

GENERALIZATION OF DEEP NEURAL NETWORKS FOR CHEST PATHOLOGY CLASSIFICATION IN X-RAYS USING GENERATIVE ADVERSARIAL NETWORKS

Hojjat Salehinejad^{*†}, Shahrokh Valaee^{*}, Tim Dowdell[†], Errol Colak[†], and Joseph Barfett[†]

^{*}Department of Electrical & Computer Engineering, University of Toronto, Toronto, Canada

[†]Department of Medical Imaging, St. Michael's Hospital, University of Toronto, Toronto, Canada
salehinejadh@smh.ca, valaee@ece.utoronto.ca, {dowdellt,colake,barfettj}@smh.ca

ABSTRACT

Medical datasets are often highly imbalanced with overrepresentation of common medical problems and a paucity of data from rare conditions. We propose simulation of pathology in images to overcome the above limitations. Using chest X-rays as a model medical image, we implement a generative adversarial network (GAN) to create artificial images based upon a modest sized labeled dataset. We employ a combination of real and artificial images to train a deep convolutional neural network (DCNN) to detect pathology across five classes of chest X-rays. Furthermore, we demonstrate that augmenting the original imbalanced dataset with GAN generated images improves performance of chest pathology classification using the proposed DCNN in comparison to the same DCNN trained with the original dataset alone. This improved performance is largely attributed to balancing of the dataset using GAN generated images, where image classes that are lacking in example images are preferentially augmented.

Index Terms— Chest X-ray, data augmentation, deep convolutional neural network (DCNN), generative adversarial network (GAN), simulated images.

1. INTRODUCTION

In the medical domain, preservation of patient privacy is paramount, and hence access to data is often intrinsically limited to research groups [1], [2]. Medical datasets, similar to financial [3] and genomics [4] datasets, are also very limited because they are often imbalanced [5]. Such imbalances in datasets potentially make the training of neural networks with equally high accuracy across classes technically challenging. Some medical problems are commonly encountered in hospital settings which leads to a substantial amount of data associated with them. However, rare conditions or syndromes such as Birt-Hogg-Dube syndrome are expected to have limited amounts of data in clinical databases [6].

The challenge of image availability across classes may be partially met by data augmentation techniques, for example,

applying transformations to images to augment the dataset. The importance of balancing datasets is highlighted by the fact that deep neural networks may be most valuable in the work up of rare or challenging diseases, which practitioners at a common skill level may fail to recognize or misinterpret [7]. Generative adversarial networks (GANs) have shown to effectively generate artificial data indiscernible from their real counterparts [8]. Some examples are statistical parametric speech synthesis [9], learning representations of emotional speech [10], noise reduction in low-dose computed tomography (CT) [11], and retinal image synthesis [12].

We propose the simulation of medical pathology in images as a means of augmenting data in a controlled fashion. Simulated data can be used to increase the number of images available and hence provide a means to balance datasets for the training of deep neural networks. An ideal data simulation scheme would be capable of generating an arbitrary number of synthetic images, which mimic the features of real images in any given class with sufficient diversity for the successful training of a deep network. In this manuscript, we propose the use of a deep convolutional generative adversarial network (DCGAN) for the generation of chest X-rays that mimic common chest pathologies. The synthetic images are used to balance and augment a labeled set of chest X-rays for the training of the proposed deep convolutional neural network (DCNN) across five pathological classes.

2. PROPOSED METHOD

We propose a DCGAN tailored for training with chest X-rays as in Figure 1. The generated artificial chest X-rays are concatenated with the real X-rays to balance and expand the training dataset to train the proposed DCNN. This DCNN performs chest pathology classification as is demonstrated in Figure 2. These models are discussed in details below.

2.1. Generating Chest X-Rays

GANs are composed of two neural networks, a Generator G and a Discriminator D , which compete with each other over the available training data to improve their performance. The

The authors thank the support of NVIDIA Corporation with the donation of the Titan X GPUs used for this project.




Fig. 1: Architecture of the DCGAN and training it with real chest X-rays.

DCGAN generates chest X-rays using DCNNs for both the G and D components of the model [13].

