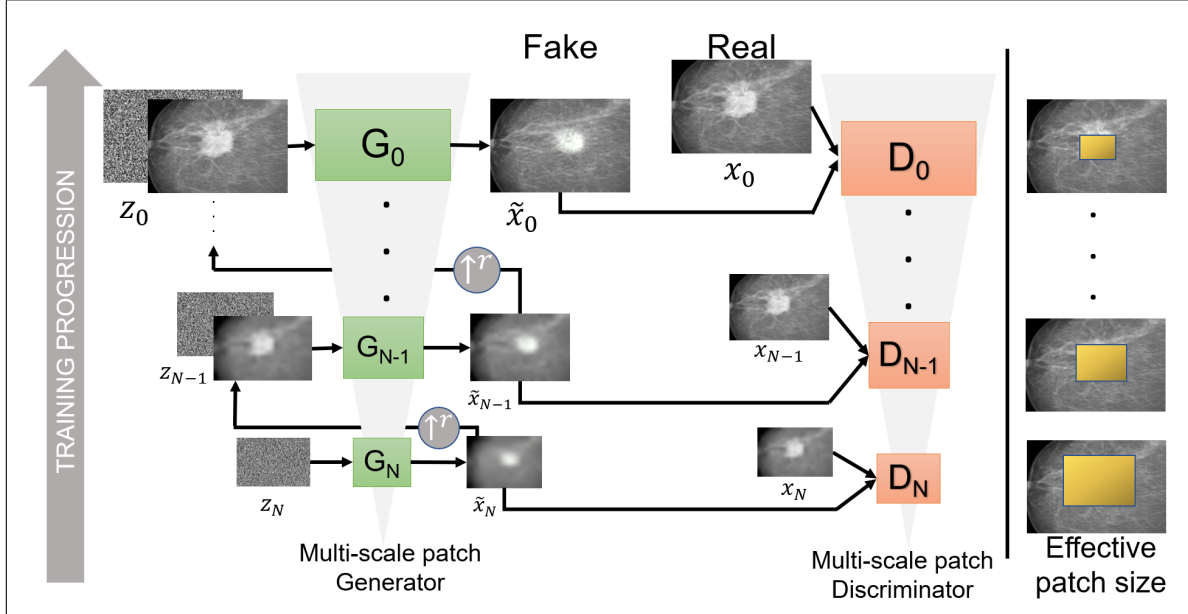


Figure 3. Implemented pipeline of the SinGAN framework adopted from Shaham et al.¹⁸ One GAN operates at each of the n different SinGAN image scales. The training process starts with the coarsest scale and progresses to the finest scale. Each GAN in the hierarchy learns to generate realistic images at its respective scale, capturing both global and local details. At each scale s_n , the image from the previous scale, \tilde{x}_{n+1} , is upsampled and added to the input noise map, z_n . The result is fed into the generator (G_n), whose output is the residual image \tilde{x}_n of scale s_n fed into discriminator D_n during training and passed to the next scale during inference.

3.3 Addressing annotation shift with synthetic data

Apart from annotation shift, our study addresses the challenge of an imbalanced dataset with biopsy-proven malignant breast lesions as minority class impacting classification performance as further elaborated in Section 4. To address this issue, we propose and empirically evaluate the usage of synthetic data generated by single-image generative models to augment the malignant samples. Therefore, our approach is two-fold. Firstly (i), we mitigate annotation shift by generating single image variations of the target annotation distribution controlled via bounding box zoom level. Secondly (ii), we further reduce the effect of dataset imbalance by balancing the dataset with synthetic images corresponding to the minority class.

We train our SinGAN models exclusively on malignant lesion patches from a specific level of zoom (e.g. G1). We note that G1 patches represent the most accurate lesion annotations and, thus, are the most challenging to obtain. For instance, G2 or G3 lesion patches can be retrieved from a G1 annotation, but not vice versa, as the lesion of a G2 or G3 annotation can be (partly) outside the boundaries of a cropped (zoomed-in) G1 patch.

