



Deep Learning Hyperparameter Optimization for Breast Mass Detection in Mammograms

Adarsh Sehgal^{1,2(✉)}, Muskan Sehgal³, Hung Manh La^{1,2}, and George Bebis²

¹ Advanced Robotics and Automation (ARA) Laboratory, Reno, USA
hla@unr.edu

² Department of Computer Science and Engineering, University of Nevada,
Reno 89557, NV, USA
asehgal@nevada.unr.edu

³ Reno, NV, USA

Abstract. Accurate breast cancer diagnosis through mammography has the potential to save millions of lives around the world. Deep learning (DL) methods have shown to be very effective for mass detection in mammograms. Additional improvements of current DL models will further improve the effectiveness of these methods. A critical issue in this context is how to pick the right hyperparameters for DL models. In this paper, we present GA-E2E, a new approach for tuning the hyperparameters of DL models for breast cancer detection using Genetic Algorithms (GAs). Our findings reveal that differences in parameter values can considerably alter the area under the curve (AUC), which is used to determine a classifier's performance.

Keywords: Breast mass detection · Genetic algorithm · GA-E2E

1 Introduction

The fast growth of machine learning, particularly deep learning, continues to pique the interest of the medical imaging community in using these approaches to increase cancer screening accuracy. Breast cancer is the second-highest cause of cancer death in women in the United States [36], and mammography screening has been shown to lower mortality [25]. Despite its advantages, screening mammography is linked to a high rate of false positives and false negatives. In the United States, the average sensitivity of digital screening mammography is 86.9%, while the average specificity is 88.9% [22]. Since the 1990s, computer-assisted detection and diagnosis (CAD) software [11] has been created to assist radiologists in improving the predicted accuracy of screening mammography. Unfortunately, evidence showed that early commercial CAD systems did not result in significant performance improvements [7, 12, 23], and that progress remained stagnant for more than a decade after their introduction. With deep learning's remarkable success in visual object recognition and detection, as well

as many other domains [20], there is a lot of interest in developing deep learning tools to help radiologists and improve screening mammography accuracy. According to recent studies [27, 28], a deep learning-based CAD system performed as well as radiologists in solo mode, and even increased radiologists' performance in assist mode.

It has been found that when deep learning is used for detecting breast cancer on screening mammography, it can improve the accuracy to as high as 96%. Because detecting subclinical breast cancer on screening mammography is a difficult image classification task, many studies [2, 17, 19] have focused on the classification of annotated lesions. [1] attempted to train neural networks using whole mammograms without any annotations, but significant lesions were difficult to find. Pre-training has been a promising method for classifier training. [15] discovered that pre-training improved both training speed and accuracy of handwritten digit recognition.

One of the interesting works [34], starts with performing patch classification on a mammogram to find the region of interest (ROI). This patch classifier is then converted to a whole image classifier by introducing a combination of the VGG network [35] and the residual network (Resnet) [14]. [34] uses CBIS-DDSM [21] dataset to perform various experiments to generate an end-to-end training process, which we will refer to as E2E. Depending on the combination of top layers used, the authors have generated various models, which can be transferred to other datasets with no ROI information. Since the learning is being transferred, it needs additional tuning of parameters, as mentioned in [34]. Even though some of the parameters can be tuned manually, it can be a tedious and inefficient task to train the parameters which have a large number of possible values. With each of the parameters assuming a large number of possible values, combinations of these values can grow exponentially. This makes it almost impossible to find near-optimal values for the parameters.

Even though the performance of the whole image classifier to detect a mass in a mammogram has been improved, the authors [34] do not justify the use of many constant parameters used in transferring the learning to a dataset with minimal/no ROI information. [30, 31] have shown that automatic tuning of such parameters can greatly enhance the performance of the overall system. [30] made use of GA with Deep Reinforcement Learning (GA-DRL), and [31] used GA with Lidar-monocular visual odometry (GA-LIMO). Both of these works proved that when GA is used for parameter tuning, it has the potential of producing encouraging results. A neural-genetic algorithm for feature selection system was put up by [39] to categorize microcalcification patterns in digital mammograms. The outcomes in [39] demonstrated that the proposed algorithm was successful in locating a suitable feature subset, which also resulted in a high classification rate. Some other related works utilizing GA include: [18, 29, 32, 33, 37].

