




Contents lists available at ScienceDirect

## Pattern Recognition Letters

journal homepage: [www.elsevier.com/locate/patrec](http://www.elsevier.com/locate/patrec)

## Attention-map augmentation for hypercomplex breast cancer classification

Eleonora Lopez<sup>\*</sup>, Filippo Betello, Federico Carmignani, Eleonora Grassucci, Danilo Comminiello

Department of Information Eng., Electronics and Telecom., Sapienza University of Rome, Italy

## ARTICLE INFO

Editor: George Azzopardi

## Keywords:

Attention mechanism  
 Attention maps  
 Hypercomplex neural networks  
 Breast cancer screening  
 Histopathological images

## ABSTRACT

Breast cancer is the most widespread neoplasm among women and early detection of this disease is critical. Deep learning techniques have become of great interest to improve diagnostic performance. However, distinguishing between malignant and benign masses in whole mammograms poses a challenge, as they appear nearly identical to an untrained eye, and the region of interest (ROI) constitutes only a small fraction of the entire image. In this paper, we propose a framework, parameterized hypercomplex attention maps (PHAM), to overcome these problems. Specifically, we deploy an augmentation step based on computing attention maps. Then, the attention maps are used to condition the classification step by constructing a multi-dimensional input comprised of the original breast cancer image and the corresponding attention map. In this step, a parameterized hypercomplex neural network (PHNN) is employed to perform breast cancer classification. The framework offers two main advantages. First, attention maps provide critical information regarding the ROI and allow the neural model to concentrate on it. Second, the hypercomplex architecture has the ability to model local relations between input dimensions thanks to hypercomplex algebra rules, thus properly exploiting the information provided by the attention map. We demonstrate the efficacy of the proposed framework on both mammography images as well as histopathological ones. We surpass attention-based state-of-the-art networks and the real-valued counterpart of our approach. The code of our work is available at <https://github.com/ispamm/AttentionBCS>.

## 1. Introduction

Breast cancer is the most common cancer in women worldwide. While mammography screenings have contributed to reducing mortality, there has been a gradual increase in the incidence rates of this disease since the mid-2000s [1]. Indeed, the mammography exam is by no means a perfect imaging test, characterized by a high rate of false positives and the subsequent need for unnecessary biopsies. Traditional computer-aided detection (CAD) algorithms fail to improve diagnostic performance and lead to high recall rates [2]. On the other hand, deep learning-based CAD systems have been shown to succeed in assisting clinicians during the reading process, reaching a higher diagnostic accuracy [3,4]. This has led to an increased interest in this research area, with a variety of open problems, from reducing the number of false positives [5] to exploiting the multi-view nature of mammography [6,7] and so on [8].

Nonetheless, applying deep learning techniques for this kind of problem still presents challenges. To begin with, the task itself is much more difficult if compared to the classification of natural images. Indeed, discriminating between benign and malignant tumors requires trained and expert radiologists, thus being a far from trivial problem

also for neural networks [9]. In addition, the region of interest (ROI) comprises a tiny portion of the entire mammogram. Because of this, it is hard to detect and, even more importantly, to concentrate on it and learn to pay particular attention to such a small patch [10]. Finally, the fine-grained detail that characterizes high-resolution mammograms is crucial to identifying and correctly diagnosing masses. However, this is lost due to image resizing, a necessary step in order to train a neural model efficiently [6].

The attention mechanism has become a successful technique to handle these kinds of challenges. It allows the network to concentrate on the most critical regions of the input [11]. Typically, it is found inside transformer-like architectures [12,13]. However, new strategies are being developed to endow convolutional neural networks (CNNs) with attention layers [14,15].

In this paper, we introduce a novel approach to exploit the information learned by attention layers to address the aforementioned challenges related to breast cancer screening. We propose a framework that consists of an attention-map augmentation step and a parameterized hypercomplex network as backbone model for breast cancer

<sup>\*</sup> Corresponding author.

E-mail addresses: [eleonora.lopez@uniroma1.it](mailto:eleonora.lopez@uniroma1.it) (E. Lopez), [filippo.betello@uniroma1.it](mailto:filippo.betello@uniroma1.it) (F. Betello), [carmignani.1845479@studenti.uniroma1.it](mailto:carmignani.1845479@studenti.uniroma1.it) (F. Carmignani), [eleonora.grassucci@uniroma1.it](mailto:eleonora.grassucci@uniroma1.it) (E. Grassucci), [daniilo.comminiello@uniroma1.it](mailto:daniilo.comminiello@uniroma1.it) (D. Comminiello).

<https://doi.org/10.1016/j.patrec.2024.04.014>

Received 12 October 2023; Received in revised form 18 February 2024; Accepted 15 April 2024

Available online 18 April 2024

0167-8655/© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. Top rows: attention maps of INbreast obtained from PatchConvNet fine-tuned on CBIS-DDSM. Bottom rows: attention maps of CBIS-DDSM obtained from PatchConvNet fine-tuned on INbreast. The left column comprises mammograms with a malignant finding, while the right presents negative/benign mammograms.