The trained G , models the underlying probability distribution \mathbf{p}_g of the training data for the set of exported chest X-rays \mathbf{x} and proposes artificial mappings $G(\mathbf{z}, \theta_g)$ from the prior input noise variable $p_z(\mathbf{z})$, where θ_g is the set of learning parameters of the DCNN in the Generator. As demonstrated in Figure 1, a 128 dimensional vector \mathbf{z} such that $z_i \sim \text{uniform}(-1, 1)$ is projected to a spatially extended convolutional representation with 1,024 feature maps. Since chest X-rays contain subtle features, high resolution images are mostly of interest for machine learning purposes in medical imaging. Therefore, a series of six fractionally-strided convolutions (instead of four convolutions [13]) convert the projected and reshaped noise vector \mathbf{z} into a 256×256 pixel chest X-ray $\hat{\mathbf{x}}$.

The Discriminator network $D(\mathbf{x}, \theta_d)$ receives a generated image $\hat{\mathbf{x}}$ or a real chest X-ray \mathbf{x} and after passing that through six convolution layers, as presented in Figure 1, produces an output \hat{o} , stating whether the input image is real or synthesized such that

$$\hat{y} = \frac{1}{1 + e^{-\hat{o}}} \text{ s.t. } \hat{y} \in [0, 1] \quad (1)$$

where $\hat{y} \rightarrow 0$ and $\hat{y} \rightarrow 1$ state that the input chest X-ray is synthesized or real, respectively. The Generator network G trains so as to propose artificial images that the Discriminator network $D(\mathbf{x})$ cannot distinguish from real images. The adversarial competition between G and D can be represented as




Fig. 2: Architecture of the DCNN and its training with real and generated chest X-rays to classify abnormalities.

$$\min_G \max_D \mathcal{L}(D, G) = \mathbb{E}_{\mathbf{x} \sim \mathbf{p}_{data}} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim \mathbf{p}_z(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))] \quad (2)$$

where the Discriminator maximizes the loss value $\mathcal{L}(D, G)$ while the Generator tries to minimize it. To train D , G works in a feed forward fashion without back propagation, and vice versa to train G [14].

2.2. Classification of Chest X-Rays

The trained Generator of the discussed DCGAN is used to generate artificial chest X-rays as demonstrated in Figure 2. The shuffled concatenation of real and synthesized chest X-rays is then fed into the proposed DCNN for detection and classification of pathology in chest X-rays.

AlexNet is a successful DCNN architecture that is composed of five convolutional layers for feature extraction followed by three fully-connected layers for classification [15]. The proposed DCNN for chest pathology classification in this paper is fundamentally similar to AlexNet, however, uses different kernel sizes, feature map sizes, and convolution layers as illustrated in Figure 2. For an input chest X-ray $\mathbf{x}_{256 \times 256}$, a convolution kernel of size 5×5 performs the convolution operation to generate the feature map m as

$$h_{u,v}^{(m)} = \sigma \left(\sum_{i=0}^4 \sum_{j=0}^4 x_{i,j} \cdot w_{i,j}^{(m)} + b_{u,v}^{(m)} \right), \quad (3)$$

for $u \in \{0, \dots, 255\}$ and $v \in \{0, \dots, 255\}$, where $\sigma(\cdot)$ is the rectified linear unit (ReLU) activation function, $\mathbf{b}^{(m)}$ is the bias vector, and $\mathbf{w}^{(m)}$ is the weight matrix. The ReLU is




Fig. 3: Samples of real (R) and synthesized (S) chest X-rays: (a) Pulmonary Edema-R; (b) Pulmonary Edema-S; (c) Normal-R; (d) Normal-S; (e) Pneumothorax-R; (f) Pneumothorax-S; (g) Cardiomegaly-R; (h) Cardiomegaly-S; (i) Pleural Effusion-R; (j) Pleural Effusion-S. The white arrow points to the pathologic condition.

