

Spatial Enhancement by Dehazing for Detection of Microcalcifications with Convolutional Nets

Alessandro Bria¹(✉), Claudio Marrocco¹, Adrian Galdran²,
Aurélio Campilho^{2,3}, Agnese Marchesi¹, Jan-Jurre Mordang⁴,
Nico Karssemeijer⁴, Mario Molinara¹, and Francesco Tortorella¹

¹ DIEI, University of Cassino and Southern Latium, Cassino, FR, Italy
{a.bria,c.marrocco,a.marchesi,m.molinara,tortorella}@unicas.it

² INESC TEC, Institute for Systems and Computer Engineering,
Technology and Science, Porto, Portugal
adrian.galdran@inesctec.pt

³ Faculdade de Engenharia, Universidade do Porto, Porto, Portugal
campilho@fe.up.pt

⁴ DIAG, Radboud University Nijmegen Medical Centre, Nijmegen, The Netherlands
{jan-jurre.mordang,nico.karssemeijer}@radboudumc.nl

Abstract. Microcalcifications are early indicators of breast cancer that appear on mammograms as small bright regions within the breast tissue. To assist screening radiologists in reading mammograms, supervised learning techniques have been found successful to detect microcalcifications automatically. Among them, Convolutional Neural Networks (CNNs) can automatically learn and extract low-level features that capture contrast and spatial information, and use these features to build robust classifiers. Therefore, spatial enhancement that enhances local contrast based on spatial context is expected to positively influence the learning task of the CNN and, as a result, its classification performance. In this work, we propose a novel spatial enhancement technique for microcalcifications based on the removal of haze, an apparently unrelated phenomenon that causes image degradation due to atmospheric absorption and scattering. We tested the influence of dehazing of digital mammograms on the microcalcification detection performance of two CNNs inspired by the popular AlexNet and VGGnet. Experiments were performed on 1,066 mammograms acquired with GE Senographe systems. Statistically significantly better microcalcification detection performance was obtained when dehazing was used as preprocessing. Results of dehazing were superior also to those obtained with Contrast Limited Adaptive Histogram Equalization (CLAHE).

Keywords: Spatial enhancement · Dehazing · Microcalcification detection · Convolutional neural networks · CAD

1 Introduction

Digital mammography is an effective and reliable method for early breast cancer detection, which is fundamental to increment the survival rate and improve

the life quality of the patients [1]. In the last few decades, Computer-Aided Detection (CADe) systems have been proposed to help radiologists in reading screening mammograms. Several studies have shown that CADe systems can improve the performance of individual radiologists [7] in detecting suspicious lesions in mammograms, such as microcalcifications (MCs) and masses. MCs are tiny deposits of calcium that appear on a mammogram as spots of size between 0.1 mm and 1 mm. They are of particular interest since they are usually associated with Ductal Carcinoma In Situ and invasive cancers [29]. Automatic MC detection is often based on supervised learning [4, 5, 14, 21], where powerful binary classifiers are applied to determine whether a MC is present at a pixel location.

Among supervised techniques, Deep Learning approaches have recently acquired great popularity thanks to their outstanding performance in computer vision [17]. In particular, Convolutional Neural Networks (CNNs) have shown to be very effective for classification of image data and also received a large consensus in medical imaging problems [9, 30], including MC detection [24, 32]. A typical CNN architecture is a sequence of feed-forward layers where convolutional filters are interlaced with nonlinear activation functions and pooling. The convolutional layers determine a set of abstract features, whereas the last fully connected layers perform the classification. When training the CNNs, image preprocessing is a fundamental step. Among preprocessing techniques, those including contrast and spatial enhancement have shown to be particularly useful to improve the CNN performance. A preprocessing contrast-extracting layer was firstly used in [6], whereas a local contrast normalization layer was proposed in [12] with the aim of normalizing the responses across all features after each convolutional layer. A layer for brightness normalization was successively introduced in [16] and local plus global contrast normalization was used in [33] to normalize brightness and color variations of RGB images.