To study the impact of using multiple SinGANs to generate a more diverse synthetic dataset, we individually train multiple SinGAN models, each on a distinct single malignant lesion image. Next, synthetic datasets are assembled, where each SinGAN contributes the same number of samples, which, put together, balance the training dataset of the classification model (i.e. same number of benign and malignant lesions). The synthetic datasets (based on a varying number of SinGANs) is used during classifier training for data augmentation. This approach allows us to systematically study the impact of using different numbers of SinGAN models in generating the synthetic dataset.

Alternatively, as baseline for experiments without synthetic samples, a weighted random sampling technique is applied to avoid class imbalance to bias the classification model. This technique assigns a sampling probability

based on class frequency to each sample in the classifier’s training and validation set. For experiments with synthetic data, weighted random sampling was only applied to the validation set.

4. EXPERIMENTS AND RESULTS

In this study, we initially assessed a binary classification task distinguishing between healthy and lesion-containing patches, yielding a test accuracy of 0.924 ± 0.009 and a test ROC-AUC (Area under the Receiver Operating Characteristic Curve) of 0.971 ± 0.009 , indicating how effectively the model classifies this binary class. Subsequent experiments described in this section extended this setup to multiclass classification, classifying patches either as healthy, benign, malignant. This extended approach enables a comprehensive analysis of detected abnormalities, aiming to predict biopsy outcomes.

Each experiment ran for 100 epochs and the model from the epoch with the best validation loss was selected. Models were evaluated based on a train-validation-test split across three folds, ensuring that each patient was present in only one of the sets. The data was partitioned into 10% for testing (344 samples), 63% for training (2167 samples), and 27% for validation (929 samples). The overall experimental process is represented in Figure 1. The experiments were conducted with consistent hyperparameters to ensure fair comparisons between methods. These included a fixed batch size of 128, utilizing the adaptive moment optimiser (Adam²¹) with default beta parameters ($\beta_1=0.9$ and $\beta_2=0.999$), and employing a learning rate scheduler that gradually decayed the learning rate which started at 10^{-2} . The scheduler had a step of 5 epochs and a gamma value of 0.1. For both the binary classification problem and the multiclass task, a binary cross-entropy loss function was employed. All experiments were run on a NVIDIA RTX 2080 Super 8GB GPU using the PyTorch library.²² Training the classifier for 100 epochs took approximately 3 hours in this setup.

4.1 Analysing model generalisation across all annotation variations

Simulating annotation shift, each experiment involved training the classifier on images from one annotation protocol (represented by a zoom group) and testing on samples from all zoom groups. As upper bound comparison, an additional model was trained using lesion patches from all three zoom groups. Table 2 summarises the empirical results measured using test set classification accuracy and ROC-AUC reported alongside their standard deviation across 3 folds. The ROC-AUC metric facilitates a thorough examination of the separability of the classes across all possible thresholds. For the three-class task, an One-vs-Rest (OvR) strategy for reporting this metric was used.

It is demonstrated that training on lesions from all three groups yields the highest accuracy and ROC-AUC for all classes. This indicates that if there are several annotations available for one sample, it is beneficial for model robustness to use all of these annotations during training.

Observing the ROC-AUC across classes, training on only G2 outperformed exclusive G1 or G3 training, indicating that G2 captures essential features for testing on both extremes. If a range of annotation options exist it can, hence, be optimal to choose the one closest to the average across annotations to generalise best across all annotation options. Conversely, exclusive G1 or G3 training results in poorer performance, emphasising the impact of annotation shift on classifier performance. Overall, the zoom level substantially influenced classifier performance across zoom levels. The model’s dependence on precise annotations is highlighted, which is prone to affect generalisation in real-world computer-aided diagnosis applications, especially in cases where annotation protocols vary.