classification. For simplicity, we refer to it as the parameterized hypercomplex attention maps (PHAM) framework. In detail, the attention-map augmentation step consists in computing attention maps for each cancer image with an already existing model, e.g., PatchConvNet [14], such as the ones displayed in Fig. 1. Then we employ them as additional input to the hypercomplex backbone. In this way, we perform a form of conditioning on the attention map during the classification step. Thus, we provide the neural network information regarding the abnormal regions of the breast cancer image, which are most significant for diagnosis. We exploit the new multi-dimensional input through parameterized hypercomplex networks. Indeed, these architectures have the ability to model correlations in multi-dimensional data while also obtaining a more lightweight network [7,16]. Certainly, quaternion and generalized hypercomplex networks have gained much interest in the past few years [17–19]. The success of quaternion neural networks (QNNs) is owed to quaternion algebra that allows to model both global and local relations in input 4D data [20,21]. This, in turn, allows to learn a more powerful representation in the latent space [22]. Thereafter, parameterized hypercomplex neural networks (PHNNs) were introduced in order to bring the advantages of QNNs to general input domains of any dimensionality  $n$  [16,19].

Thus, in this work, we employ parameterized hypercomplex (PH) ResNets as the backbone of our framework. In this way, we are able to capture the correlations between the original breast cancer image and the corresponding attention map thanks to hypercomplex algebra properties. Through an experimental analysis conducted on public benchmark datasets of mammograms, i.e. CBIS-DDSM [23] and INbreast [24], and histopathological microscopic images, i.e., BreakHis [25], we show how the proposed method is able to outperform the real-valued counterpart as well as attention-based state-of-the-art models.

The rest of the paper is organized as follows. Section 2 lays out an overview of the current related works. Section 3 gives a detailed description of the proposed framework and the theory behind hypercomplex networks. Section 4 provides technical details regarding data, training, and experimental results. Finally, conclusions are drawn in Section 5.

## 2. Related works

Prior works adopt several different approaches for the task of breast cancer classification. Owing to the aforementioned problems, many studies focus on the classification of single patches instead of the whole mammogram [26–29]. Certainly, by considering the single patch containing the ROI, there is no need to devise a method for the network to focus on it since it becomes the main object in the image. Moreover, the details would be clearly visible even after image resizing, thus making the classification task much easier compared to whole-mammogram classification [10]. On the other hand, methods that directly process the whole mammography image, usually adopt a pertaining strategy based on either patch-level classification or on natural images in order to alleviate these problems [5–7,10,30–32]. Instead, recent approaches for breast cancer classification from histopathology images involve graph neural networks [33], architectures specifically tailored for histopathology [3] and transfer learning [34].

More recently, with the success of the transformer [11] and vision transformer (ViT) [35] architectures, also the medical imaging community has started to develop architectures based on the former models [12,13,36]. Indeed, these methods aim to focus the attention of the neural model on the ROI, which is often small in the medical scenario, through the self-attention mechanism. However, transformer-based models are characterized by much less inductive biases which instead are inherent to convolutions, i.e., locality and translation equivariance [37,38]. For this reason, many works started to investigate strategies to incorporate self-attention into convolutional neural networks (CNNs). The goal is to maintain its intrinsic inductive biases and additionally gain the advantages of the attention mechanism. One of the first popular works introduced a convolutional bottleneck attention module (CBAM), which can be easily plugged into any CNN to refine feature maps through the inferred channel and spatial attention maps [15]. Then, a more recent study proposed an extension of this module specifically designed for breast cancer classification that is able to exploit cross-view information [39]. The same strategy is also being employed for various medical areas such as EEG decoding [40] and Alzheimer’s disease classification [41]. Finally, Touvron et al. [14] introduced an attention-based aggregation layer to augment any CNN. Furthermore, they proposed a patch-based architecture, PatchConvNet, that allows to obtain higher quality attention maps by keeping the input resolution constant throughout the network. In this paper, we take a different approach and employ the pretrained PatchConvNet to obtain attention maps for mammography images. Then, we exploit them through parameterized hypercomplex layers in order to overcome the problems presented in Section 1.

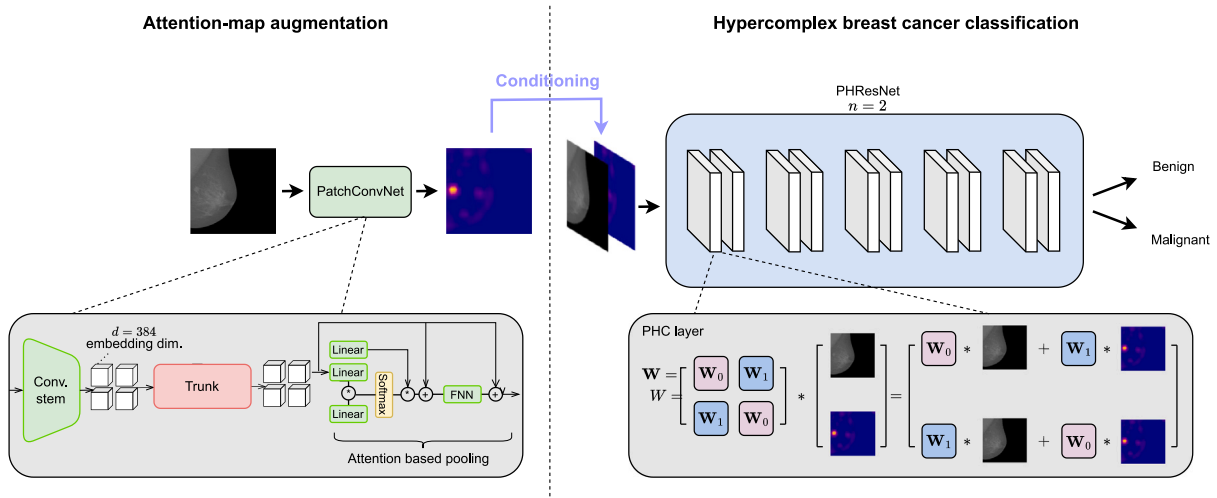
## 3. Proposed method

In this section, we expound on the proposed method, parameterized hypercomplex attention maps (PHAM) framework for breast cancer detection, depicted in Fig. 2. Initially, we present the theoretical background of hypercomplex neural networks, followed by a detailed description of the framework we introduce.

