DeepMiner: Discovering Interpretable Representations for Mammogram Classification and Explanation

Jimmy Wu^{1,*}, Bolei Zhou², Diondra Peck³, Scott Hsieh⁴, Vandana Dialani, MD⁵,
Lester Mackey⁶, and Genevieve Patterson^{7,*}

¹ Department of Computer Science, Princeton University, Princeton, USA

² The Chinese University of Hong Kong, Hong Kong, China

³ Microsoft, Redmond, USA

⁴ Department of Radiological Sciences, UCLA, Los Angeles, USA

⁵ Beth Israel Deaconess Medical Center, Cambridge, USA

⁶ Microsoft Research New England, Cambridge, USA

⁷ VSCO, Oakland, USA

ABSTRACT. We propose DeepMiner, a framework to discover interpretable representations in deep neural networks and to build explanations for medical predictions. By probing convolutional neural networks (CNNs) trained to classify cancer in mammograms, we show that many individual units in the final convolutional layer of a CNN respond strongly to diseased tissue concepts specified by the Breast Imaging-Reporting and Data System (BI-RADS) lexicon. After expert annotation of the interpretable units, our proposed method is able to generate explanations for CNN mammogram classification that are consistent with ground truth radiology reports on the Digital Database for Screening Mammography. We show that DeepMiner not only enables better understanding of the nuances of CNN classification decisions but also possibly discovers new visual knowledge relevant to medical diagnosis.

Keywords: deep learning, interpretability, human-in-the-loop machine learning, mammography

MEDIA SUMMARY

Deep learning algorithms are often criticized for producing uninterpretable, black-box results. This lack of interpretability is especially concerning for medical decisions. Here we introduce DeepMiner, a human-in-the-loop framework for interpreting the medical predictions of deep neural networks. After training a neural network to detect malignancy in mammograms, DeepMiner identifies the visual phenomena most influential to the network’s decision-making and calls upon human experts to label those phenomena (for example as “calcified vessels,” “masses with smooth edges,” or

*jw60@cs.princeton.edu, genevieve@vSCO.co

“normal breast tissue”). These expert annotations enable DeepMiner to explain its mammogram classifications by summarizing the named phenomena driving each prediction.

1. INTRODUCTION

Deep convolutional neural networks (CNNs) have made great progress in visual recognition challenges such as object classification (Krizhevsky et al., 2012) and scene recognition (Zhou et al., 2014), even reaching human-level image understanding in some cases (He et al., 2015b). Recently, CNNs have been widely used in medical image understanding and diagnosis (Esteva et al., 2017; Rajpurkar et al., 2017; Wang et al., 2017). However, with millions of model parameters, CNNs are often treated as ‘black-box’ classifiers, depriving researchers of the opportunity to investigate what is learned inside the network and explain the predictions being made. Especially in the domain of automated medical diagnosis, it is crucial to have interpretable and explainable machine learning models.

Several visualization methods have previously been proposed for investigating the internal representations of CNNs. For example, internal units of a CNN can be represented by reverse-mapping features to the input image regions that activate them most (Zeiler & Fergus, 2014) or by using backpropagation to identify the most salient regions of an image (Mahendran & Vedaldi, 2014; Simonyan et al., 2014). Our work is inspired by recent work that visualizes and annotates interpretable units of a CNN using Network Dissection (Bau et al., 2017).