In this work, we propose a novel algorithm (called GA-E2E), which uses a Genetic Algorithm (GA) on the E2E model to tune the parameters for transferring the learning to Inbreast database [24]. We are using the pre-trained model files from E2E. These pre-trained models were generated for whole-mammogram classification. Each run of a fitness function performs E2E training on the

Inbreast dataset using the parameter values decided by the GA. The output of the fitness function is Area Under the Curve (AUC). We are taking a GA reference from an open-source implementation [4]. Our tests revealed that GA-E2E is computationally demanding; hence we propose employing a standalone cluster of two PCs connected by a network switch. We used Spark to create this configuration. Our findings demonstrate that GA-E2E contributed to a higher AUC. GA focuses search on the promising portion of the search space, which saves time, computational power, and manual efforts. Open source code is available at <https://github.com/aralab-unr/GA-mammograms>. In the next section, we will describe some closely related works.

2 Background and Motivation

2.1 End-to-End Pipeline




Fig. 1. E2E pipeline.

Converting a Classifier from Recognizing Patches to Whole Images. A common technique for classifying or segmenting large complicated images is to utilize a classifier in sliding window mode to recognize local patches of an image and provide a grid of probabilistic outputs. This is followed by a process that summarizes the outputs of the patch classifier to get the final classification or segmentation result. Using entire slide images of sentinel lymph node biopsies, such approaches have been utilized to detect metastatic breast cancer [41] and to segment neuronal membranes in microscopic images [6]. This technique, however, necessitates two phases, each of which must be tuned separately. Here, training on entire images is performed by combining the two phases into a single step (Fig. 1).

The h function takes complete photos as input and outputs labels for the entire image. As a result, it may be trained from the beginning to the end, giving it two advantages over the two-step technique. First, the entire network may be trained simultaneously, eliminating sub-optimal solutions at each stage; second, the learned network can be transferred to another dataset without relying on ROI annotations explicitly. Large mammography databases with ROI

annotations are difficult to come by and very expensive. DDSM, the world’s largest public library of ROI annotations for digitized film mammography, contains thousands of pictures with pixel-level annotations that can be used to train a patch classifier. Once the patch classifier has been converted into a whole image classifier h , it can be fine-tuned using only image-level labels from other databases. This method significantly reduces the need for ROI annotations, and it has a wide range of applications in medical imaging, including breast cancer diagnosis on screening mammograms.

Network Design. The E2E pipeline is depicted in Fig. 1. A separate class of neural networks is represented in this pipeline by function f . The heatmap [5] is of size $u \times v \times c$. A $u \times v$ grid of probabilistic outputs of c classes that were obtained after modifying the input is referred to as a heatmap. A simple way to make an entire picture classifier from a patch classifier is to flatten the heat map and connect it to the image classification output with fully-connected layers. After the heatmap, a max-pooling layer can be employed to boost the model’s translational invariance to the patch classifier’s output. It is also possible to construct a shortcut between the heatmap and the output to make training simple. The heatmap is produced directly by the patch classifier’s output, which makes use of the softmax activation:

$$f(z)_j = \frac{e^{z_j}}{\sum_{i=1}^c e^{z_i}} \text{ for } j = 1, \dots, c. \quad (1)$$

Convolutional layers, which preserve spatial information, are also employed as top layers. On top of the patch classifier layers, two blocks of convolutional layers (VGG or Resnet) can be added, followed by a global average pooling layer, and finally the image’s classification output (Fig. 1). As a result, for full picture classification, this approach provides an “all convolutional” network. The heatmap reduces the depth of the feature map between the patch classifier layers and the top layers, as shown in Fig. 1, potentially causing information loss in the entire picture classification. As a result, when the heatmap was completely removed from the overall image classifier, the upper layers were able to fully utilize the patch classifier’s capabilities.