Preprocessing techniques are commonly applied in digital mammograms [19, 20, 22, 23], and recently, the effect of contrast enhancement techniques on a CNN has been studied for medical imaging problems [18, 25]. In this work, we propose a novel spatial enhancement method for MCs based on the removal of haze, an apparently unrelated phenomenon usually present in outdoor images that causes image degradation due to atmospheric absorption and scattering. Since CNNs automatically learn and extract low-level features that capture contrast and spatial information, spatial enhancement is expected to positively influence its classification performance. We show that applying an image dehazing approach on mammograms we enhance the contrast of MCs with respect to the surrounding tissue, thus obtaining statistically significantly better MC detection performance when dehazing is used as preprocessing for two different CNNs.

2 Dataset

For this study, we collected a database consisting of 1,066 mammograms acquired with GE Senographe systems (GE, Fairfield, Connecticut,

United States) in Radboud University Medical Center (Nijmegen, The Netherlands). All mammograms were acquired with standard clinical settings at a pixel resolution of 0.1 mm. A total of 7,579 individual MCs were annotated by an experienced reader who marked the center of each microcalcification based on the diagnostic reports. To feed the CNNs, we extracted a dataset of patches of size 12×12 pixels from the mammograms. The patches containing MCs (*positive* samples) were taken by centering the detector window at the groundtruth microcalcification centers, yielding the same number of samples as the individually labeled MCs. The background patches (*negative* samples) were randomly extracted from the remaining regions of the images, totalizing 27,017,503 samples.

3 Spatial Enhancement by Dehazing

3.1 Image Dehazing

The goal of image dehazing is to remove degradation in outdoor images caused by atmospheric absorption and scattering. This physical effect was modelled in [15] as being directly proportional to the distance of the object from the observer, according to the following light propagation law:

$$I(x) = t(x)R(x) + A(1 - t(x)), \quad (1)$$

where x represents a pixel location, $I(x)$ is the intensity captured, $R(x)$ is the radiance in a hypothetical haze-free scene, A is the predominant color of the atmosphere, and $t(x)$ is the transmission of light in the atmosphere. Following [31], the input image can be assumed to have intensities normalized in $[0, 1]$ and be white-balanced, so that the highest intensity in the image is white and A can be approximated by $A \approx (1, 1, 1)$. The haze degradation model simplifies in:

$$I(x) = t(x)R(x) + 1 - t(x), \quad (2)$$

from which, assuming that an estimate of $t(x)$ is available, we can factor the true radiance $R(x)$ by:

$$R(x) = 1 + \frac{I(x) - 1}{t(x)}. \quad (3)$$

Many methods have been proposed for an accurate and robust estimate of $t(x)$. Below we provide a detailed analysis of one of the most successful techniques, namely the Dark Channel Prior [10], that will reveal the direct link between haze removal and spatial enhancement of MCs.

3.2 Dark Channel Prior

The Dark Channel technique is probably the most popular method for image dehazing, partly due to its simplicity. It is based on the observation that most local patches in haze-free images contain some pixels with very low intensity in at least one color channel (the so-called Dark Channel Prior). Thus, given a

pixel x and a local spatial neighborhood $\Omega(x)$ centered on it, the dark channel of the true radiance $R(x)$ contains mostly low values:

$$R^{\text{dark}}(x) = \min_{c \in \{R, G, B\}} \left(\min_{z \in \Omega(x)} R(z) \right) \rightarrow 0. \quad (4)$$

On the other hand, due to the additive degradation component in Eq. 1, the dark channel of the haze-degraded image $I(x)$ can be approximated by:

$$I^{\text{dark}}(x) = \min_{c \in \{R, G, B\}} \left(\min_{z \in \Omega(x)} I(z) \right) \approx A(1 - t(x)). \quad (5)$$

Using this prior in the simplified haze imaging model of Eq. 2, it is possible to directly estimate $t(x)$ as:

$$t(x) \approx 1 - \omega I_{\text{dark}}(x) \quad (6)$$

where $\omega \in (0, 1)$ is a parameter controlling the amount of contrast introduced in the final dehazed image. Due to the implicit local depth constancy made in Eq. 6, the estimated transmission map will usually suffer from a characteristic block artifact, that would lead to halos in the output image unless removed. This can be accomplished with different specialized refining filters, being the typical choice for this task the Guided Filter [11].