Meanwhile, recent work in automated diagnosis methods has shown promising progress toward interpreting models and explaining model predictions. Wu et al. (2018) show that CNN internal units learn to detect medical concepts that match the vocabulary used by practicing radiologists. Rajpurkar et al. (2017) and Wang et al. (2017) use the class activation map defined by Zhou et al. (2016) to explain informative regions relevant to final predictions. Zhang et al. (2017) propose a hybrid CNN and LSTM (long short-term memory) network capable of diagnosing bladder pathology images and generating radiological reports if trained on sufficiently large image and diagnostic report data sets. However, their method requires training on full medical reports. In contrast, our approach can be used to discover informative visual phenomena spontaneously with only coarse training labels. Jing et al. (2017) successfully created a visual and semantic network that directly generates long-form radiological reports for chest X-rays after training on a data set of X-ray images and associated ground truth reports. However, even with these successes, many challenges remain. Wu et al. (2018) only show that interpretable internal units are correlated with medical events without exploring ways to explain the final prediction. The heatmaps generated by Rajpurkar et al. (2017) and Wang et al. (2017) qualitatively identify important *locations* in an image but fail to identify *specific concepts*. Jing et al. (2017) train their models on large-scale medical report data sets; however, large text corpora associated with medical images are not easily available in other scenarios. Additionally, Zhang et al. (2017) acknowledge that their current classification model produces false alarms from which it cannot yet self-correct.

In this article, we propose a general framework called *DeepMiner* for discovering medical phenomena in coarsely labeled data and generating explanations for final predictions, with the help of a few human expert annotators. We apply our framework to mammogram classification, an already well-characterized domain, in order to provide confidence in the capabilities of deep neural networks for discovery, classification, and explanation.

To the best of our knowledge, our work is the first automated diagnosis CNN that can both discover discriminative visual phenomena for breast cancer classification and generate interpretable, radiologist-collaborative explanations for its decision-making. Our main contribution is two-fold: (1) we propose a human-in-the-loop framework to enable medical practitioners to explore the behavior of CNN models and annotate the visual phenomena discovered by the models, and (2) we leverage the internal representations of CNN models to explain their decision making, without the

use of external large-scale report corpora. Our data, results, and open source code replicating all experiments can be found at <https://github.com/jimmyhwu/ddsm-visual-primitives>. A copy of our data and pretrained CNN weights is also available through the Harvard Dataverse (Wu et al., 2021).

2. THE DEEPMINER FRAMEWORK

The DeepMiner framework consists of three phases, as illustrated in Figure 1. In the first phase, we train standard neural networks for classification on patches cropped from full mammograms. Then, in the second phase, we invite human experts to annotate the top class-specific internal units of the trained networks. Finally, in the third phase, we use the trained network to generate explainable predictions by ranking the contributions of individual units to each prediction.

Figure 1. The DeepMiner framework: Illustration of our framework for mammogram classification and explanation.

In this work, we select mammogram classification as the testing task for our DeepMiner framework. The classification task for the network is to correctly classify mammogram patches as normal (containing no findings of interest), benign (containing only noncancerous findings), or malignant (containing cancerous findings). Our framework can be further generalized to other medical image classification tasks. Note that we refer to convolutional filters in our CNNs as ‘units,’ as opposed to ‘neurons,’ to avoid conflation with the biological entities.

2.1. Data Set and Training. We choose ResNet-152 pretrained on ImageNet (He et al., 2015a) as our reference network due to its outstanding past performance on object classification, image segmentation, and fine-grained localization across a variety of image domains (He et al., 2017). We replaced the final layer of the network with a three-class classification layer and fine-tuned all network weights to classify mammogram patches as normal, benign, or malignant using data from the Digital Database for Screening Mammography (DDSM) (Heath et al., 2000). DDSM is a data set compiled to facilitate research in computer-aided breast cancer screening. It consists of 2,620 cases, each including two images of each breast, a BI-RADS rating of 0-5 for cancer risk, a radiologist’s subjective subtlety rating for each finding, and a BI-RADS keyword description of

abnormalities. Labels include image-wide designations (e.g., malignant, benign, and normal) and pixel-wise segmentations of lesions (Heath et al., 2000).

In our experiments, we train our classifier on 80% of the DDSM cases, reserve 10% of the cases as a hold-out set to select the learning rate for training, and use the final 10% of the images as a test set for evaluating CNN classification performance. We highlight that all images belonging to the same case are placed in the same data set partition to mimic the standard use-case of evaluating a trained system on a previously unseen case.