4.2 Assessing cancer classification robustness under specific annotation shifts

Focusing on the classification class of interest (i.e., ‘Malignant’) to indicate the presence of a cancerous lesion, we assess classifier performance for training and testing on specific varying zoom levels. As expected, the highest accuracy and ROC-AUC was achieved when testing on lesion patches from the same zoom level group as the training data. On the contrary, the lowest accuracy occurred when the classifier trained on group 1 (G1) was tested on group 3 (G3), which also corresponded to the second lowest ROC-AUC result. The lowest ROC-AUC was observed when the classifier trained on group 3 was tested on group 1. With results summarised in Table 2 and in Figure 4, we note that training on the intermediate zoom level of group 2 yields better performance for the

Experiment		Overall accuracy (Mean \pm SD)	ROC-AUC per class (Mean \pm SD)		
Train-Val	Test		Healthy	Malignant	Benign
G1:G2:G3	G1:G2:G3	0.837 \pm 0.045	0.977 \pm 0.009	0.901 \pm 0.037	0.942 \pm 0.021
G1	G1:G2:G3	0.748 \pm 0.044	0.938 \pm 0.022	0.787 \pm 0.046	0.931 \pm 0.027
G2	G1:G2:G3	0.804 \pm 0.035	0.969 \pm 0.008	0.875 \pm 0.046	0.938 \pm 0.017
G3	G1:G2:G3	0.807 \pm 0.035	0.968 \pm 0.008	0.804 \pm 0.055	0.927 \pm 0.021

Table 1. Mean classification accuracy and ROC-AUC for the three classes alongside their standard deviations (SD) across 3 runs. While trained on only on one group of zoom level (G1, G2, G3 in row 2, 3, 4, respectively), the models are tested on samples from all the groups (G1, G1, G3) The upper-bound baseline model in row 1 is trained and tested on all zoom level groups.

malignant class when testing on groups 1 or 3 compared to training on those groups. Overall, it is demonstrated that annotation shift considerably impacts classification performance and further measures to ensure model robustness in such scenarios are needed.

Experiment		Overall accuracy (Mean \pm SD)	ROC-AUC: Malignant class (Mean \pm SD)
Train-Val	Test		
G1	G1	0.929 \pm 0.034	0.948 \pm 0.043
G2	G2	0.899 \pm 0.024	0.930 \pm 0.020
G3	G3	0.887 \pm 0.031	0.928 \pm 0.036
G3	G1	0.865 \pm 0.088	0.677 \pm 0.076
G2	G1	0.877 \pm 0.009	0.840 \pm 0.041
G1	G3	0.780 \pm 0.021	0.709 \pm 0.066
G2	G3	0.835 \pm 0.025	0.893 \pm 0.052

Table 2. Results indicating model robustness under annotation shift. Accuracy across classes and ROC-AUC metrics for the class of interest ('Malignant') at testing time are reported alongside their standard deviation (SD) across three folds. The ROC-AUC metrics for the three classes in these experiments are plotted in Figure 4.

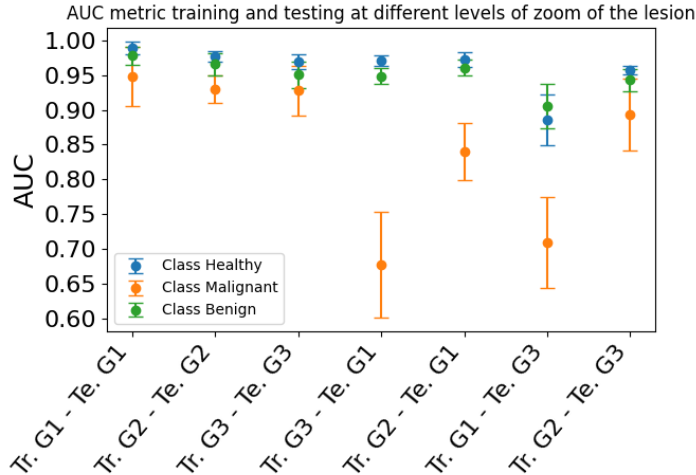


Figure 4. Area under the Receiver Operating Characteristic (ROC) Curve for each class computed using the One-vs-Rest (OvR) strategy for training and testing on specific varying zoom levels. The vertical bars in the plot representing the standard deviation. The results reveal varying model robustness under annotation shift.