3.3 Spatial Enhancement of Microcalcifications

We applied the Dark Channel Prior on mammograms to selectively enhance the contrast of MCs with respect to the surrounding tissue. To show this, let us write $t(x)$ for a grayscale image:

$$t(x) = 1 - \omega \min_{z \in \Omega(x)} I(z) \quad (7)$$

which inserted into Eq. 3, and after simple algebraic manipulations, yields:

$$R(x) = 1 - \frac{1 - I(x)}{1 - \omega \left(\min_{z \in \Omega(x)} I(z) \right)} \quad (8)$$

The key factor in our case is the selection of a neighborhood $\Omega(x)$ slightly bigger than the MC size. In our case, since MCs have typical dimensions well below 1 mm and mammograms have a pixel resolution of 0.1 mm, we chose a squared neighborhood of size 11×11 pixels. This leads us to establish two key observations as explained in the following.

1. *The intensity of MCs is slightly reduced by dehazing.*

If x belongs to a MC, then $\exists \epsilon \in \mathbb{R}^+, \epsilon \ll 1$ so that:

$$I(x) = 1 - \epsilon \quad (9)$$

since MCs have a high intensity in the image. Moreover, in the neighborhood $\Omega(x)$ there will be a background pixel that has the lowest intensity μ within $\Omega(x)$. Then, we can rewrite Eq. 8 as:

$$R(x) = 1 - \frac{\epsilon}{1 - \omega\mu} \tag{10}$$

Let $\Delta = I(x) - \mu$ be the difference between the intensity of the MC pixel under consideration and the lowest-intensity background pixel in the neighborhood $\Omega(x)$. Recalling that $0 < \omega < 1$, and after simple algebraic manipulations, we can bound $R(x)$ as:

$$I(x) > R(x) > \frac{\Delta}{\epsilon + \Delta} \tag{11}$$

Then, combining Eqs. 9 and 11 yields:

$$0 < I(x) - R(x) < 1 - \epsilon - \frac{\Delta}{\epsilon + \Delta} \tag{12}$$

which after simple algebraic manipulations rewrites as:

$$0 < I(x) - R(x) < \frac{\mu}{1 + \frac{\Delta}{\epsilon}} \tag{13}$$

Since $\mu, \epsilon, \Delta > 0$, this provides an upper bound to the difference in intensity between the MC pixels before and after dehazing. Specifically, since μ is the lowest-intensity background pixel in $\Omega(x)$, then $\Delta \gg \epsilon$ and the fraction in Eq. 13 yields a small value. In other words, independently from the choice of ω , the intensity of the MC pixels will only be slightly reduced by dehazing.

- 2. *The intensity of the background around MCs is greatly reduced by dehazing.*
 If x is a background pixel close to a MC so that part of the MC is within $\Omega(x)$, and $\Omega(x)$ is small, then we can approximate the lowest-intensity pixel in $\Omega(x)$ with $I(x)$:

$$\min_{z \in \Omega(x)} I(z) \approx I(x) \tag{14}$$

which combined with Eq. 8 yields:

$$R(x) \approx 1 - \frac{1 - I(x)}{1 - \omega I(x)} \tag{15}$$

This acts as a power-law gamma correction transform controlled by ω (see Fig. 1). Since $I(x)$ is supposed to have mid-low intensity, this transform will greatly darken $I(x)$. The closer ω to 1, the stronger the darkening of $I(x)$.

Following the above observations it is possible to conclude that the contrast between MCs and background tissue is enhanced by dehazing. This can be seen in Fig. 2 where we show a close-up of MCs before and after dehazing with $\omega = 0.9$ and $\Omega(x)$ of size 11×11 pixels. These parameters were fixed at the beginning of our experiments and were not varied afterwards.




Fig. 1. Intensity transformations induced by dehazing on the background surrounding MCs for different values of ω . The closer ω to 1, the stronger the darkening of $I(x)$.




Fig. 2. A mammogram before (left) and after (right) dehazing. In the close-ups, two microcalcifications clusters are shown.