To increase the number of training examples for fine-tuning, we extract smaller image patches from our mammograms in a sliding window fashion. The dimensions of each image patch are 25% of the width of the original mammogram, and overlapping patches are extracted using a stride of 50% of the patch width. We use a background texture classifier to discard any patch containing less than 50% breast tissue. We create three class labels (normal, benign, malignant) for each image patch based on (1) whether at least 30% of the patch contains benign or malignant tissue and (2) whether at least 30% of a benign or malignant finding is located in that patch. These patch labels were determined automatically by calculating pixel overlap with the ground-truth lesion segmentation.

We fine-tune our reference network using the PyTorch (Paszke et al., 2019) implementation of stochastic gradient descent (Bottou et al., 2018; Kiefer, Wolfowitz, et al., 1952; Robbins & Monro, 1951). The network was trained for five epochs with a median runtime of 85 minutes per epoch on a single Titan X Pascal GPU. We choose a batch size of 32 to ensure that all images in the batch fit into GPU memory, select a learning rate of 0.0001 to balance overfitting and underfitting as judged by differences in training set and hold-out set accuracies, and set the momentum and weight decay hyperparameters to the default values for ImageNet training in PyTorch (0.9 and 0.0001, respectively). The final test set performance of the trained network is presented in Section 3.1, and we provide a brief introduction to CNNs and their training in Appendix A.

2.2. Human Annotation of Visual Primitives Used by CNNs. We use our hold-out split of DDSM to create visualizations for units in the final convolutional layer of our fine-tuned ResNet-152. We choose the final layer since it is most likely to contain high-level semantic concepts due to the hierarchical structure of CNNs.

It would be infeasible to annotate all 2048 units in the last convolutional layer. Instead, we select a subset of the units deemed most frequently ‘influential’ to classification decisions. Given a classification decision for an image, we define the influence of a unit toward that decision as the unit’s maximum activation score on that image multiplied by the weight of that unit for a given output class in the final fully connected layer.

We selected the 20 most frequently influential units for each of the three classes and asked human experts to annotate the resulting 60 units. For the normal tissue class, if the 20 units we selected were annotated, those annotations would account for 59.27% of the per-image top eight units over all of the hold-out set images. The corresponding amount for the benign class is 69.77% and for the malignant class is 75.82%.

We create visualizations for each individual unit by passing every image patch from all mammograms in our hold-out set through our classification network. For each unit in the final convolutional layer, we record the unit’s maximum activation value as well as the receptive field from the image patch that caused the measured activation. To visualize each unit (see Figures 2 and 3), we display the top activating image patches sorted by their activation score and further segmented by the binarized and upsampled response map of that unit.

A radiologist and a medical physicist specializing in mammography annotated the 60 most frequently influential units we selected. We compare the named phenomena detected by these units

to the BI-RADS lexicon (Orel et al., 1999). The experts used the annotation interface shown in Figure 2. Our survey displays a table of dozens of the top scoring image patches for the unit being visualized. When the expert mouses over a given image patch, the mammogram that the patch came from is displayed on the right with the patch outlined in red. This gives the expert some additional context. From this unit preview, experts are able to formulate an initial hypothesis of what phenomena a unit detects.

Figure 2. The interface of the DeepMiner survey: Experts used this survey form to label influential units. The survey asks questions such as: “Do these images show recognizable phenomena?” and “Please describe each of the phenomena you see. For each phenomenon please indicate its association with breast cancer.” In the screenshot above, the radiologist who was our expert-in-the-loop has labeled the unit’s phenomena as ‘edge of mass with circumscribed margins.’