Developing Patch and Whole Image Classifiers on CBIS-DDSM. The pixel-level annotations for the ROIs are stored in the CBIS-DDSM database, along with their pathologically confirmed labels: benign or malignant. Each ROI is also labeled as a calcification or mass. The majority of mammograms only had one ROI. All mammograms were converted to PNG format and reduced using interpolation to 1152×896 , with no picture cropping. The limitation of GPU memory size prompted the downsizing. By sampling picture patches from ROIs and background regions, two patch datasets were constructed. The size of all patches was 224×224 , which was large enough to cover the majority of the ROIs annotated. The first dataset (S1) was made up of two sets of patches, one of which was centered on the ROI and the other was a random background patch from the same image.

Transfer Learning for Whole Image Classification on Inbreast. The dataset was set up and processed. The Inbreast [24] dataset is a public database that contains FFDM images that were acquired more recently. These images exhibit different intensity profiles than digitized film mammograms from the CBIS-DDSM. As a result, Inbreast is a great way to see if a full picture classifier can be used across different mammography platforms. There are 115 patients and 410 mammograms in the Inbreast database, with both CC and MLO views. Each view is considered a separate image. The radiologists' BI-RADS [38] assessment categories are defined as follows in the Inbreast database: 0: insufficient examination; 1: no findings; 2: benign; 3: probably benign; 4: suspicious; 5: highly indicative of malignancy; 6: known biopsy-proven cancer. [34] provides a platform to perform to transfer the learning to any dataset. In this work, we use [34] on Inbreast, and show that when the parameters are tuned in an optimal way, it can greatly improve the AUC.

2.2 Genetic Algorithm

GAs [8, 13, 16] were developed to investigate poorly-understood areas [9], where a comprehensive search is impossible and conventional search methods perform poorly. When GAs are used as function optimizers, their goal is to maximize fitness that is related to the optimization goal. Evolutionary computing approaches in general, and GAs in particular, have had a lot of empirical success on a variety of difficult design and optimization problems. They start with a population of randomly started candidate solutions, which are frequently encoded as a string (chromosome). A selection operator reduces the search space to the most plausible spots, but crossover and mutation operators generate new possibilities.

2.3 Binary Tournament Selection

A tournament selection (TS) operator is based on a competition among a group of parents. The fitness of a potential solution is calculated for each parent, and the parent with the highest fitness is chosen. The term "binary tournament" refers to a tournament with a field size of two players, which is the most basic kind of tournament selection [3]. The binary tournament selection (BTS) process begins with the selection of two people at random. The fitness levels of these people are then assessed. The person with the best fitness is then chosen. The tournament selection has the advantage of being able to handle either minimization or maximization problems without requiring any structural changes. Furthermore, the use of a negative value is unrestricted [40].

2.4 Simulated Binary Crossover (SBX)

The search power of the simulated binary crossover (SBX) [10] is similar to that of the single-point crossover. The distinction between real-coded GAs with SBX and binary-coded GAs with single-point crossover implementations is that the former method eliminates variable coding and creates a child string from a probability distribution based on the location of the parent strings.

2.5 Polynomial Mutation

A polynomial mutation operator with a user-defined index parameter (η_m) was proposed by Deb and Agrawal (1999). They determined from theoretical research that η_m causes a perturbation of $O((b-a)/\eta_m)$ in a variable, where a and b are the variable's lower and upper bounds. They also discovered that on most of the issues they tried, a value of $\eta_m \in [20, 100]$ was sufficient. A polynomial probability distribution is utilized in this operator to perturb a solution in the neighborhood of a parent. The mutation operator adjusts the probability distribution to the left and right of a variable value so that no value outside the specified range $[a, b]$ is formed. The mutated solution p' for a certain variable is constructed for a random number u formed within $[0, 1]$ for a specified parent solution $p \in [a, b]$, as follows:

$$p' = \begin{cases} p + \bar{\delta}_L(p - x_i^{(L)}), & \text{for } u \leq 0.5, \\ p + \bar{\delta}_R(x_i^{(U)} - p), & \text{for } u > 0.5. \end{cases} \quad (2)$$

3 Approach

3.1 Pre-processing Images

Inbreast dataset contains images in DICOM format. The dataset comes with a mapping document that maps mammogram image file names to other information such as patient ID, BI-RADS assessment categories, etc. It also provides sample MATLAB code to read the images. We used this code along with our code to split the DICOM images into two classes: positive and negative. We, then, converted these images to PNG format to be used with the E2E algorithm. Only a limited number of converted images were used in our experiments. The images were downsized to 1152×896 . Training, validation, and test sub-set splits are described in the following sub-section. Two phases are used in the experiments: GA training, and testing. GA training uses training images to generate a new model built by training the last two layers in E2E. The generated model is then used on validation sub-sets to compute AUC. The output of the fitness function used for GA is the average of AUCs from all epochs. The testing phase is used to compare the GA-found parameters with the original values of the parameters. This phase uses train and validation sub-sets to generate a model, which is then used with the test sub-set to compute the final AUC.