4 Convolutional Neural Networks

A CNN is an ensemble of neurons each featuring several weighted inputs and one output, performing convolution of inputs with weights and transforming the outcome according to a nonlinear activation function. Neurons are arranged in layers and usually share the same weights so as to produce a feature map and reduce the number of parameters. In a typical CNN architecture, convolutional layers are equipped with the Rectified Linear Units (ReLUs) and are intertwined with max-pooling layers. ReLUs apply a nonsaturating activation function $f(x) = \max(0, x)$ which allows the network to easily obtain sparse representations. Max-pooling layers aggregate the outputs of multiple neurons and

Table 1. AlexNet-based architecture

#	Type	Output	Kernel	Stride	Pad
0	Input	$1 \times 12 \times 12$			
1	Conv	$96 \times 12 \times 12$	3×3	1	1
2	LRN	$96 \times 12 \times 12$			
3	Maxpool	$96 \times 6 \times 6$	2×2	2	1
4	Conv	$256 \times 6 \times 6$	3×3	1	1
5	LRN	$256 \times 6 \times 6$			
6	Maxpool	$256 \times 3 \times 3$	2×2	2	1
7	Conv	$384 \times 3 \times 3$	3×3	1	1
8	Conv	$384 \times 3 \times 3$	3×3	1	1
9	Conv	$256 \times 5 \times 5$	3×3	1	1
10	Maxpool	$256 \times 3 \times 3$	3×3	1	1
11	FC	1,024	1×1		
12	Dropout	1,024			
13	FC	512	1×1		
14	Dropout	512			
15	FC	2	1×1		

Table 2. VGGnet-based architecture

#	Type	Output	Kernel	Stride	Pad
0	Input	$1 \times 12 \times 12$			
1	Conv	$32 \times 12 \times 12$	3×3	1	1
2	Conv	$32 \times 12 \times 12$	3×3	1	1
3	Maxpool	$32 \times 6 \times 6$	2×2	2	1
4	Conv	$32 \times 6 \times 6$	3×3	1	1
5	Conv	$32 \times 6 \times 6$	3×3	1	1
6	Maxpool	$32 \times 3 \times 3$	2×2	2	1
7	FC	256	1×1		
8	Dropout	256			
9	FC	256	1×1		
10	Dropout	256			
11	FC	2	1×1		

return the maximum, which results in less training time and lower network complexity. The final decision is made through one or more fully connected layers where each neuron is fed with the outputs of all the neurons of the previous layer. Dropout layers usually follow a fully connected layer to reduce overfitting. The term dropout indicates that, at each training stage, a fixed percentage of outputs coming from the previous layer is ignored in the training of the successive layer.

In this study, we implemented two CNNs inspired by the AlexNet [16] and the VGGnet [28]. The first model is composed by five convolutional and three fully connected layers. Local Response Normalization (LRN) layers follow the first and second convolutional layers, whereas max-pooling layers follow both LRN layers and the last convolutional layer. The ReLU nonlinearity is applied to the output of every convolutional and fully connected layer. The parameters of each layer are reported in Table 1. The second model consists of two stacks of two convolutional layers followed by one max-pooling layer. ReLU is used as activation function for each convolutional layer. The final layers are three fully connected layers. The parameters of each layer are shown in Table 2.

5 Experiments

We applied the two CNNs to the unprocessed mammograms and to the mammograms processed with dehazing and with CLAHE [26], which is a well-known method for spatial enhancement, also applied on mammograms [2]. The parameters of CLAHE were clip limit = 0.01 and block size = 8×8 pixels [34]. We used 2-fold cross validation to train and test the networks. In each cross validation step, the CNN was trained on the 50% of the samples and tested on the other 50%. Before training, positive and negative samples were balanced by means of data augmentation using flipping, rotation, and replication. Each network was trained to minimize the Softmax loss function by means of backpropagation and Mini-Batch Stochastic Gradient Descent, with mini-batches of 32 samples. Standardization was applied to the inputs by mean subtraction and normalization

to unit variance. Weights of each learning layer were initialized using the algorithm of Glorot and Bengio [8]. The learning rate was set to the initial value of 10^{-3} and decreased during training by a factor of 10 every 6 epochs. Momentum and weight decay were set respectively to 0.9 and $5 \cdot 10^{-4}$. The dropout was performed with a probability of 0.5. For the LRN layers of the AlexNet we set the following parameters: $k = 1$, $n = 5$, $\alpha = 10^{-4}$, and $\beta = 0.75$. The learning was stopped after 30 epochs (1 epoch = 844,297 iterations), i.e. when the loss function did not decrease significantly. We used the Caffe framework [13] for the implementation of both networks, and all the experiments were performed on a computer with 2 Intel Xeon e5-2609 processors, 256 GB of RAM and 2 GPU NVIDIA TitanX Pascal.