Of the 60 units selected, 46 were labeled by at least one expert as detecting a nameable medical phenomenon. Figure 3 shows five of the annotated units. In this figure, each row illustrates a different unit. The table lists the unit ID number, the BI-RADS category for the concept the unit is detecting, the expert-provided unit annotation, and a visual representation of the unit. We visualize each unit by displaying the top four activating image patches from the hold-out set. The unit ID number is listed to uniquely identify each labeled unit in the network used in this article, which will be made publicly available upon publication.

Figure 3 demonstrates that the DeepMiner framework discovers significant medical phenomena, relevant to mammogram-based diagnosis. Because breast cancer is a well-characterized disease, we are able to show the extent to which discovered unit detectors overlap with phenomena deemed to be important by the radiological community. For diseases less well understood than breast cancer, DeepMiner could be a useful method for revealing unknown discriminative visual features helpful in diagnosis and treatment planning.

2.3. Explaining Network Decisions. We further use the annotated units to build an explanation for single image prediction. We first convert our trained network into a fully convolutional network (FCN) using the method described in Long et al. (2015) and remove the global average pooling layer. The resulting network is able to take in full mammogram images and output probability maps aligned with the input images.

As illustrated in Figure 1, given an input mammogram, we output a classification as well as the Class Activation Map (CAM) proposed in Zhou et al. (2016). We additionally extract the activation maps for the units most influential toward the classification decision. By looking up the corresponding expert annotations for those units, we are able to see which nameable visual phenomena contributed to the network’s final classification. For examples of the DeepMiner explanations, please refer to Section 3.2.

Unit ID	BI-RADS Lexicon	Expert Annotation	Top Activated Images
1814	Mass, Associated Features	Edge of a mass with circumscribed margins
1152	Calcification	Malignant pleomorphic calcifications
860	Calcification	Benign vascular calcifications
1299	Associated Features	Spiculation
1468	Mass	Masses with smooth edges

Figure 3. Interpretable units discovered by DeepMiner: The table above illustrates five annotated units from the last convolutional layer of our reference network, with associated categories from the Breast Imaging-Reporting and Data System (BI-RADS) lexicon. Even though the CNN presented in this article was only trained to classify normal, benign, and malignant tissue, these internal units detect a variety of recognizable visual events. Both benign and malignant calcifications are identified, as well as features related to the margins of masses. These details are significant factors in planning interventions for breast cancer.

2.4. Annotation Efficiency. Annotating a single unit through the survey shown in Figure 2 takes approximately 30 s. As an expert only needs to annotate 60 units per network, this consumes 30 min of the expert’s time in total. This level of annotation efficiency is one of the strengths of the DeepMiner framework and should be contrasted with the standard approach to annotation which processes each image in isolation by annotating the image regions with polygons and labeling each region with a medical term. Typically, annotating a single image would require 1 min of an expert’s time, and annotating the entire DDSM data set of 10,000 images would therefore consume 10,000 min (166 hr) of expert time.

2.5. Incorporating Auxiliary Features. Genetic information is sometimes used to personalize breast cancer screening. Mutations in BRCA1 and BRCA2 are associated with an estimated lifetime risk of breast cancer of approximately 70% (Kuchenbaecker et al., 2017), and patients carrying these mutations are recommended annual contrast-enhanced breast MRI exams in addition to screening mammography as well as the option to undergo prophylactic mastectomy. Polygenic risk scores, which aggregate the risk of several individually less influential gene mutations, are routinely used in the management of detected breast cancer (Sparano et al., 2018). Similar scores have been proposed for breast cancer screening (Mavaddat et al., 2019) but are not yet routinely used. When available, these genetic scores can be incorporated into the DeepMiner classifier as auxiliary input features accompanying the mammogram. The interactions between genetic information and the radiologist interpretations can also be exposed using an attention-based mechanism. A similar approach could be applied to other risk factors, such as smoking status, age, or body mass index.