defined as $f(x) = \max(0, x)$, which takes advantage of its non-saturating and non-linear properties as well as the tendency to enable more efficient learning than \tanh or sigmoid activation functions [15]. A max-pooling layer after the convolution layer down-samples the latent representation by a constant factor, usually taking the maximum value over non-overlapping sub-regions such as

$$O_{i,j} = \max\{h_{q,r}^{(m)}\}, \quad (4)$$

for $q, r \in \{(2i, 2j), (2i+1, 2j), (2i, 2j+1), (2i+1, 2j+1)\}$, where the max-pooling kernel is a square with side length of $L = 2$. This operation helps to obtain translation-invariant representations. The max-pooled features, after the multiple convolutional and max-pooling layers, are reshaped as a vector, \mathbf{f} , and are fed to a multi-layer perceptron (MLP) network with one hidden layer as illustrated in Figure 2. The output layer of MLP for the input vector \mathbf{f} is

$$y_c = \phi\left(\sum_{j=0}^{|\mathbf{f}|-1} f_j \cdot w_{j,c} + b_c\right) \quad (5)$$

where $|\mathbf{f}| = 4,096$ is the length of vector \mathbf{f} and \mathbf{w} is the weight of connections between the layers containing \mathbf{f} and the output layer \mathbf{y} . The *softmax* function $\phi(\cdot)$ assigns a probability to each output unit, which corresponds to each class c , such as

$$\phi(h_c) = \frac{e^{h_c}}{\sum_{j=1}^P e^{h_j}} \quad (6)$$

where $c \in \{\text{Pneumothorax, Pulmonary Edema, Pleural Effusion, Normal, Cardiomegaly}\}$, h_c is the input to the output layer, and $C = 5$ is the number of chest X-ray classes [16].

3. EXPERIMENTS

We discuss the details of implementation and our obtained results in this section.

3.1. Data

With the approval from our institutional ethics review board, search of our hospital's Radiology Information System (RIS) was undertaken using the Montage Search and Analytics engine. The dataset contained 15,781 Normal exams, 17,098 examples of Cardiomegaly, 14,510 Pleural Effusions, 5,018 examples of Pulmonary Edema, and 4,013 examples of Pneumothorax. These images were exported and anonymized as

PNG files after down-sampling to 256×256 pixels to have a balance between preserving resolution and computational complexity (i.e. number of free parameters) of the models. For all the experiments, 1,000 real chest X-rays with equal contribution from each class are selected for validation and the same number for testing of the model.

3.2. Technical Details of Training

A team of radiologists removed inappropriate generated chest X-rays from the respective class directories and in order to keep the training dataset balanced across different classes, the DCGAN was trained with a dataset of 2,000 chest X-rays per class. The parameters of the models were as follows: the mini-batch size was 64, the number of training iterations was 20, an Adam optimizer was implemented with adaptive learning rate, starting between 2×10^{-4} and 2×10^{-3} depending on the dataset size, and momentum of 0.5. The proposed DCNN for pathology classification was trained over 100 iterations using a mini-batch size of 128, an Adam optimizer with sigmoid decay adaptive learning rate starting at 1×10^{-3} and a momentum term of 0.5. The weights of the convolutional layers were selected using a normal distribution and biases of 0.1. The ReLU activation functions were implemented before the max-pooling layers. L_2 regularization was set to 1×10^{-4} with early-stopping. Cross validation was performed 10 times prior to reporting results.

3.3. Performance Evaluation

Generated images were evaluated by qualitative and quantitative means. First, a board certified radiologist reviewed the images for features appropriate to the five defined classes. Second, the proportion of real and artificial chest X-rays were varied to create concatenated datasets and the results of the trained DCNN model were assessed on the test data.