3.2 Building Spark Cluster

This study uses Spark [42] to run GA in a distributed setup. This is important because each run of the fitness function is computationally intensive. We used two computers to build a cluster. Advantages of cluster setup include scalability and the ability for a program to finish in the real time possible. We will describe the details of the GA-E2E algorithm in the next sub-section. The word *cluster* is used to describe the setup of two or more computers running in parallel

to achieve a common goal. We chose Spark for our work because it provides easy-to-use configurations, and online resources/documentation. The high-level architecture of Spark is shown in Fig. 2. A cluster manager connects *SparkContext* and *Worker Nodes*. For our work, the cluster was built as shown in Fig. 3. The computers are connected via a network switch. For our implementation specifically, we ran one worker on each of the machines.




Fig. 2. Architecture of Spark cluster.

3.3 GA-E2E Algorithm

Now, we set up GA from [4]. As described in the background section, GA uses BTS, SBX crossover, and polynomial mutation. Since the total training epochs are set at ten, the approximate time for each fitness function evaluation is about 70s for each of the workers. Considering this, we have set the population size to 70, to be run for 70 generations. GA was run on the model from [34] using transfer learning on the Inbreast dataset. For each evaluation of the fitness function, GA selects values of parameters, uses them to train the network, computes average validation AUC, and finally uses this trained model on the test sub-set. The




Fig. 3. Example representation of master and worker setup.

AUC returned from the fitness function is the AUC computed on the test sub-set. GA training phases use 30 images in the training sub-set, and 12 images in the validation sub-set. Algorithm 1 describes the high-level overview of the GA-E2E algorithm.

Algorithm 1. Proposed GA-E2E Algorithm

- 1: Choose population of n chromosomes
 - 2: Set the values of parameters into the chromosome
 - 3: **for** all chromosome values **do**
 - 4: Run the code to fine-tune the existing model using independent dataset (Inbreast)
 - 5: **return** AUC
 - 6: **end for**
 - 7: Perform SBX Crossover
 - 8: Perform Polynomial Mutation
 - 9: Repeat for required number of generations to find optimal solution
-

3.4 Evaluation

The experiments are broken into two sections, as previously stated: GA training and testing. In both phases, the training and validation sets are identical. However, the final AUC is computed using 100 mammogram images in the testing phase.

4 Experimental Results

4.1 Experimental Setup

It is important to plan the experiments to be executed to test the algorithm, some details of which we have described in the previous sections. The overall algorithm is described in Algorithm 1. The workspace was set up with [34]. This code link provides further links to the pre-trained models on whole image classification using the CBIS-DDSM dataset.

Spark requires that the master and worker be able to communicate key-less using SSH. So, we established password-less SSH between our machines. Both machines can SSH as a user/root into other machines. Both computers have the same username, folder structure, GA-E2E code, same images used in train/validation sub-sets, and the .jar files used by Spark. The number of parallel processes was set to two, one for each worker.

Table 1 shows the parameters we used in the GA-E2E algorithm. Each set of parameters decided by the GA was used to perform transfer learning on the Inbreast dataset. The reward for the fitness function is the validation AUC averaged over epochs.

Table 1. Parameters used in the GA-E2E, along with their description.

Parameter	Description
pos-cls-weight	Weight of positive class
neg-cls-weight	Weight of negative class
weight-decay	Stage 1
weight-decay2	Stage 2
init-learningrate	Learning rate for stage 1
all-layer-multiplier	Multiplier for stages other than stage 1

Each of the parameters is binary encoded into a single chromosome, an example of which is shown in Fig. 4.




Fig. 4. Chromosome representation for the GA.