### 3.1. Parameterized hypercomplex models

Quaternion neural networks (QNNs) are models that operate in an extension of complex numbers  $\mathbb{C}$ , namely the quaternion domain  $\mathbb{Q}$ . A quaternion is defined as  $q = q_0 + q_1i + q_2j + q_3k$ , in which  $q_i \in \mathbb{R}$ , with  $i = 0, \dots, 3$ , are the real coefficients and  $i, j, k \in \mathbb{Q}$  are the imaginary units. The product between two imaginary units is not commutative, thus the Hamilton product has been introduced to properly model the multiplication of two quaternions. Thanks to this product, the weight matrix and the input can be encapsulated into a quaternion as

$$\begin{aligned} \mathbf{W} &= \mathbf{W}_0 + \mathbf{W}_1i + \mathbf{W}_2j + \mathbf{W}_3k \\ \mathbf{x} &= \mathbf{x}_0 + \mathbf{x}_1i + \mathbf{x}_2j + \mathbf{x}_3k. \end{aligned} \quad (1)$$



**Fig. 2.** PHAM framework. On the left, the attention-map augmentation step is depicted. Herein, attention maps are computed offline with the fine-tuned PatchConvNet model. Then, they are used to perform a form of conditioning on the hypercomplex model. On the right, a PHResNet with  $n = 2$  for mammography images ( $n = 4$  for histopathology images) is employed as the backbone to perform breast cancer classification. It is able to model relations between breast cancer images and attention maps, as can be seen in the visualization of the PHC layer.

Then the convolution between them is defined as follows:

$$\mathbf{W} * \mathbf{x} = \begin{bmatrix} \mathbf{W}_0 & -\mathbf{W}_1 & -\mathbf{W}_2 & -\mathbf{W}_3 \\ \mathbf{W}_1 & \mathbf{W}_0 & -\mathbf{W}_3 & \mathbf{W}_2 \\ \mathbf{W}_2 & \mathbf{W}_3 & \mathbf{W}_0 & -\mathbf{W}_1 \\ \mathbf{W}_3 & -\mathbf{W}_2 & \mathbf{W}_1 & \mathbf{W}_0 \end{bmatrix} * \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \end{bmatrix}. \quad (2)$$

In this way, the filter submatrices  $\mathbf{W}_i$ ,  $i = 0, \dots, 3$  in Eq. (2) are shared between each dimension of the input  $\mathbf{x}$ . This property brings two main advantages. Firstly, the number of parameters is reduced by  $1/4$ , thus yielding a much lighter model with respect to the real-valued counterpart. Secondly, by sharing weights between input dimensions, the neural model is endowed with the ability to model local relations. Therefore, QNNs, in addition to modeling global relations as standard neural networks are able to do, also capture correlations among channels by treating the input as a unique entity. Instead, real-valued networks assign different weights to different dimensions, thus treating them independently when they are actually correlated. Indeed, this additional information, which networks in the real domain fail to grasp, allows quaternion models to learn a more powerful representation of the data and yield more accurate predictions as a result.

However, in order to extend this approach to inputs of any dimensionality  $nD$ , instead of being limited to 4D inputs, a generalization in the hypercomplex domain  $\mathbb{H}$  was introduced, i.e., parameterized hypercomplex neural networks (PHNNs) [16,19]. In this case, the filter weight matrix  $\mathbf{W}$  is expressed as a parameterized sum of Kronecker products [16]:

$$\mathbf{W} = \sum_{i=0}^n \mathbf{A}_i \otimes \mathbf{F}_i, \quad (3)$$

where  $n$  defines the number domain in which the network operates (e.g.,  $n = 4$  corresponds to the quaternion domain) and can be chosen freely in order to best represent the input data. While matrices  $\mathbf{A}_i$  and  $\mathbf{F}_i$  encode the algebra rules and the weight filters, respectively, both learned during training. With this formulation, the advantages of quaternion models are maintained, thus still leveraging hypercomplex algebra properties to model latent relations among channels, while reducing the number of free parameters by  $1/n$ .

### 3.2. Parameterized hypercomplex attention maps (PHAM)

We design the parameterized hypercomplex attention maps (PHAM) framework as follows. We define an augmentation operation based on the computation of attention maps for each input image. Then we use

them to condition the hypercomplex model during training to improve the classification performance of breast cancer. The relations among the original images and the relative attention maps are exploited through parameterized hypercomplex models thanks to their aforementioned properties.