3. RESULTS

3.1. Classifying Mammogram Regions. We benchmark our reference network on the test set patches using the area under the receiver operating characteristic curve (AUC) score. Our network achieves AUCs of 0.862 for the normal class (pAUC @ TPR of 0.8 was 0.142), 0.844 for the benign class (pAUC of 0.136), and 0.872 for the malignant class (pAUC of 0.146). This performance is comparable to the state-of-the-art AUC score of 0.88 (Shen, 2017) for single network malignancy on DDSM. For comparison, positive detection rates of human radiologists range from 0.745 to 0.923 (Elmore et al., 2009). Note that achieving state-of-the-art performance for mammogram classification is not the focus of this work.

3.2. Explanation for Predictions. Using the DeepMiner framework, we next create explanations for the classifications of our reference network on the test set. Figures 4 and 5 show sample DeepMiner explanations for malignant and benign classifications, respectively. In these figures, the left-most image is the original mammogram with the benign or malignant lesion outlined in maroon. The ground truth radiologist’s report from the DDSM data set is printed beneath each mammogram. The heatmap directly on the right of the original mammogram is the class activation map for the detected class.

In Figures 4 and 5, the four or five images on the right-hand side show the activation maps of the units most influential to the prediction. In all explanations, the DeepMiner explanation units are among the top eight most influential units overall, but we only print up to five units that have been annotated as part of the explanation.

In these examples, the DeepMiner explanation gives context and depth to the final classification. For the true positive classifications in Figures 4a and 5a, the explanation further describes the finding in a manner consistent with a detailed BI-RADS report. For the false positive cases in Figures 4b and 5b, the explanation helps to identify why the network is confused or what conflicting evidence there was for the final classification.

To test whether the DeepMiner explanations enable an expert to better distinguish malignant cases from benign cases, we carried out the following human-in-the-loop experiment. We selected 165 benign cases uniformly at random from all benign and benign-without-callback cases in the test set and 165 malignant cases uniformly at random from all malignant cases in the test set. For each of the resulting $n = 330$ cases, we outputted the expert annotations associated with the three units most influential in the classification decision for that case. An expert was then tasked with classifying each case as ‘malignant’ or ‘benign’ based only on the knowledge provided by our generated explanations. As an example, Table 1 displays the first three cases presented to the expert. Equipped with only the DeepMiner explanations, the expert correctly classified 182 of the cases, a significant improvement over the baseline of random guessing (the associated p -value was 0.0346 for the exact test of a binomial proportion exceeding 0.5).

Table 1. DeepMiner explanations presented for the first three cases in the malignant vs. benign expert classification experiment of Section 3.2.

DeepMiner Explanation
(‘Calcified vessels’, ‘Calcified vessels’, ‘Spiculation’)
(‘Calcified vessels’, ‘Edge of the mass, speculations associated with cancer’, ‘Malignant pleomorphic calcifications’)
(‘Spiculation’, ‘Calcified vessels’, ‘Normal Breast Tissue’)

(A) The mammogram above is labeled Breast Imaging-Reporting and Data System (BI-RADS) assessment 4 (high risk), Digital Database for Screening Mammography (DDSM) subtlety 2 (not obvious). Our network correctly classifies the mammogram as containing malignancy. Then, DeepMiner shows the most influential units for that classification, which correctly identify the finding as a mass with spiculations.

(B) This mammogram is falsely classified by our network as containing a malignant mass, when it in fact contains a benign mass. However, the DeepMiner explanation lists the most influential unit as detecting calcified vessels, a benign finding, in the same location as the malignant class activation map. The most influential units shown here help explain how the network both identifies a benign event and misclassifies it as a malignant event.

Figure 4. Malignant DeepMiner explanations: Sample DeepMiner explanations of mammograms classified as malignant. Best viewed in color.

4. CONCLUSION

We proposed the DeepMiner framework, which uncovers interpretable representations in deep neural networks and builds explanations for deep network predictions. We trained a network for mammogram classification and showed with human expert annotation that interpretable units emerge to detect different types of medical phenomena even though the network is trained using only coarse labels. We further used the expert annotations to automatically build explanations for final network predictions. We believe our proposed framework is applicable to many other domains, potentially enabling discovery of previously unknown discriminative visual features relevant to medical diagnosis.