From the various manual experiments we did, we would like to mention a few of them. In one such experiment, we used 8 images in a training sub-set,

out of which four were negative, and four were positive. By positive we mean that the mammogram image of the breast is malignant, and by negative we mean that the mammogram image of the breast is benign. A similar setup was done for the test sub-set. For the validation sub-set, two images each of the positive, and negative classes were used. Training epochs were fixed to four, and the number of cores of the GPU to be used was fixed to six. The sub-sets chosen had multiple views (CC/MLO) of the same patient. Each epoch took about 160s to complete. On using the learned model on a test sub-set, we got the AUC as 0.5. We experimented with a various number of images in different sub-sets and found that the program ran for a very long time when more images were used. Since GA needs to run the training a few hundred times to find the near-optimal parameter values, for the rest of our implementation, we decided to use the number of images as in the sub-set division described in this section. However, to run GA-E2E, we used only CC views from different patients for training, validation, and testing.

Finally, we used 30 images for training, 12 for validation, and 100 for test sub-sets. Each GA fitness function was run for ten epochs, while the testing phase had 100 epochs. The E2E uses an early stopping approach; hence the training stopped between 30–50 epochs each time the training was run. This was applicable for both original and GA-found parameters. Our experiment assumes each view of the mammogram as a separate image.

4.2 Running GA-E2E

The system has now been set up to run the algorithm GA-E2E from 1. Based on the previously described system setup, the GA completed running in about two days. We captured the logs of each run of the fitness function and the corresponding values for each of the parameters. Table 2 shows the comparison of values of parameters for E2E and GA-E2E. E2E parameters are the ones originally used in transfer learning with Inbreast in [34]. It can be observed that changing parameter values has a significant effect on the AUC of the test sub-set. The increase in AUC is as high as 11% when compared the to the original AUC.

Table 2. E2E vs. GA-E2E values of parameters.

Parameter	Range	E2E, AUC = 0.51	GA-E2E AUC = 0.57
pos-cls-weight	0–0.999	1.0	0.948
neg-cls-weight	0–0.999	1.0	0.319
weight-decay	0–0.999	0.0001	0.269
weight-decay2	0–0.999	0.0001	0.225
init-learningrate	0–0.999	0.01	0.0008
all-layer-multiplier	0–0.999	0.1	0.475

4.3 Training Evaluation

It is significantly important to monitor the progress of a GA as it is running to optimize the performance of the system. From the way GA's work, some of the chromosomes will outperform others. Ultimately, the average fitness would increase towards the maximum fitness value. Figure 5 shows the best and average fitness value (AUC) over generations.




Fig. 5. Fitness value evaluation over GA generations.

4.4 Analysis

It can be observed that all generations except the first two, saw the best fitness value of 1.0. However, the average fitness value started at 0.5, and ultimately, the GA was able to have the entire population at 1.0 AUC. GA took about 18 generations to reach the maximum success rate. Even though the GA reached its maximum success rate, the system was left to run until completion.

5 Experimental Issues

Configuring the Spark cluster posed major challenges. Each configuration parameter needs to be appropriately set for the standalone cluster to work using GPUs. The GPU makes program execution much faster, and running the program on a cluster of two computers, cuts the running time to half. Any combination of incorrectly set-up configuration values can result in an unequal load on the machines, which can further lead to some machines being under-utilized/over-utilized.

6 Conclusion and Future Work

6.1 Conclusion

This study demonstrates that there is no well-defined correlation between parameter values and AUC. The AUC may not necessarily rise as the parameter values rise. As a result, a straightforward hill climber [26] may be useless, and the role of GA in fine-tuning the parameter values becomes significant. We also show that GA focuses search in the promising search space. GA started with an AUC close to zero. As the GA progressed, the AUC got higher and ultimately surpassed the originally achieved AUC of 0.51. This points to an increase of around 11% in comparison to the original AUC.

6.2 Future Work

We believe that GA may be able to reach the AUC of 1.0 much earlier if GA operators are designed specifically for E2E. Also, different convolution layers can be tested with E2E for better adaptability of the learned model with an independent dataset. Adding more computers to the cluster may reduce execution time further. Modified GA-E2E could be run on a supercomputer, which would enable the ability to perform various experiments using a variety of datasets.