Indeed, any neural method can be easily defined in the hypercomplex domain [16], therefore in this paper, we deploy parameterized hypercomplex ResNets (PHResNets). In such a manner, the standard residual block  $\mathbf{y} = \mathcal{F}(\mathbf{x}) + \mathbf{x}$ , where  $\mathcal{F}$  is composed of interleaving convolutional layers, batch normalization (BN) and ReLU activation function, becomes:

$$\mathcal{F}(\mathbf{x}) = \text{BN}(\text{PHC}(\text{ReLU}(\text{BN}(\text{PHC}(\mathbf{x}))))). \quad (4)$$

Herein, PHC refers to parameterized hypercomplex convolutions, with the weight matrix defined as in Eq. (3), and  $\mathbf{x}$  is the multidimensional input composed of breast cancer images and the corresponding attention map.

Attention maps are a visualization of what the attention layer has learned during training [14]. We propose to exploit this information in image form, i.e., we augment the dataset with the attention maps, inspired by the similar utilization of heatmaps in recent works [6,42]. More in detail, to compute the attention maps we deploy the recent PatchConvNet [14], since it allows to obtain high-resolution attention maps thanks to its non-hierarchical design. As the network is initially trained on ImageNet, direct utilization for medical images is not feasible. Thus, we first perform a fine-tuning step on a breast cancer dataset and thereafter we apply the fine-tuned model to infer attention maps on other breast cancer databases. Then, we construct the augmented dataset in which each sample is composed of the original mammograms or histopathological images and the corresponding attention maps, considering them as a single multi-dimensional input. In this way, the attention map conditions the training process by emphasizing the most critical portion of the image. This allows the neural network to focus on these crucial areas, thereby enhancing its predictive capabilities and leading to improved performance.

To conclude, when processing images corresponding to mammograms, the input  $\mathbf{x}$  in Eq. (4) has two dimensions. The first represents the mammogram and the second corresponds to the attention map. We set the hyperparameter  $n = 2$  and operate in the complex domain  $\mathbb{C}$  in order to capture relations between them. On the other hand, when considering histopathological images, we set  $n = 4$ , thus operating in the quaternion domain. This choice aligns with the fact that histology

**Table 1**

Results on the test sets of the INbreast and CBIS-DDSM datasets. Attention maps (AM) are obtained from PatchConvNet fine-tuned on CBIS-DDSM (top) and INbreast (bottom), respectively. Results in bold and underlined correspond to the best and second best.

Dataset	Model	Params	PH	AM	AUC
INbreast	ResNet18	11M	✗	✗	0.480 ± 0.070
	ResNet18+CBAM [39]	11M	✗	✗	0.491 ± 0.135
	PatchConvNet [14]	24M	✗	✗	0.806 ± 0.218
	ResNet18	11M	✗	✓	0.840 ± 0.027
	ResNet18+CBAM [15]	11M	✗	✓	0.847 ± 0.018
	PHResNet18 ( $n = 2$ )	5M	✓	✓	<b>0.852 ± 0.037</b>
CBIS-DDSM	ResNet18	11M	✗	✗	0.659 ± 0.008
	ResNet18	11M	✗	✓	0.694 ± 0.034
	PHResNet18 ( $n = 2$ )	5M	✓	✓	<b>0.725 ± 0.027</b>

images are stored in RGB format, comprising 3 channels, plus the additional attention map. According to the aforementioned discussions, by endowing the architecture with PHC layers, we can better leverage the attention maps. Indeed, thanks to hypercomplex algebra properties, the hypercomplex network has the ability to capture local relations between original breast cancer images and the respective maps. In this way, we truly exploit the additional information contained in the attention maps, unlike real-valued networks which fail to model local relations [20].

## 4. Experimental evaluation

### 4.1. Data

We validate our proposed approach on publicly available datasets of mammography, CBIS-DDSM [23] and INbreast [24], and histopathology, BreakHis [25], widely used in literature [3,10].

The Curated Breast Imaging Subset of the Digital Database for Screening Mammography (CBIS-DDSM) [23] consists of scanned film mammography images standardized in the DICOM format. It provides biopsy-proven pathology labels, which can be either benign or malignant. We utilize the official training and testing data splits and obtain the validation split by performing an additional stratified partition of the training set, for a total of 991 images for training, 240 for validation and 361 for testing.

On the other hand, INbreast [24] is a much smaller database consisting of 410 full-filed digital mammography images. Pathological labels are not available but instead, it provides BI-RADS classification from which we extract binary labels by considering categories 1 and 2 as negative, and 4, 5 and 6 as positive, discarding category 3 [10]. We split the dataset in a stratified fashion and take 20% of images for validation and testing, respectively.

The third dataset, the Breast Cancer Histopathological Image Classification (BreakHis) [25], consists of 9109 microscopic images of breast tissue at four different magnifying factors, i.e., 40X, 100X, 200X and 400X. BreakHis is divided into two main groups, that is benign and malignant tumors. We create different sets for training, validation and testing by splitting the dataset patient-wise in order to avoid information-leaking, taking 20% for the two latter sets.

The same preprocessing is applied for all datasets: images are resized to  $384 \times 384$  and standardized. Data augmentation operations are applied for training images only, i.e., random horizontal and vertical flips and a random rotation of degrees taken from  $(-10^\circ, +10^\circ)$ .

### 4.2. Validation metrics

To validate the experiments conducted on datasets of mammogram images we utilize the area under the ROC curve (AUC), as it is the primary metric employed in the literature [5,6,10]. It provides a measure of the predictive ability of a classifier at different probability thresholds, taking into account the trade-off between true positive




Fig. 3. ROC curve corresponding to the best run of models with AM for experiments conducted on INbreast and CBIS-DDSM datasets.