3.3.1. Qualitative Evaluation by a Radiologist

Real and artificial images were presented to a radiologist as shown in Figure 3 and were visually examined. The artificial and real chest X-rays show similar characteristics, though the synthetic ones are of comparatively low resolution. For example in Figure 3 we demonstrate increased attenuation of the lung parenchyma in an example of pulmonary edema, a

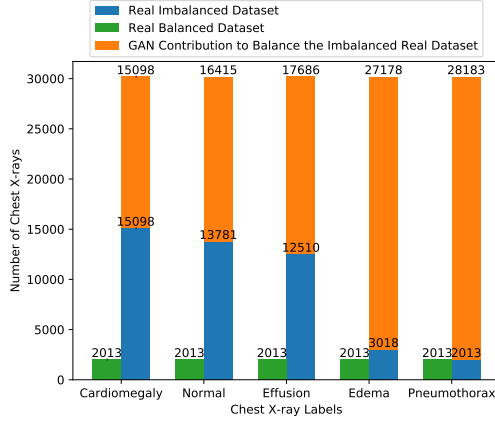


Fig. 4: Number of samples for imbalanced dataset of real chest X-rays, balanced dataset of real chest X-rays, and balanced dataset of real chest X-rays using DCGAN synthesized chest X-rays. The total number of chest X-rays per label in the balanced dataset is the summation of corresponding blue and orange bars (i.e. 30,196).

pleural line in the case of pneumothorax, an enlarged cardiac silhouette in the case of cardiomegaly, as well as bilateral pleural effusions. The generated images convincingly demonstrate the range of pathology under examination and can be classified for the training of a deep neural network.

3.3.2. Chest Pathology with DCGAN Generated X-Rays

Deep neural networks trained with a combination of real and artificial data have potential advantage over networks trained with real data alone, including a larger quantity of data as well as a better diversified dataset. Our original hospital dataset of chest X-rays, like many clinical datasets, is highly imbalanced and dominated by normal cases and common clinical conditions. Rare conditions by virtue of their low prevalence are underrepresented in such datasets.

We trained the proposed DCNN using a real imbalanced dataset (DS1), real balanced dataset (DS2), and augmented dataset with DCGAN synthesized chest X-rays to balance the imbalanced real dataset (DS3). As Figure 4 shows, each image class in the balanced dataset had 30,196 samples, which was twice the maximum number of available samples in the real imbalanced dataset (i.e. for Cardiomegaly). The accuracy of predictions for the DCNN is presented in Table 1. The augmented dataset DS3 outperforms the original dataset across all classes. A mean classification accuracy of 92.10% was achieved using the proposed method, where data was augmented by DCGAN generated images to train the proposed DCNN. This performance is more than 20% higher than that obtained by the same DCNN architecture trained only by the original data. The greatest performance improvement was seen in the pneumothorax class, which was the class with the fewest native images.

Accuracy and standard deviation of the proposed DCNN on the validation dataset for training iterations over DS1,

Table 1: Classification accuracy of the DCNN is improved by balancing the original dataset with DCGAN generated chest X-rays. DS1: Imbalanced dataset of real radiographs; DS2: Balanced dataset of real X-rays; DS3: Balanced dataset of real images with synthesized X-rays.

Accuracy (%)	DS1	DS2	DS3
Cardiomegaly	79.15	71.73	95.31
Normal	77.75	72.53	95.02
Pleural Effusion	73.64	51.23	91.19
Pulmonary Edema	65.86	50.12	89.68
Pneumothorax	57.99	48.92	88.84
Total	70.87±0.47	58.90±0.48	92.10±0.41




Fig. 5: Accuracy (solid line) and standard deviation (shaded region) of the DCNN on the validation dataset during training iterations over the DS1, DS2, and DS3 datasets.

DS2, and DS3 datasets are presented in Figure 5. The DCNN trained with DS3 has achieved higher accuracy comparing with the DS1 and DS2. The plot show that the DCNN trained on DS3 has almost converged after 90 iterations. The same DCNN trained on DS1 and DS2 is overfitted after 22 and 31 iterations, respectively, to a lower accuracy. The plot clearly shows effectiveness of the added diversity to the dataset by the generated chest X-rays using the DCGAN, which helps to improve generalization performance of the DCNN for chest pathology classification in X-rays and avoid over-fitting.

4. CONCLUSION

In this paper, we have shown that artificial data generated by DCGAN can augment real datasets to provide both a greater quantity of data for training of large neural networks and can balance the dataset, resulting in substantial improvement in classification performance in the most anemic classes. We obtained best results with a combination of real and artificial data used to train the DCNN. The reported results in this paper suggest that data augmentation using synthesized images increases the diversity of training dataset and therefore, improves generalization performance of deep learning for classification of unseen data.