4.3 Enhancing model robustness with SinGAN-based data augmentation

In this set of experiments, we investigate data augmentation using SinGAN models to enhance classifier performance under annotation shift for the 'malignant' class. Not only showed this class consistently the lowest classification performance in previous experiments, but it was also heavily impacted by annotation shift, as

shown in Table 2. For each experiment, the classifier was trained on zoom group 3 (G3). Each of the SinGAN models used for generating new data was trained on a different malignant (G1) lesion image. The malignant lesion images used to train each of the SinGAN models, along with two examples generated with each of them, are shown in Table 5.


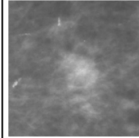
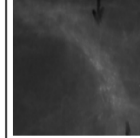
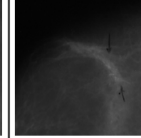
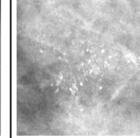
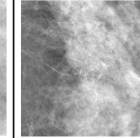
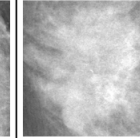
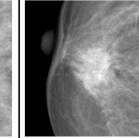
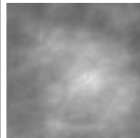
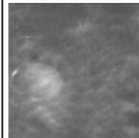
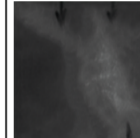
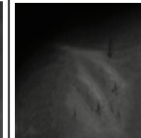
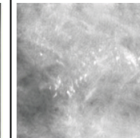
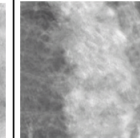
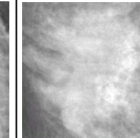
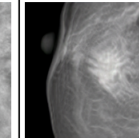
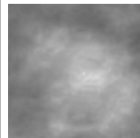
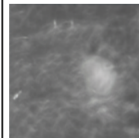
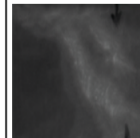
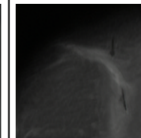
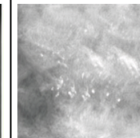
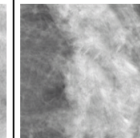
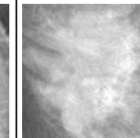
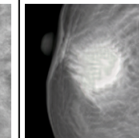
	Film				Digital			
	Sample (i)		Sample (ii)		Sample (i)		Sample (ii)	
	G1	G3	G1	G3	G1	G3	G1	G3
Original								
SinGAN generated								
								

Figure 5. The figure presents the original samples utilized for training each SinGAN model, alongside two synthetic samples generated by the respective trained model. There are two samples for each format (film and digital), and for each sample, its corresponding patch from zoom groups G1 and G3. For augmenting the dataset only samples generated from models trained on a sample from group 1 (G1) were used. This presentation aims to demonstrate the realism achieved by the SinGAN models in generating synthetic content.

To quantitatively assess synthetic data quality, we used the Single Image FID (SiFID) metric¹⁸ measuring the similarity between the real and synthetic images. The SiFID values obtained are presented in Figure 6. To establish a benchmark, these metrics were compared against an upper bound calculated from a dataset of random noise images, resulting in a SiFID of 39.282 ± 0.143 . The SiFID metrics, ranging from 0.14 to 0.49, illustrate that the synthetic images closely match the feature distributions of real images. This alignment indicates a high level of quality and fidelity in the process of image generation.

In each experiment, synthetic data were generated from 1, 2, and 4 different SinGAN models, each trained on distinct single malignant lesion images. The number of total generated samples is the same for all the experiments. This approach allowed for the assessment of the impact of augmented data diversity. The OvR ROC-AUC scores per class for each experiment are presented in Table 3 and Figure 7. The most favorable outcome for the malignant class in terms of ROC-AUC was achieved when generating new data from 4 different SinGAN models. Despite using as few as 1 to 4 lesion images with in-domain annotations, SinGAN data augmentation improves upon the baseline for the underrepresented malignant class.