Fig. 4. Confusion matrices of models with AM for experiments on INbreast and CBIS-DDSM.

and false positive rates. For the second set of experiments, the metric adopted in the original paper is the accuracy [25], thus we also utilize it for evaluating the performance of the proposed approach on the BreakHis dataset.

### 4.3. Training and architecture details

ResNet-based architectures are trained with the Adam optimizer [43], which is a variant of Stochastic Gradient Descent [44] with adaptive moment estimation, a learning rate of  $10^{-5}$  and weight decay at  $5 \times 10^{-4}$ . The number of epochs is set to 100 and we early stop the training when the AUC on the validation data does not improve for 20 epochs. Instead, PatchConvNet is fine-tuned with the Lamb optimizer [45], a variant of Adam that provides a strategy for adapting the learning rate in large batch settings. We set the learning rate at  $5 \times 10^{-4}$  and weight decay at  $10^{-2}$ , following the recipe for fine-tuning experiments of the original paper. Moreover, we employ the s60 configuration of PatchConvNet, which consists of an embedding dimension of 384 and 60 repeated blocks in the trunk [14]. Finally, for hypercomplex models, we employ ResNet18 and ResNet50 in the hypercomplex domain for the first and second sets of experiments,



**Table 2**

Results on the test set of the BreakHis dataset at magnification factors 100X, 200X, and 400X. Attention maps (AM) are obtained from PatchConvNet fine-tuned at magnification factor 40X. Results in bold and underlined correspond to the best and second best, respectively.

Magnifying factor	Model	Params	PH	AM	Accuracy (%)
100X	ResNet50	16M	✗	✗	0.696 ± 0.108
	PHResNet50 ( $n = 3$ )	5M	✓	✗	0.767 ± 0.074
	ResNet50	16M	✗	✓	0.809 ± 0.016
	PHResNet50 ( $n = 4$ )	4M	✓	✓	<b>0.821 ± 0.027</b>
200X	ResNet50	16M	✗	✗	0.743 ± 0.020
	PHResNet50 ( $n = 3$ )	5M	✓	✗	0.753 ± 0.057
	ResNet50	16M	✗	✓	0.766 ± 0.007
	PHResNet50 ( $n = 4$ )	4M	✓	✓	<b>0.781 ± 0.016</b>
400X	ResNet50	16M	✗	✗	0.714 ± 0.043
	PHResNet50 ( $n = 3$ )	5M	✓	✗	0.730 ± 0.027
	ResNet50	16M	✗	✓	0.746 ± 0.018
	PHResNet50 ( $n = 4$ )	4M	✓	✓	<b>0.765 ± 0.008</b>

respectively, as BreakHis is a much larger dataset compared to INbreast and CBIS-DDSM.

#### 4.4. Experiments and results

The first set of experiments is conducted on datasets of mammography exams. Specifically, in the first step, we fine-tune the PatchConvNet architecture on CBIS-DDSM utilizing the available ImageNet weights [14]. As a second step, we operate the model at inference time to obtain the attention maps on INbreast. Finally, we train the hypercomplex network, i.e. PHResNet18 ( $n = 2$ ) with the conditioning provided by the attention map, which from here on we denote with AM. We compare the results of our approach against a baseline ResNet18 without AM, state-of-the-art methods, and the real-valued counterpart of the proposed framework, i.e., ResNet18 with AM. Notably, we do not need to test PHResNet18 without AM, as in that case the parameter  $n$  would be set to 1, which is equivalent to its real-valued counterpart, i.e. ResNet18 [16]. The first state-of-the-art method we consider is the cross-view attention module (CvAM) designed for multi-view analysis of mammography, which is integrated inside a ResNet architecture [39]. In order to compare it with our method, we utilize the single-view equivalent of this approach, which is the CBAM module integrated in the same fashion as described in the original paper [39]. The second state-of-the-art network for comparison is PatchConvNet itself, directly fine-tuned on INbreast. The mean AUC over 5 runs is reported in Table 1 together with the number of parameters of each network. Evidently, our approach in both hypercomplex and real domains outperforms both baseline and state-of-the-art methods. This shows that the proposed strategy for exploiting such knowledge is effective, thus resulting in more performant classifiers. Moreover, the proposed framework produces the most discriminant model, i.e. PHResNet with AM, which yields an AUC of 0.852. It is also important to note that it achieves such result with just 5M parameters, that is 1/5 of PatchConvNet. This is thanks to hypercomplex algebra rules which allow modeling local relations between input dimensions, thus grasping correspondences among mammogram and attention map. To conclude, even though the main gain is attained by including attention maps, by introducing hypercomplex algebra the performance is still slightly improved but with a much lighter model. Thereafter, we also perform the same experimental evaluation as described above but with the two datasets switched. Thus, PatchConvNet is fine-tuned on INbreast, then the attention maps for CBIS-DDSM are inferred and used to train the different networks. The results reported in the bottom part of Table 1 support our theory, this time even attaining much more gain from the introduction of hypercomplex algebra, i.e., from 0.694 to 0.725. Furthermore, the ROC curve for experiments in both datasets