References

1. Aboutalib, S.S., Mohamed, A.A., Berg, W.A., Zuley, M.L., Sumkin, J.H., Wu, S.: Deep learning to distinguish recalled but benign mammography images in breast cancer screening deep learning in mammography. *Clin. Cancer Res.* **24**(23), 5902–5909 (2018)
2. Arevalo, J., González, F.A., Ramos-Pollán, R., Oliveira, J.L., Lopez, M.A.G.: Convolutional neural networks for mammography mass lesion classification. In: 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pp. 797–800. IEEE (2015)
3. Bäck, T., Fogel, D.B., Michalewicz, Z.: *Evolutionary Computation 1: Basic Algorithms and Operators*. CRC Press, Boca Raton (2018)
4. Benitez-Hidalgo, A., Nebro, A.J., Garcia-Nieto, J., Oregi, I., Del Ser, J.: jmetalpy: a python framework for multi-objective optimization with metaheuristics. *Swarm Evol. Comput.* **51**, 100598 (2019)
5. Bulat, A., Tzimiropoulos, G.: How far are we from solving the 2D& 3D face alignment problem?(and a dataset of 230,000 3D facial landmarks). In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1021–1030 (2017)
6. Ciresan, D., Giusti, A., Gambardella, L., Schmidhuber, J.: Deep neural networks segment neuronal membranes in electron microscopy images. In: *25th Proceedings of the Conference on Advances in Neural Information Processing Systems* (2012)
7. Cole, E.B., Zhang, Z., Marques, H.S., Hendrick, R.E., Yaffe, M.J., Pisano, E.D.: Impact of computer-aided detection systems on radiologist accuracy with digital mammography. *AJR Am. J. Roentgenol.* **203**(4), 909 (2014)
8. Davis, L.: *Handbook of Genetic Algorithms*. Van Nostrand Reinhold, Washington, DC (1991)

9. De Jong, K.: Learning with genetic algorithms: an overview. *Mach. Learn.* **3**(2–3), 121–138 (1988)
10. Deb, K., Agrawal, R.B., et al.: Simulated binary crossover for continuous search space. *Complex Syst.* **9**(2), 115–148 (1995)
11. Elter, M., Horsch, A.: Cadx of mammographic masses and clustered microcalcifications: a review. *Med. Phys.* **36**(6Part1), 2052–2068 (2009)
12. Fenton, J.J., et al.: Influence of computer-aided detection on performance of screening mammography. *N. Engl. J. Med.* **356**(14), 1399–1409 (2007)
13. Goldberg, D.E., Holland, J.H.: Genetic algorithms and machine learning. *Mach. Learn.* **3**(2), 95–99 (1988)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778 (2016)
15. Hinton, G.E., Osindero, S., Teh, Y.W.: A fast learning algorithm for deep belief nets. *Neural Comput.* **18**(7), 1527–1554 (2006)
16. Holland, J.H.: Genetic algorithms. *Sci. Am.* **267**(1), 66–73 (1992)
17. Jamieson, A.R., Drukker, K., Giger, M.L.: Breast image feature learning with adaptive deconvolutional networks. In: *Medical Imaging 2012: Computer-Aided Diagnosis*, vol. 8315, pp. 64–76. SPIE (2012)
18. Jiang, J., Yao, B., Wason, A.: A genetic algorithm design for microcalcification detection and classification in digital mammograms. *Comput. Med. Imaging Graph.* **31**(1), 49–61 (2007)
19. Kooi, T., et al.: Large scale deep learning for computer aided detection of mammographic lesions. *Med. Image Anal.* **35**, 303–312 (2017)
20. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436–444 (2015)
21. Lee, R.S., Gimenez, F., Hoogi, A., Miyake, K.K., Gorovoy, M., Rubin, D.L.: A curated mammography data set for use in computer-aided detection and diagnosis research. *Sci. Data* **4**(1), 1–9 (2017)
22. Lehman, C.D., et al.: National performance benchmarks for modern screening digital mammography: update from the breast cancer surveillance consortium. *Radiology* **283**(1), 49–58 (2017)
23. Lehman, C.D., et al.: Diagnostic accuracy of digital screening mammography with and without computer-aided detection. *JAMA Int. Med.* **175**(11), 1828–1837 (2015)
24. Moreira, I.C., Amaral, I., Domingues, I., Cardoso, A., Cardoso, M.J., Cardoso, J.S.: Inbreast: toward a full-field digital mammographic database. *Acad. Radiol.* **19**(2), 236–248 (2012)
25. Oeffinger, K.C., et al.: Breast cancer screening for women at average risk: 2015 guideline update from the American cancer society. *J. Am. Med. Assoc.* **314**(15), 1599–1614 (2015)
26. Prügel-Bennett, A.: When a genetic algorithm outperforms hill-climbing. *Theoret. Comput. Sci.* **320**(1), 135–153 (2004)
27. Rodríguez-Ruiz, A., et al.: Detection of breast cancer with mammography: effect of an artificial intelligence support system. *Radiology* **290**(2), 305–314 (2019)
28. Rodríguez-Ruiz, A., et al.: Stand-alone artificial intelligence for breast cancer detection in mammography: comparison with 101 radiologists. *J. Natl. Cancer Inst.* **111**(9), 916–922 (2019)
29. Sehgal, A.: Genetic algorithm as function optimizer in reinforcement learning and sensor odometry. Master’s thesis, University of Nevada, Reno (2019)