Additionally, an interesting observation was made: In the experiment utilizing 2 SinGAN models, a performance degradation was observed compared to using either 1 or 4 models. This suggests that the selection of these few lesion images accurately annotated for training SinGAN has a direct impact on the performance enhancement of the classifier. The choice of images used to train the SinGAN model and, therefore, the data employed for data augmentation, likely substantially influences the classifier’s performance. We are currently developing a method to carefully select the most suitable images for data augmentation, aiming to refine the process and enhance classification outcomes.

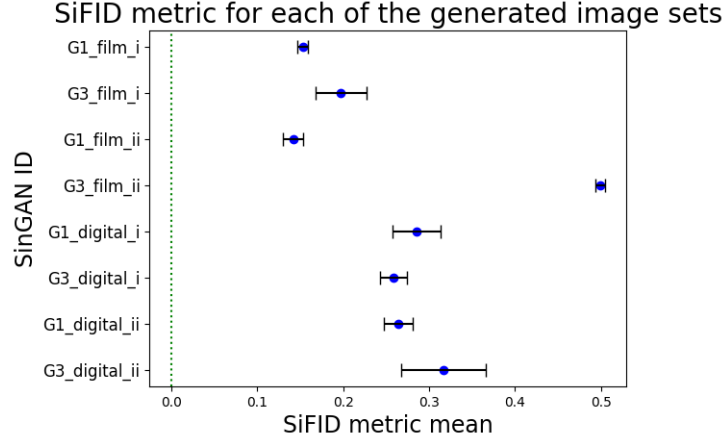


Figure 6. SiFID metric values distribution is presented for datasets generated by SinGAN models trained on different single images. The y-axis represents the ID of the training image used to generate each dataset, indicating the group of zoom level (G1 or G3) and the format (digital or film). The dots represent the mean value, and the horizontal black line represents the standard deviation. For reference, the metric was computed with respect to a dataset of random noise images as an upperbound, which had a SiFID of 39.282 ± 0.143 , so the metrics obtained demonstrate the excellent quality of the generated images. The corresponding original images and the generated ones can be found in Figure 5.

		Number of SinGAN models used for data augmentation			
		No synthetic data generated	1 SinGAN model	2 SinGANs models	4 SinGAN models
Class	Healthy	0.971 ± 0.008	0.968 ± 0.004	0.970 ± 0.001	0.976 ± 0.009
	Malign	0.677 ± 0.076	0.737 ± 0.053	0.679 ± 0.059	0.771 ± 0.045
	Benign	0.949 ± 0.012	0.944 ± 0.007	0.938 ± 0.001	0.953 ± 0.018

Table 3. Table displays experiment results using synthetic data generated from 1, 2, and 4 SinGAN models trained on distinct malignant lesion images from zoom level G1, known for their accurate annotations. Total generated samples remain consistent across experiments. Original classifier training samples are from zoom level G3, augmented with SinGAN-generated data. OvR ROC-AUC scores for classifier testing on G1 samples per class presented, evaluating data diversity impact. SinGAN augmentation enhances underrepresented malignant class; however, using only 2 SinGAN models leads to performance degradation, highlighting the importance of accurately annotated lesion images for SinGAN training. Original malignant lesion images used for SinGAN training are listed in Table 5.

4.4 Combination with traditional data augmentation techniques

Experiments were conducted to compare data augmentation using SinGAN-generated samples with a traditional data augmentation method, specifically single-image (G1) oversampling. The experiments with SinGAN-augmented data involved generating a synthetic dataset from 4 SinGAN models trained on different lesion images from group G1, as it yielded the best results in section 4.3. For consistency, the same four training images utilized for training each SinGAN model were also employed for the oversampling method. Additionally, experiments were repeated with a random resized cropping applied to the test samples. This random resized cropping allows us to additionally measure a classification robustness in the case where the lesion is not at the center of the region-of-interest image. As the assumption of a centered lesion likely does not hold in clinical settings, we analyse whether SinGAN’s capability of varying lesion position observable in Figure 5 is beneficial for robustness in our downstream lesion classification task.