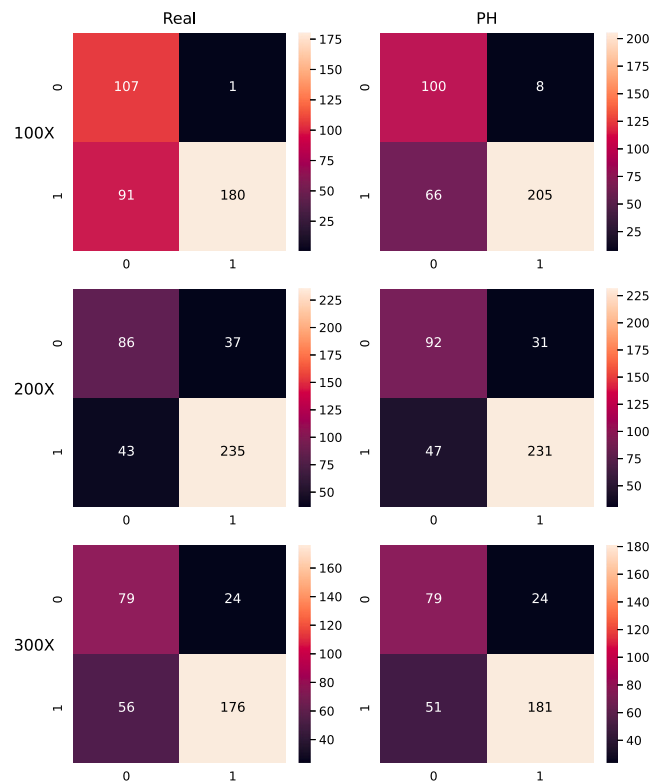


Fig. 5. Confusion matrices of models with AM for experiments on BreakHis.

is depicted in Fig. 3 and is in accordance with the values presented in Table 1.

The second set of experiments is conducted on histopathological microscopic images of tumor tissue at different magnifying factors. In detail, we fine-tune PatchConvNet on images at a magnification factor of 40X and then we utilize it at inference time to compute attention maps for the remaining magnifying factors, i.e. 100X, 200X, and 400X. Finally, we train PHResNet50 with AM on the latter datasets and compare it against a vanilla ResNet50, PHResNet50 with  $n = 3$  (since they are RGB images, i.e. with three channels), and the real-valued respective of our method, i.e., ResNet50 with AM. The results are reported in Table 2, showing the average of the accuracy over 3 runs and the standard deviation. Firstly, as expected, the advantage brought by hypercomplex algebra can be seen in both scenarios with and without AM. In both cases, the mean accuracy is improved and the models are comprised of only 5M and 4M parameters, for  $n = 3$  and  $n = 4$ , respectively. Secondly, the experiments demonstrate how attention maps are generalized for different magnifying factors, thus improving the performance in every case. The hypercomplex network yields the best accuracy for each scenario with a quarter of the parameters with respect to its counterpart in the real domain. Thus, we further demonstrate the efficacy of the proposed framework on microscopic images at different magnification factors of tumor tissue in addition to X-ray mammogram exams. Finally, for both sets of experiments, we illustrate the respective confusion matrices in Figs. 4 and 5. We can see that in every case either the number of false negatives or false positives is reduced by the PH model, aligning with the results of Tables 1 and 2.

To conclude, Table 1 also includes different ablation experiments as the gain from each proposed component, i.e. conditioning on attention maps (AM) and hypercomplex algebra (PH) is shown. In fact, we test a baseline ResNet, then ResNet with AM and finally, we add hypercomplex algebra with PHResNet with AM. Table 2 also presents the experiment with hypercomplex algebra and without AM, i.e. PHResNet with  $n = 3$ .

## 5. Conclusions

In this paper, we have proposed a novel framework to exploit the information learned by the attention mechanism. We build an augmented dataset where each sample comprises the original image and the respective attention map. This new multi-dimensional input is used to condition the training of a PHResNet which handles it as a single unit. Thanks to hypercomplex algebra properties, the PH model has the capacity to capture latent relations between the original breast cancer image and the attention map. In this way, we effectively exploit the additional information provided by conditioning on the attention map regarding the location of the tumor region, shifting the focus of the network on it. We demonstrate the validity of the proposed framework on breast cancer datasets comprising mammography exams and histopathological microscopic images. Our approach outperforms attention-based state-of-the-art architectures and the real-valued counterpart of the proposed technique. In future works, we aim to overcome the limitations of the proposed framework which requires a fine-tuning step of a network such as PatchConvNet due to the domain difference between natural and medical images.

## CRedit authorship contribution statement

**Eleonora Lopez:** Investigation, Methodology, Project administration, Supervision, Validation, Writing – original draft, Writing – review & editing, Software. **Filippo Betello:** Investigation, Methodology, Software, Writing – original draft, Writing – review & editing. **Federico Carmignani:** Software. **Eleonora Grassucci:** Conceptualization, Formal analysis, Validation. **Danilo Comminiello:** Conceptualization, Formal analysis, Funding acquisition, Methodology, Project administration, Supervision.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Data availability

The data used is publicly available.

## Acknowledgments

This work was partially supported by the Italian Ministry of University and Research (MUR) within the PRIN 2022 Program for the project “EXEGETE: Explainable Generative Deep Learning Methods for Medical Signal and Image Processing”, under grant number 2022ENK9LS, CUP C53D23003650001, and in part by the European Union under the National Plan for Complementary Investments to the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, Project PNC 0000001 D3 4 Health - SPOKE 1 - CUP B53C22006120001.