Disclosure Statement. The authors have no conflicts of interest to declare.

(A) The above image sequence explains a true positive classification of a benign mammogram. The benign mass is quite small, but several unit detectors identify the location of the true finding as ‘mass with smooth edges’ (likely benign) and ‘large isolated calcification.’

(B) The above image sequence shows a false positive for benign classification. The mammogram actually contains a malignant calcification. However, the fifth most influential unit detected a ‘mass with calcification clusters inside [...] very suspicious’ just below the location of the ground truth finding.

Figure 5. Benign DeepMiner explanations: Sample DeepMiner explanations of mammograms classified as benign. Best viewed in color.

ACKNOWLEDGMENTS

We thank the editor and anonymous referees for their role in improving this manuscript.

APPENDIX A. AN INTRODUCTION TO CONVOLUTIONAL NEURAL NETWORKS

Convolutional neural networks (CNNs) are a popular method for extracting high-level information from images. In this appendix, we will provide a concise and targeted introduction to CNNs and not a complete reference. Interested readers can consult Goodfellow et al. (2016) for more information.

CNNs apply multiple layers of processing to the original image. We will denote each layer as $f_l(x, y, c)$, where x and y denote 2D Cartesian image coordinates, l denotes the layer index, and c denotes the channel index. The first layer, $f_1(x, y, c)$, is the original image. Most images produced using digital cameras provide three channels for the RGB (red, green, and blue) colors. The grayscale images captured using mammography do not have colors, so the single-channel grayscale is replicated across the three color channels for the first layer. For an N -layer CNN, subsequent

convolutional layers are defined as follows:

$$f_l(x, y, c) = h(\sum_j f_{l-1}(x, y, j) \otimes n_{lc}(x, y, j))$$

The $\otimes$ represents the convolution operator, and the $n_{lc}(x, y, j)$ are the filter weights for the l^{th} layer and c^{th} channel. These filter weights are typically nonzero only in a limited region. Many architectures specify that the filter weights are 3×3 or 5×5 in size, making the convolution computationally efficient. Here, $h(x)$ is the nonlinear activation function. A popular choice for $h(x)$ is the rectified linear (ReLU) function, $h(x) = \max(x, 0)$.

In a binary classification task that distinguishes cancer patches from noncancer patches, to produce an estimate of the probability of cancer, the contents of the final convolutional layer $f_N(x, y, c)$ are rearranged into a single one-dimensional vector, $\tilde{f}_N(t)$. We then compute the cancer probability using the sigmoid function:

$$p_{\text{cancer}} = (1 + e^{-w^\top \tilde{f}_N})^{-1}$$

Here, w is a weighting vector and $w^\top$ is its transpose so that $w^\top \tilde{f}_N$ is the dot product between w and $\tilde{f}_N$.

We have described a simple CNN that only uses convolutional layers followed by a simple fully connected layer to calculate the cancer probability. Most CNNs used today also include more than two output classes (e.g., malignant vs. benign vs. normal) and other kinds of layers, including downsampling layers that reduce the size of the image, allowing future layers to efficiently extract information across larger spatial scales. Skip connections, which connect early layers with later layers, are also widely used, allowing efficient training of deep networks. Deep networks have found much more empirical success than shallow networks in many domains. The ResNet architecture used in this work is best known for introducing skip connections to the CNN literature.

CNNs are typically trained by presenting the network with thousands to millions of labeled datapoints. In the context of our binary classification example, each datapoint is an image patch of a mammogram that is labeled to be either positive (cancer) or negative (not cancer). For each patch, both the filter weights $n_{lc}(x, y, j)$ and the final weighting factor w are updated so that the final probability p_{cancer} is increased for cancerous patches and decreased for negative patches.