The classifier was trained, as in the experiments of section 4.3, on samples from group zoom G3, while the data augmentation was performed with samples from group zoom G1. Furthermore, the classifier was evaluated on samples from both zoom groups G1, characterized by accurate lesion annotation with minimal healthy tissue inclusion in the patch, and G3, featuring less zoom to the lesion.

In all experiments, an ensemble architecture was utilized. For both augmentation techniques, SinGAN-based and with oversampling, and for the baseline (with no data augmentation), three models were trained with a

AUC metrics for different experiments with SinGAN-based data augmentation

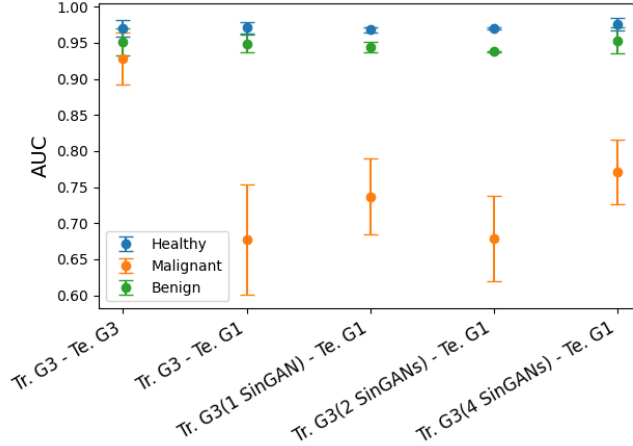


Figure 7. ROC-AUC for each class across experiments where the classifier was trained on zoom group 3 (G3), optionally augmented with synthetic SinGAN generated training images balancing the training dataset. Synthetic images are generated by 1, 2, or 4 SinGAN models trained on single images that had been randomly sampled from the G1 zoom group and result in a performance increase for the malignant class.

train-validation split across three folds, ensuring that each patient was present in only one of the sets. For evaluation, a fixed test set comprising images from patients exclusively allocated to this set was utilized. The ensemble prediction was determined as the mean of the predictions generated by each of the three models within the ensemble. In the case of the ensemble combining SinGAN augmentation and oversampling, a total of six models were configured for the ensemble. It is important to note that the test set was not altered across any of the described experiments. The test samples are sourced from patients who are not represented in the validation or training sets.

While SinGAN-based results show improvement compared to the baseline without data augmentation, they are not substantially superior to our implemented traditional oversampling method indicating further potential of the latter. To this end, we further extend our method to an ensemble architecture combining both SinGAN and oversampling methods. Respective experiments demonstrated that our ensemble further increased classification performance, e.g., compared to applying each method independently. This suggests that SinGAN-based annotation shift mitigation can be further enhanced when used in combination with other methods such as traditional data augmentation techniques.

ROC-AUC (Mean \pm SD) of the Malignant class for data augmentation techniques comparison					
Experiment configuration		Data augmentation method			
Test group	Random resized crop	Baseline	SinGAN	Oversampling	SinGAN and oversampling
G1	No	0.636	0.667	0.708	0.692
G1	Yes	0.404	0.623	0.637	0.666
G3	No	0.917	0.942	0.915	0.942
G3	Yes	0.777	0.816	0.831	0.841

Table 4. ROC-AUC of malignant class comparing SinGAN-generated sample augmentation with traditional single-image (G1) oversampling. Classifier trained on G3 samples, augmented with G1 samples, and evaluated on both G1 and G3 samples. Experiments were also done including random resized crop on test samples to simulate annotation shift impact. An ensemble architecture was employed for all experiments, comprising three models for each augmentation technique and six models for the combined ensemble. Evaluation was conducted on a fixed test set with patients exclusively allocated to this set, ensuring consistent testing conditions across all experiments.