5. REFERENCES

- [1] Elisa Bertino, Beng Chin Ooi, Yanjiang Yang, and Robert H Deng, "Privacy and ownership preserving of outsourced medical data," in *Data Engineering, 2005. ICDE 2005. Proceedings. 21st International Conference on*. IEEE, 2005, pp. 521–532.
- [2] Hojjat Salehinejad, Shahrokh Valaee, Tim Dowdell, and Joseph Barfett, "Image augmentation using radial transform for training deep neural networks," in *Acoustics, Speech and Signal Processing (ICASSP), 2018 IEEE International Conference on*. IEEE, 2018.
- [3] Hojjat Salehinejad and Shahryar Rahnamayan, "Customer shopping pattern prediction: A recurrent neural network approach," in *Computational Intelligence (SSCI), 2016 IEEE Symposium Series on*. IEEE, 2016, pp. 1–6.
- [4] Farhad Pouladi, Hojjat Salehinejad, and Amir Mohammad Gilani, "Recurrent neural networks for sequential phenotype prediction in genomics," in *Developments of E-Systems Engineering (DeSE), 2015 International Conference on*. IEEE, 2015, pp. 225–230.
- [5] Der-Chiang Li, Chiao-Wen Liu, and Susan C Hu, "A learning method for the class imbalance problem with medical data sets," *Computers in biology and medicine*, vol. 40, no. 5, pp. 509–518, 2010.
- [6] Michael L Nickerson, Michelle B Warren, Jorge R Toro, Vera Matrosova, Gladys Glenn, Maria L Turner, Paul Duray, Maria Merino, Peter Choyke, Christian P Pavlovich, et al., "Mutations in a novel gene lead to kidney tumors, lung wall defects, and benign tumors of the hair follicle in patients with the birt-hogg-dube syndrome," *Cancer cell*, vol. 2, no. 2, pp. 157–164, 2002.
- [7] Maciej A Mazurowski, Piotr A Habas, Jacek M Zurada, Joseph Y Lo, Jay A Baker, and Georgia D Tourassi, "Training neural network classifiers for medical decision making: The effects of imbalanced datasets on classification performance," *Neural networks*, vol. 21, no. 2, pp. 427–436, 2008.
- [8] Weixuan Tang, Shunquan Tan, Bin Li, and Jiwu Huang, "Automatic steganographic distortion learning using a generative adversarial network," *IEEE Signal Processing Letters*, vol. 24, no. 10, pp. 1547–1551, 2017.
- [9] Takuhiro Kaneko, Hirokazu Kameoka, Nobukatsu Hojo, Yusuke Ijima, Kaoru Hiramatsu, and Kunio Kashino, "Generative adversarial network-based postfilter for statistical parametric speech synthesis," in *Proc. 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP2017)*, 2017, pp. 4910–4914.
- [10] J. Chang and S. Scherer, "Learning representations of emotional speech with deep convolutional generative adversarial networks," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 2746–2750.
- [11] Jelmer M Wolterink, Tim Leiner, Max A Viergever, and Ivana Isgum, "Generative adversarial networks for noise reduction in low-dose ct," *IEEE Transactions on Medical Imaging*, 2017.
- [12] Pedro Costa, Adrian Galdran, Maria Ines Meyer, Meindert Niemeijer, Michael Abràmoff, Ana Maria Mendonça, and Aurélio Campilho, "End-to-end adversarial retinal image synthesis," *IEEE Transactions on Medical Imaging*, 2017.
- [13] Alec Radford, Luke Metz, and Soumith Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [15] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [16] Mark Cicero, Alexander Bilbily, Errol Colak, Tim Dowdell, Bruce Gray, Kuhan Perampaladas, and Joseph Barfett, "Training and validating a deep convolutional neural network for computer-aided detection and classification of abnormalities on frontal chest radiographs," *Investigative radiology*, vol. 52, no. 5, pp. 281–287, 2017.