Several algorithms are available to perform the training. Most of the algorithms are a variation on gradient descent, also called the Newton-Raphson method. Stochastic gradient descent (SGD) is one of the most popular choices and uses a limited number of datapoints selected uniformly at random from the full data set. The number of samples used for each iteration of SGD is known as the minibatch size. The use of a minibatch rather than the entire collection of samples can be viewed as a type of regularization, to prevent the network from collapsing into a local minimum. The step size in SGD must be empirically selected according to a preset schedule. Besides SGD, other optimization algorithms are available that automatically select an adaptive step size, including AdaGrad (adaptive gradient descent) or Adam (adaptive momentum). These often reduce training time, but the selection of the best optimizer for a given data set remains empirical.

CNNs often have millions of unknown parameters that must be learned. As the CNN is trained, the filter weights n_{lc} and the final weighting vector w are progressively improved.

REFERENCES

- Bau, D., Zhou, B., Khosla, A., Oliva, A., & Torralba, A. (2017). Network dissection: Quantifying interpretability of deep visual representations. *CVPR*. <https://doi.org/10.1109/CVPR.2017.354>

- Bottou, L., Curtis, F. E., & Nocedal, J. (2018). Optimization methods for large-scale machine learning. *Siam Review*, 60(2), 223–311. <https://doi.org/10.1137/16M1080173>
- Elmore, J. G., Jackson, S. L., Abraham, L., Miglioretti, D. L., Carney, P. A., Geller, B. M., Yankaskas, B. C., Kerlikowske, K., Onega, T., Rosenberg, R. D., et al. (2009). Variability in interpretive performance at screening mammography and radiologists’ characteristics associated with accuracy. *Radiology*, 253(3), 641–651. <https://doi.org/10.1148/radiol.2533082308>
- Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115–118. <https://doi.org/10.1038/nature21056>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning* [<http://www.deeplearningbook.org>]. MIT Press.
- He, K., Gkioxari, G., Dollar, P., & Girshick, R. (2017). Mask r-cnn. *Proceedings of the IEEE international conference on computer vision*, 2961–2969. <https://doi.org/10.1109/ICCV.2017.322>
- He, K., Zhang, X., Ren, S., & Sun, J. (2015a). Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*. <https://doi.org/10.1109/CVPR.2016.90>
- He, K., Zhang, X., Ren, S., & Sun, J. (2015b). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *ICCV*, 1026–1034. <https://doi.org/10.1109/ICCV.2015.123>
- Heath, M., Bowyer, K., Kopans, D., Moore, R., & Kegelmeyer, W. P. (2000). The digital database for screening mammography. *Proceedings of the 5th international workshop on digital mammography*, 212–218.
- Jing, B., Xie, P., & Xing, E. (2017). On the automatic generation of medical imaging reports. *arXiv preprint arXiv:1711.08195*.
- Kiefer, J., Wolfowitz, J. et al. (1952). Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 23(3), 462–466. <https://doi.org/10.1214/aoms/1177729392>
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 1097–1105. <https://doi.org/10.1145/3065386>
- Kuchenbaecker, K. B., Hopper, J. L., Barnes, D. R., Phillips, K.-A., Mooij, T. M., Roos-Blom, M.-J., Jervis, S., van Leeuwen, F. E., Milne, R. L., Andrieu, N., Goldgar, D. E., Terry, M. B., Rookus, M. A., Easton, D. F., Antoniou, A. C., the BRCA1, & Consortium, B. C. (2017). Risks of breast, ovarian, and contralateral breast cancer for brca1 and brca2 mutation carriers. *JAMA*, 317(23), 2402–2416. <https://doi.org/10.1001/jama.2017.7112>
- Long, J., Shelhamer, E., & Darrell, T. (2015). Fully convolutional networks for semantic segmentation. *Proceedings of CVPR*. <https://doi.org/10.1109/CVPR.2015.7298965>
- Mahendran, A., & Vedaldi, A. (2014). Understanding deep image representations by inverting them. *CoRR*, abs/1412.0035. <https://doi.org/10.1109/CVPR.2015.7299155>
- Mavaddat, N., Michailidou, K., Dennis, J., Lush, M., Fachal, L., Lee, A., Tyrer, J. P., Chen, T.-H., Wang, Q., Bolla, M. K., et al. (2019). Polygenic risk scores for prediction of breast cancer and breast cancer subtypes. *The American Journal of Human Genetics*, 104(1), 21–34. <https://doi.org/10.1016/j.ajhg.2018.11.002>
- Orel, S. G., Kay, N., Reynolds, C., & Sullivan, D. C. (1999). Bi-rads categorization as a predictor of malignancy. *Radiology*, 211(3), 845–850.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., ... Chintala, S. (2019). Pytorch: An imperative style,