5. DISCUSSION AND CONCLUSION

In this study, we focused on the development of a high-accuracy malignancy classification model trained to distinguish cancerous from benign breast lesions. Our investigation uncovered several critical insights into the challenges and potential solutions in this domain. Firstly, we quantified the impact of annotation shift on classification performance through simulations with varying bounding box tightness representing different annotation protocols. To simulate annotation shift, patches from the lesions were extracted from more and less tightly fitting bounding boxes surrounding the lesion, i.e., with different levels of zoom. Our experiments on sections 4.1 and 4.2 revealed that the zoom level or annotation accuracy substantially influenced classifier performance, emphasizing the model’s dependence on precise annotations, which could affect generalization in real-world computer-aided diagnosis applications where annotation protocols vary.

Moreover, we observed a performance disparity in the classifier’s prediction of malignant cases compared to benign and healthy images. This class consistently exhibited the lowest classification performance and was heavily impacted by annotation shift. To address this, we explored data augmentation techniques, particularly focusing on enhancing performance for the ‘malignant’ class. Our explorations with synthetic data from single-image generative AI models demonstrated the potential to enhance performance with as few as 1 to 4 in-domain annotations. However, we also found that the choice of images used to train the SinGAN model and, therefore, the data employed for data augmentation, significantly influences the classifier’s performance. We are currently working on refining a method for selecting the optimal images for the task of data augmentation, aiming to further improve classification outcomes.

Additionally, we conducted a comparative analysis between SinGAN models and a traditional augmentation technique. Specifically, we employed a traditional oversampling method by augmenting the training dataset with four accurately annotated malignant images. In parallel, four SinGAN models were individually trained on each of these images to generate and augment the data. Our observations revealed that while SinGAN-based augmentation showed improvement in the classifier performance compared to the baseline without augmentation, it did not exhibit substantial superiority over our traditional oversampling approach, indicating the continued potential of the latter. To further explore, we conducted an extended experiment where we combined both methods in an ensemble architecture. This experiment demonstrated that SinGAN models could enhance traditional data augmentation techniques when integrated into an ensemble architecture.

In conclusion, our study reveals multiple insights into breast cancer risk assessment using deep-learning classification models. These findings shed light into the potential of one-shot generative AI models in mitigating annotation shift challenges and collectively contribute to advancing computer-assisted breast cancer diagnosis and detection.

This work is not being, or has been, submitted for publication or presentation elsewhere.

ACKNOWLEDGMENTS

This project has received funding from the European Union’s Horizon Europe and Horizon 2020 research and innovation programme under grant agreement No 101057699 (RadioVal) and No 952103 (EuCanImage), respectively. Also, this work was partially supported by the project FUTURE-ES (PID2021-126724OB-I00) from the Ministry of Science and Innovation of Spain.