## References

- [1] R.L. Siegel, K.D. Miller, N.S. Wagle, A. Jemal, Cancer statistics, 2023, CA: Cancer J. Clin. 73 (1) (2023) 17–48.
- [2] D. Abdelhafiz, C. Yang, R. Ammar, S. Nabavi, Deep convolutional neural networks for mammography: advances, challenges and applications, BMC Bioinform. 20 (11) (2019) 1–20.
- [3] M.M. Eltoukhy, et al., Classification of multiclass histopathological breast images using residual deep learning, Comput. Intell. Neurosci. 2022 (2022).
- [4] S.M. Naguib, H.M. Hamza, K.M. Hosny, M.K. Saleh, M.A. Kassem, Classification of cervical spine fracture and dislocation using refined pre-trained deep model and saliency map, Diagnostics 13 (7) (2023) 1273.
- [5] N. Wu, et al., Reducing false-positive biopsies using deep neural networks that utilize both local and global image context of screening mammograms, J. Digit. Imaging 34 (2021) 1414–1423.

- [6] N. Wu, et al., Deep neural networks improve radiologists' performance in breast cancer screening, IEEE Trans. Med. Imaging 39 (4) (2020).
- [7] E. Lopez, E. Grassucci, M. Valleriani, D. Comminiello, Hypercomplex neural architectures for multi-view breast cancer classification, 2022, arXiv preprint arXiv:2204.05798.
- [8] G. Zhao, Q. Feng, C. Chen, Z. Zhou, Y. Yu, Diagnose like a radiologist: Hybrid neuro-probabilistic reasoning for attribute-based medical image diagnosis, IEEE Trans. Pattern Anal. Mach. Intell. (2021).
- [9] G. Murtaza, et al., Deep learning-based breast cancer classification through medical imaging modalities: state of the art and research challenges, Artif. Intell. Rev. 53 (2020).
- [10] L. Shen, L. Margolies, J. Rothstein, E. Fluder, R. McBride, W. Sieh, Deep learning to improve breast cancer detection on screening mammography, Sci. Rep. 9 (2019).
- [11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, Adv. Neural Inf. Process. Syst. (NeurIPS) 30 (2017).
- [12] O. Dalmaz, M. Yurt, T. Çukur, ResViT: Residual vision transformers for multimodal medical image synthesis, IEEE Trans. Med. Imaging 41 (10) (2022) 2598–2614.
- [13] Y. Mo, et al., HoVer-Trans: Anatomy-aware HoVer-Transformer for ROI-free breast cancer diagnosis in ultrasound images, IEEE Trans. Med. Imaging (2023).
- [14] H. Touvron, M. Cord, A. El-Nouby, P. Bojanowski, A. Joulin, G. Synnaeve, H. Jégou, Augmenting convolutional networks with attention-based aggregation, 2021, arXiv preprint arXiv:2112.13692.
- [15] S. Woo, J. Park, J.-Y. Lee, I.S. Kweon, CBAM: Convolutional block attention module, in: Proc. of the Europ. Conf. on Computer Vision, ECCV, 2018.
- [16] E. Grassucci, A. Zhang, D. Comminiello, PHNNs: Lightweight neural networks via parameterized hypercomplex convolutions, IEEE Trans. Neural Netw. Learn. Syst. (2022) 1–13.
- [17] D. Comminiello, M. Lella, S. Scardapane, A. Uncini, Quaternion convolutional neural networks for detection and localization of 3D sound events, in: IEEE Int. Conf. on Acoust., Speech and Signal Process., ICASSP, Brighton, UK, 2019, pp. 8533–8537.
- [18] T. Parcollet, M. Morchid, G. Linares, Quaternion convolutional neural networks for heterogeneous image processing, in: IEEE Int. Conf. on Acoust., Speech and Signal Process., ICASSP, Brighton, UK, 2019, pp. 8514–8518.
- [19] A. Zhang, Y. Tay, S. Zhang, A. Chan, A.T. Luu, S.C. Hui, J. Fu, Beyond fully-connected layers with quaternions: Parameterization of hypercomplex multiplications with  $1/n$  parameters, in: Int. Conf. on Machine Learning, ICML, 2021.
- [20] T. Parcollet, M. Morchid, G. Linares, A survey of quaternion neural networks, Artif. Intell. Rev. (2019).
- [21] T. Parcollet, M. Ravanelli, M. Morchid, G. Linares, C. Trabelsi, R. De Mori, Y. Bengio, Quaternion recurrent neural networks, in: Int. Conf. on Learning Representations, ICLR, New Orleans, LA, 2019, pp. 1–19.
- [22] C. Brignone, G. Mancini, E. Grassucci, A. Uncini, D. Comminiello, Efficient sound event localization and detection in the quaternion domain, IEEE Trans. Circuits Syst. II: Express Brief 69 (5) (2022) 2453–2457.
- [23] R. Lee, F. Gimenez, A. Hoogi, M. Kanae, M. Gorovoy, D. Rubin, A curated mammography data set for use in computer-aided detection and diagnosis research, Sci. Data 4 (2017).
- [24] I.C. Moreira, I. Amaral, I. Domingues, A. Cardoso, M.J. Cardoso, J.S. Cardoso, INbreast: Toward a full-field digital mammographic database, Acad. Radiol. 19 (2) (2012) 236–248.
- [25] F.A. Spanhol, L.S. Oliveira, C. Petitjean, L. Heutte, A dataset for breast cancer histopathological image classification, IEEE Trans. Biomed. Eng. 63 (7) (2016) 1455–1462.
- [26] T. Kooi, N. Karssemeijer, Classifying symmetrical differences and temporal change for the detection of malignant masses in mammography using deep neural networks, J. Med. Imag. 4 (4) (2017) 044501.
- [27] T. Kooi, G. Litjens, B. Van Ginneken, A. Gubern-Mérida, C.I. Sánchez, R. Mann, A. den Heeten, N. Karssemeijer, Large scale deep learning for computer aided detection of mammographic lesions, Med. Image Anal. 35 (2017) 303–312.
- [28] R.M. Al-Tam, A.M. Al-Hejri, S.M. Narangale, N.A. Samee, N.F. Mahmoud, M.A. Al-Masni, M.A. Al-Antari, A hybrid workflow of residual convolutional transformer encoder for breast cancer classification using digital X-ray mammograms, Biomedics 10 (11) (2022) 2971.
- [29] A.A. Alhussan, A.A. Abdelhamid, S. Towfek, A. Ibrahim, L. Abualigah, N. Khodadadi, D.S. Khafaga, S. Al-Otaibi, A.E. Ahmed, Classification of breast cancer using transfer learning and advanced al-biruni earth radius optimization, Biomimetics 8 (3) (2023) 270.
- [30] H. Li, D. Chen, W.H. Nailon, M.E. Davies, D.I. Laurenson, Dual convolutional neural networks for breast mass segmentation and diagnosis in mammography, IEEE Trans. Med. Imaging 41 (1) (2021) 3–13.
- [31] J. Wang, Y. Zheng, J. Ma, X. Li, C. Wang, J. Gee, H. Wang, W. Huang, Information bottleneck-based interpretable multitask network for breast cancer classification and segmentation, Med. Image Anal. 83 (2023) 102687.