- high-performance deep learning library. *Advances in neural information processing systems* 32 (pp. 8024–8035). Curran Associates, Inc.
- Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., Ding, D., Bagul, A., Langlotz, C., Shpanskaya, K., et al. (2017). Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint arXiv:1711.05225*.
- Robbins, H., & Monro, S. (1951). A stochastic approximation method. *The annals of mathematical statistics*, 400–407. <https://doi.org/10.1214/aoms/1177729586>
- Shen, L. (2017). End-to-end training for whole image breast cancer diagnosis using an all convolutional design. *arXiv preprint arXiv:1708.09427*.
- Simonyan, K., Vedaldi, A., & Zisserman, A. (2014). Deep inside convolutional networks: Visualising image classification models and saliency maps. *Workshop at International Conference on Learning Representations*.
- Sparano, J. A., Gray, R. J., Makower, D. F., Pritchard, K. I., Albain, K. S., Hayes, D. F., Geyer, C. E., Dees, E. C., Goetz, M. P., Olson, J. A., Lively, T., Badve, S. S., Saphner, T. J., Wagner, L. I., Whelan, T. J., Ellis, M. J., Paik, S., Wood, W. C., Ravdin, P. M., ... Sledge, G. W. (2018). Adjuvant chemotherapy guided by a 21-gene expression assay in breast cancer [PMID: 29860917]. *New England Journal of Medicine*, 379(2), 111–121. <https://doi.org/10.1056/NEJMoa1804710>
- Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., & Summers, R. M. (2017). Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3462–3471. <https://doi.org/10.1109/CVPR.2017.369>
- Wu, J., Peck, D., Hsieh, S., Dialani MD, V., Lehman MD, C. D., Zhou, B., Syrgkanis, V., Mackey, L., & Patterson, G. (2018). Expert identification of visual primitives used by CNNs during mammogram classification. *SPIE Medical Imaging*. <https://doi.org/10.1117/12.2293890>
- Wu, J., Zhou, B., Peck, D., Hsieh, S., Dialani, V., Mackey, L., & Patterson, G. (2021). Replication Data for: DeepMiner: Discovering Interpretable Representations for Mammogram Classification and Explanation. <https://doi.org/10.7910/DVN/U39HOQ>
- Zeiler, M. D., & Fergus, R. (2014). Visualizing and understanding convolutional networks. *European conference on computer vision*, 818–833. https://doi.org/10.1007/978-3-319-10590-1_53
- Zhang, Z., Xie, Y., Xing, F., McGough, M., & Yang, L. (2017). Mdnnet: A semantically and visually interpretable medical image diagnosis network. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. <https://doi.org/10.1109/CVPR.2017.378>
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., & Torralba, A. (2016). Learning deep features for discriminative localization. *Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference on*, 2921–2929. <https://doi.org/10.1109/CVPR.2016.319>
- Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., & Oliva, A. (2014). Learning deep features for scene recognition using places database. *Advances in neural information processing systems*, 487–495.