REFERENCES

- [1] “Public github repository of this study,” (2024). All code used for this study is made publicly available at <https://github.com/MartaBuetas/EnhancingBreastCancerDiagnosis>.
- [2] Sharma, N., Ng, A. Y., James, J. J., Khara, G., Ambrózay, É., Austin, C. C., Forrai, G., Fox, G., Glocker, B., Heindl, A., et al., “Multi-vendor evaluation of artificial intelligence as an independent reader for double reading in breast cancer screening on 275,900 mammograms,” *BMC cancer* **23**(1), 1–13 (2023).
- [3] Dembrower, K., Crippa, A., Colón, E., Eklund, M., and Strand, F., “Artificial intelligence for breast cancer detection in screening mammography in sweden: a prospective, population-based, paired-reader, non-inferiority study,” *The Lancet Digital Health* **5**(10), e703–e711 (2023).
- [4] Richard Osuala, Kaisar Kushibar, L. G. A. L. Z. S. S. K. B. G. O. D. K. L., “Data synthesis and adversarial networks: A review and meta-analysis in cancer imaging,” in [*Medical Image Analysis 84*, p. 102704],
- [5] Salim, M., Wåhlin, E., Dembrower, K., Azavedo, E., Foukakis, T., Liu, Y., Smith, K., Eklund, M., and Strand, F., “External evaluation of 3 commercial artificial intelligence algorithms for independent assessment of screening mammograms,” *JAMA oncology* **6**(10), 1581–1588 (2020).
- [6] Garrucho, L., Kushibar, K., Osuala, R., Diaz, O., Catanese, A., Del Riego, J., Bobowicz, M., Strand, F., Igual, L., and Lekadir, K., “High-resolution synthesis of high-density breast mammograms: Application to improved fairness in deep learning based mass detection,” *Frontiers in Oncology* **12**, 1044496 (2023).
- [7] Joann G. Elmore, Carolyn K. Wells, C. H. L. D. H. H. and Feinstein., A. R., “Variability in radiologists. interpretations of mammograms,” *New England Journal of Medicine* (PMID: 7969300, pp. 1493– 1499).
- [8] Castro, D.C., W. I. G. B., “Causality matters in medical imaging,” *Nature Communications* **11**(3673) (2020).
- [9] Guevara Lopez, M. A., Posada, N., Moura, D., Pollán, R., Franco-Valiente, J., Ortega, C., Del Solar, M., Díaz-Herrero, G., Ramos, I., Loureiro, J., Fernandes, T., and Araújo, B., “Bcdr: A breast cancer digital repository,” 1065–1066 (01 2012).
- [10] Szafranowska, Z., Osuala, R., Breier, B., Kushibar, K., Lekadir, K., and Diaz, O., “Sharing generative models instead of private data: a simulation study on mammography patch classification,” in [*16th International Workshop on Breast Imaging (IWBI2022)*], **12286**, 169–177, SPIE (2022).
- [11] He, K., Zhang, X., Ren, S., and Sun, J., “Identity mappings in deep residual networks,” in [*Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*], 630–645, Springer (2016).
- [12] Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L., “Imagenet: A large-scale hierarchical image database,” in [*2009 IEEE Conference on Computer Vision and Pattern Recognition*], 248–255 (2009).
- [13] Garrucho, L., Kushibar, K., Jouide, S., Diaz, O., Igual, L., and Lekadir, K., “Domain generalization in deep learning based mass detection in mammography: A large-scale multi-center study,” *Artificial Intelligence in Medicine* **132**, 102386 (2022).
- [14] Xu, I. R., Van Booven, D. J., Goberdhan, S., Breto, A., Porto, J., Alhousseini, M., Algohary, A., Stoyanova, R., Punnen, S., Mahne, A., et al., “Generative adversarial networks can create high quality artificial prostate cancer magnetic resonance images,” *Journal of personalized medicine* **13**(3), 547 (2023).
- [15] Sallé, G., Conze, P.-H., Bert, J., Boussion, N., Visvikis, D., and Jaouen, V., “Cross-modal tumor segmentation using generative blending augmentation and self training,” *arXiv preprint arXiv:2304.01705* (2023).
- [16] Thambawita, V., Salehi, P., Sheshkal, S. A., Hicks, S. A., Hammer, H. L., Parasa, S., Lange, T. d., Halvorsen, P., and Riegler, M. A., “Singan-seg: Synthetic training data generation for medical image segmentation,” *PloS one* **17**(5), e0267976 (2022).
- [17] Osuala, R., Skorupko, G., Lazrak, N., Garrucho, L., García, E., Joshi, S., Jouide, S., Rutherford, M., Prior, F., Kushibar, K., Díaz, O., and Lekadir, K., “medigan: a python library of pretrained generative models for medical image synthesis,” *Journal of Medical Imaging* **10** (feb 2023).
- [18] Shaham, T. R., Dekel, T., and Michaeli, T., “Singan: Learning a generative model from a single natural image,” in [*2019 IEEE/CVF International Conference on Computer Vision (ICCV)*], 4569–4579 (2019).
- [19] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y., “Generative adversarial nets,” *Advances in neural information processing systems* **27** (2014).

- [20] Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C., “Improved training of wasserstein gans,” *Advances in neural information processing systems* **30** (2017).
- [21] Kingma, D. P. and Ba, J., “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980* (2014).
- [22] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S., “Pytorch: An imperative style, high-performance deep learning library,” in [*Advances in Neural Information Processing Systems 32*], Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R., eds., 8024–8035, Curran Associates, Inc. (2019).