- [32] K. Jabeen, M.A. Khan, J. Balili, M. Alhaisoni, N.A. Almujaally, H. Alrashidi, U. Tariq, J.-H. Cha, BC2NetRF: breast cancer classification from mammogram images using enhanced deep learning features and equilibrium-jaya controlled regula falsi-based features selection, *Diagnostics* 13 (7) (2023) 1238.
- [33] V. Patel, V. Chaurasia, R. Mahadeva, S.P. Patole, GARL-Net: Graph based adaptive regularized learning deep network for breast cancer classification, *IEEE Access* 11 (2023) 9095–9112.
- [34] F. Hamedani-KarAzmoddehFar, R. Tavakkoli-Moghaddam, A.R. Tajally, S.S. Aria, Breast cancer classification by a new approach to assessing deep neural network-based uncertainty quantification methods, *Biomed. Signal Process. Control* 79 (2023) 104057.
- [35] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, N. Houlsby, An image is worth 16x16 words: Transformers for image recognition at scale, in: *Proc. Int. Conf. on Learning Representations, ICLR*, 2021.
- [36] H.-Y. Zhou, J. Guo, Y. Zhang, X. Han, L. Yu, L. Wang, Y. Yu, nnFormer: volumetric medical image segmentation via a 3D transformer, *IEEE Trans. Image Process.* (2023).
- [37] X. Chen, C.-J. Hsieh, B. Gong, When vision transformers outperform ResNets without pre-training or strong data augmentations, in: *Int. Conf. on Learning Representations, ICLR*, 2021.
- [38] Z. Peng, W. Huang, S. Gu, L. Xie, Y. Wang, J. Jiao, Q. Ye, Conformer: Local features coupling global representations for visual recognition, in: *Proc. IEEE/CVF Int. Conf. on Computer Vision, ICCV*, 2021, pp. 367–376.
- [39] X. Zhao, L. Yu, X. Wang, Cross-view attention network for breast cancer screening from multi-view mammograms, in: *IEEE Int. Conf. on Acoust., Speech and Signal Process.*, ICASSP, 2020, pp. 1050–1054.
- [40] X. Ma, W. Chen, Z. Pei, J. Liu, B. Huang, J. Chen, A temporal dependency learning CNN with attention mechanism for MI-EEG decoding, *IEEE Trans. Neural Syst. Rehabil. Eng.* (2023).
- [41] Z. Hu, et al., Conv-Swinformer: Integration of CNN and shift window attention for Alzheimer’s disease classification, *Comput. Biol. Med.* 164 (2023) 107304.
- [42] R. Singh, P. Zinemanas, X. Serra, J.P. Bello, M. Fuentes, Flowgrad: Using motion for visual sound source localization, in: *IEEE Int. Conf. on Acoust., Speech and Signal Process.*, ICASSP, 2023, pp. 1–5.
- [43] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: Y. Bengio, Y. LeCun (Eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015, URL <http://arXiv.org/abs/1412.6980>.
- [44] S. Ruder, An overview of gradient descent optimization algorithms, 2016, arXiv preprint [arXiv:1609.04747](https://arxiv.org/abs/1609.04747).
- [45] Y. You, et al., Large batch optimization for deep learning: Training BERT in 76 minutes, in: *Int. Conf. Learn. Represent.*, ICLR, 2020.