6 Results

The CNN-based microcalcification detectors without and with the two spatial enhancement methods have been evaluated in terms of Receiver Operating Characteristics (ROC) curve by plotting True Positive Rate (TPR) against False Positive Rate (FPR) for a series of thresholds on the CNN output associated to each sample. Furthermore, the mean sensitivity of the ROC curve in the specificity range on a logarithmic scale was calculated and compared. The mean sensitivity is defined as [24]:

$$\bar{S}(a, b) = \frac{1}{\ln(b) - \ln(a)} \int_a^b \frac{s(f)}{f} df \tag{16}$$

where a and b are the lower and upper bound of the false positive fraction and were set, respectively, to 10^{-6} and 10^{-1} and $s(f)$ is the sensitivity at the false positive fraction f . Statistical comparisons were performed by means of bootstrapping [27] as in [3]. On the test set, average ROC curves were calculated over 1,000 bootstraps, and are reported in Fig. 3. Additionally, the mean sensitivity was calculated for each bootstrap and p -values were computed for

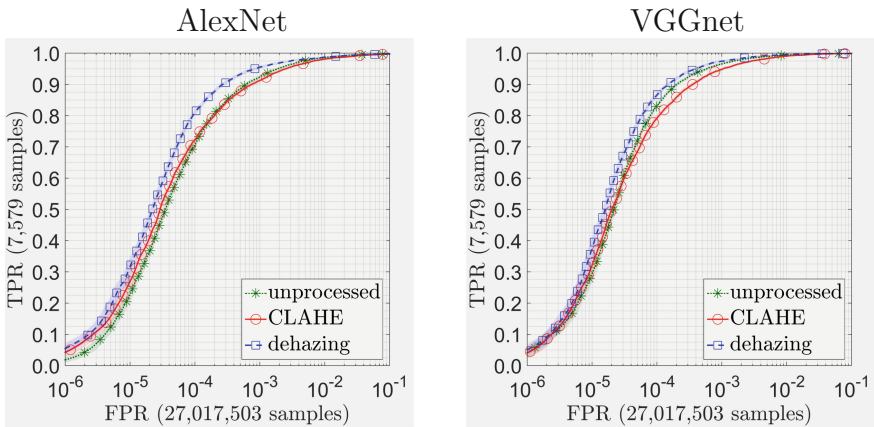


Fig. 3. ROC curves of the CNN detectors averaged from 1,000 bootstrap iterations.

Table 3. Comparative results of mean sensitivity \bar{S} in the FPR range $[10^{-6}, 10^{-1}]$ for different methods (UN = unprocessed, DH = dehazing, CL = CLAHE).

CNN	\bar{S}_{UN}	\bar{S}_{CL}	\bar{S}_{DH}	$\bar{S}_{CL} - \bar{S}_{UN}$	$\bar{S}_{DH} - \bar{S}_{UN}$	$\bar{S}_{DH} - \bar{S}_{CL}$
<i>AlexNet</i>	68.47	69.64	73.29	+1.17 ($p < 0.001$)	+4.82 ($p < 0.001$)	+3.65 ($p < 0.001$)
<i>VGGnet</i>	73.88	73.07	76.26	-0.81 ($p < 0.001$)	+2.38 ($p < 0.001$)	+3.19 ($p < 0.001$)

testing significance. The statistical significance level was chosen as $\alpha = 0.05$ but, due to the number of comparisons $m = 3$, we applied the Bonferroni correction, so that performance differences were considered statistically significant if $p < 0.017$. Results are reported in Table 3. The mean sensitivities obtained on unprocessed images were 68.47 and 73.88 for the AlexNet- and the VGGnet-based CNNs, respectively. Results with dehazing were statistically significantly better than on unprocessed images (+4.82 with AlexNet and +2.38 with VGGnet) and also superior to those of CLAHE (+3.65 with AlexNet and +3.19 with VGGnet).