30. Sehgal, A., La, H., Louis, S., Nguyen, H.: Deep reinforcement learning using genetic algorithm for parameter optimization. In: 2019 Third IEEE International Conference on Robotic Computing (IRC), pp. 596–601. IEEE (2019)
31. Sehgal, A., Singandhupe, A., La, H.M., Tavakkoli, A., Louis, S.J.: Lidar-monocular visual odometry with genetic algorithm for parameter optimization. In: Bebis, G., et al. (eds.) ISVC 2019. LNCS, vol. 11845, pp. 358–370. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-33723-0_29
32. Sehgal, A., Ward, N., La, H., Louis, S.: Automatic parameter optimization using genetic algorithm in deep reinforcement learning for robotic manipulation tasks. arXiv preprint [arXiv:2204.03656](https://arxiv.org/abs/2204.03656) (2022)
33. Sehgal, A., Ward, N., La, H.M., Papachristos, C., Louis, S.: GA-DRL: genetic algorithm-based function optimizer in deep reinforcement learning for robotic manipulation tasks. arXiv preprint [arXiv:2203.00141](https://arxiv.org/abs/2203.00141) (2022)
34. Shen, L., Margolies, L.R., Rothstein, J.H., Fluder, E., McBride, R., Sieh, W.: Deep learning to improve breast cancer detection on screening mammography. *Sci. Rep.* **9**(1), 1–12 (2019)
35. 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)
36. Smith, R.A., et al.: American cancer society guidelines for breast cancer screening: update 2003. *Cancer j. Clin.* **53**(3), 141–169 (2003)
37. Sun, Y., Babbs, C., Delp, E.: A comparison of feature selection methods for the detection of breast cancers in mammograms: adaptive sequential floating search vs. genetic algorithm. In: 2005 IEEE Engineering in Medicine and Biology 27th Annual Conference, pp. 6532–6535. IEEE (2006)
38. Timmers, J., et al.: The breast imaging reporting and data system (BI-RADS) in the Dutch breast cancer screening programme: its role as an assessment and stratification tool. *Eur. Radiol.* **22**(8), 1717–1723 (2012)
39. Verma, B., Zhang, P.: A novel neural-genetic algorithm to find the most significant combination of features in digital mammograms. *Appl. Soft Comput.* **7**(2), 612–625 (2007)
40. Wahde, M.: *Biologically Inspired Optimization Methods: An Introduction*. WIT Press (2008)
41. Wang, D., Khosla, A., Gargeya, R., Irshad, H., Beck, A.H.: Deep learning for identifying metastatic breast cancer. arXiv preprint [arXiv:1606.05718](https://arxiv.org/abs/1606.05718) (2016)
42. Zaharia, M., Chowdhury, M., Franklin, M.J., Shenker, S., Stoica, I.: Spark: Cluster computing with working sets. In: 2nd USENIX Workshop on Hot Topics in Cloud Computing, (HotCloud 2010) (2010)