7 Conclusions

In the present study, we have established a novel connection between the problem of spatial enhancement of MCs in mammograms and the apparently unrelated problem of haze removal in outdoor images. We have shown that the performance of MC detection with CNNs can be greatly improved if mammograms are preprocessed with a dehazing technique to enhance local contrast of MCs. Indeed, it is known that the first layers of a CNN automatically learn and extract low-level features. We can suppose that improving the local contrast of MCs is beneficial for these layers that capture contrast and spatial information in the salient regions. Consequently, this may positively influence the learning task of the subsequent layers that are aimed at capturing more complex features.

Future works will be focused on analyzing the impact of image dehazing also on other CNN architectures and MC detectors. In addition, we will experiment other existing dehazing methods. If successful, this would lead to an entire new family of simple and effective alternative spatial enhancement methods for MCs.

Acknowledgment. The authors from University of Cassino gratefully acknowledge the support of NVIDIA Corporation for the donation of the Titan X Pascal GPUs.

References

1. Cancer Facts and Figures 2016. American Cancer Society (2016)
2. Bick, U., Diekmann, F.: Digital Mammography. Springer Science & Business Media, Heidelberg (2010)
3. Bria, A., Marrocco, C., Molinara, M., Tortorella, F.: An effective learning strategy for cascaded object detection. *Inf. Sci.* **340**, 17–26 (2016)

4. Bria, A., Marrocco, C., Karssemeijer, N., Molinara, M., Tortorella, F.: Deep cascade classifiers to detect clusters of microcalcifications. In: Tingberg, A., Lång, K., Timberg, P. (eds.) IWDM 2016. LNCS, vol. 9699, pp. 415–422. Springer, Cham (2016). doi:[10.1007/978-3-319-41546-8_52](https://doi.org/10.1007/978-3-319-41546-8_52)
5. Bria, A., Marrocco, C., Molinara, M., Tortorella, F.: A ranking-based cascade approach for unbalanced data. In: 2012 21st International Conference on Pattern Recognition (ICPR), pp. 3439–3442. IEEE (2012)
6. Cireşan, D.C., Meier, U., Masci, J., Maria Gambardella, L., Schmidhuber, J.: Flexible, high performance convolutional neural networks for image classification. In: IJCAI Proceedings-International Joint Conference on Artificial Intelligence, Barcelona, Spain, vol. 22, p. 1237 (2011)
7. Eadie, L.H., Taylor, P., Gibson, A.P.: A systematic review of computer-assisted diagnosis in diagnostic cancer imaging. *Eur. J. Radiol.* **81**(1), e70–e76 (2012)
8. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: AISTATS, vol. 9, pp. 249–256 (2010)
9. Greenspan, H., van Ginneken, B., Summers, R.M.: Guest editorial deep learning in medical imaging: overview and future promise of an exciting new technique. *IEEE Trans. Med. Imaging* **35**(5), 1153–1159 (2016)
10. He, K., Sun, J., Tang, X.: Single image haze removal using dark channel prior. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(12), 2341–2353 (2011)
11. He, K., Sun, J., Tang, X.: Guided image filtering. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(6), 1397–1409 (2013)
12. Jarrett, K., Kavukcuoglu, K., LeCun, Y., et al.: What is the best multi-stage architecture for object recognition? In: 2009 IEEE 12th International Conference on Computer Vision, pp. 2146–2153. IEEE (2009)
13. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: convolutional architecture for fast feature embedding. arXiv preprint [arXiv:1408.5093](https://arxiv.org/abs/1408.5093) (2014)
14. Jing, H., Yang, Y., Nishikawa, R.M.: Detection of clustered microcalcifications using spatial point process modeling. *Phys. Med. Biol.* **56**(1), 1–17 (2011)
15. Koschmieder, H.: Theorie der horizontalen Sichtweite: Kontrast und Sichtweite. Keim & Nennich (1925)
16. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems, pp. 1097–1105 (2012)
17. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436–444 (2015)
18. Liskowski, P., Krawiec, K.: Segmenting retinal blood vessels with deep neural networks. *IEEE Trans. Med. Imaging* **35**(11), 2369–2380 (2016)
19. Maitra, I.K., Nag, S., Bandyopadhyay, S.K.: Technique for preprocessing of digital mammogram. *Comput. Methods Programs Biomed.* **107**(2), 175–188 (2012)
20. Marchesi, A., Bria, A., Marrocco, C., Molinara, M., Mordang, J.J., Tortorella, F., Karssemeijer, N.: The effect of mammogram preprocessing on microcalcification detection with convolutional neural networks. In: 2017 IEEE 30th International Symposium on Computer-Based Medical Systems (CBMS), pp. 118–121. IEEE (2017)
21. Marrocco, C., Molinara, M., Tortorella, F., Rinaldi, P., Bonomo, L., Ferrarotti, A., Aragno, C., lo Moriello, S.S.: Detection of cluster of microcalcifications based on watershed segmentation algorithm. In: 2012 25th International Symposium on Computer-Based Medical Systems (CBMS), pp. 1–5. IEEE (2012)

22. Mirzaalian, H., Ahmadzadeh, M.R., Sadri, S., Jafari, M.: Pre-processing algorithms on digital mammograms. In: MVA, pp. 118–121 (2007)
23. Molinara, M., Marrocco, C., Tortorella, F.: Automatic segmentation of the pectoral muscle in mediolateral oblique mammograms. In: 2013 IEEE 26th International Symposium on Computer-Based Medical Systems (CBMS), pp. 506–509. IEEE (2013)
24. Mordang, J.-J., Janssen, T., Bria, A., Kooi, T., Gubern-Mérida, A., Karssemeijer, N.: Automatic microcalcification detection in multi-vendor mammography using convolutional neural networks. In: Tingberg, A., Lång, K., Timberg, P. (eds.) IWDM 2016. LNCS, vol. 9699, pp. 35–42. Springer, Cham (2016). doi:[10.1007/978-3-319-41546-8_5](https://doi.org/10.1007/978-3-319-41546-8_5)
25. Orlando, J.I., Prokofyeva, E., del Fresno, M., Blaschko, M.B.: Convolutional neural network transfer for automated glaucoma identification. In: 12th International Symposium on Medical Information Processing and Analysis, vol. 10160, p. 101600U (2017)
26. Pizer, S.M., Amburn, E.P., Austin, J.D., Cromartie, R., Geselowitz, A., Greer, T., ter Haar Romeny, B., Zimmerman, J.B., Zuiderveld, K.: Adaptive histogram equalization and its variations. *Comput. Vis. Graph. Image Process.* **39**(3), 355–368 (1987)
27. Samuelson, F.W., Petrick, N.: Comparing image detection algorithms using resampling. In: International Symposium on Biomedical Imaging, pp. 1312–1315 (2006)
28. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014)
29. Stomper, P.C., Geradts, J., Edge, S.B., Levine, E.G.: Mammographic predictors of the presence and size of invasive carcinomas associated with malignant microcalcification lesions without a mass. *Am. J. Roentgenol.* **181**(6), 1679–1684 (2003)
30. Tajbakhsh, N., Shin, J.Y., Gurudu, S.R., Hurst, R.T., Kendall, C.B., Gotway, M.B., Liang, J.: Convolutional neural networks for medical image analysis: full training or fine tuning? *IEEE Trans. Med. Imaging* **35**(5), 1299–1312 (2016)
31. Tarel, J.P., Hautiere, N.: Fast visibility restoration from a single color or gray level image. In: 2009 IEEE 12th International Conference on Computer Vision, pp. 2201–2208. IEEE (2009)
32. Wang, J., Yang, X., Cai, H., Tan, W., Jin, C., Li, L.: Discrimination of breast cancer with microcalcifications on mammography by deep learning. *Sci. Rep.* **6**, 27327 (2016)
33. Zeiler, M., Fergus, R.: Stochastic pooling for regularization of deep convolutional neural networks. In: Proceedings of the International Conference on Learning Representation (ICLR) (2013)
34. Zuiderveld, K.: Contrast limited adaptive histogram equalization. In: Graphics Gems IV, pp. 474–485. Academic Press Professional, Inc. (1994)