

Using deep learning for detecting gender in adult chest radiographs

Zhiyun Xue, Sameer Antani, L. Rodney Long, George R. Thoma
National Library of Medicine, National Institutes of Health, Bethesda, MD

ABSTRACT

In this paper, we present a method for automatically identifying the gender of an imaged person using their frontal chest x-ray images. Our work is motivated by the need to determine missing gender information in some datasets. The proposed method employs the technique of convolutional neural network (CNN) based deep learning and transfer learning to overcome the challenge of developing handcrafted features in limited data. Specifically, the method consists of four main steps: pre-processing, CNN feature extractor, feature selection, and classifier. The method is tested on a combined dataset obtained from several sources with varying acquisition quality resulting in different pre-processing steps that are applied for each. For feature extraction, we tested and compared four CNN architectures, viz., AlexNet, VggNet, GoogLeNet, and ResNet. We applied a feature selection technique, since the feature length is larger than the number of images. Two popular classifiers: SVM and Random Forest, are used and compared. We evaluated the classification performance by cross-validation and used seven performance measures. The best performer is the VggNet-16 feature extractor with the SVM classifier, with accuracy of 86.6% and ROC Area being 0.932 for 5-fold cross validation. We also discuss several misclassified cases and describe future work for performance improvement.

Keywords: Chest radiograph, Annotation, Classification, Gender

1. INTRODUCTION

Automatic labeling or annotation for medical images is an important research topic. The annotated tags can be used in the fields of medical image analysis, diagnosis, retrieval/search, classification, and visual question answering [1]. Among medical images, chest radiographs (CXR) are a common imaging tool used to diagnose diseases in and around the thorax. There has been recent interest in developing automatic computer algorithms for detecting disease in CXRs. Some of our prior work has been developing algorithms to locate the major organs in a digital CXR, such as lungs [2], heart [3], and ribs [4], or if the CXR view is frontal or lateral [5]. While these are also useful for annotating CXRs for image retrieval or visual question-answering, we have used this step as a preprocessing to detect presence of pulmonary abnormalities, with a particular focus on detecting manifestations consistent with tuberculosis infection [6]. In this paper, we will present our work on identifying gender from the CXR image.

Chest x-ray image datasets can have images in the clinical DICOM format, or in common image formats (JPEG, PNG, etc.). In contrast to JPEG, PNG (and several other image) formats, DICOM images have image headers that permit inclusion of gender information. However, a lack of proper acquisition quality-assurance (QA) procedures or aggressive de-identification processes often result in images absent gender data. For example, this information is not available in the Indiana University chest x-ray collection available in U.S. National Library of Medicine's (NLM) open-access biomedical image search engine Open-i® [7]. Therefore, the goal of this work is to automatically detect the gender of the imaged person given their frontal chest x-ray. To this end, we create a dataset which consists of 2066 images combined from several datasets (some of which are publicly available) in which the gender information is available.

In recent years, there has been increased interest in applying deep learning, especially using convolutional neural networks (CNN), to the medical image domain. Deep learning can be used to automatically extract complex discriminative descriptors (features) through a learning procedure and reduces the need for developing handcrafted features. However, to achieve good performance, it may require a significant amount of labeled/annotated data which is much more challenging to get for medical images than for general images. To deal with this challenge while still taking advantage of deep learning, the technique of transfer learning has been applied for small labeled medical image datasets [8, 9]. In this technique, the CNNs which have been trained using another large scale, labeled, general image dataset (such as ImageNet [10]) can be used as a feature extractor for the images in the small dataset. Then these extracted CNN-based features and corresponding

labels of these images are used to train/test a conventional classifier. We use this approach for our gender identification in the chest x-rays application.

In the following sections, we first describe the image data we use and introduce the visual hints for gender differentiation. We then present the method for feature extraction, feature selection and classification, and discuss the experimental testing and results. At last, we point directions for future work and conclude the paper.

2. IMAGE DATA

In this work, we combine chest radiographs from four datasets (which have the gender information) to generate one dataset. We use this mixed dataset to test and quantitatively evaluate our method. The four datasets are the Shenzhen hospital dataset, the Montgomery county dataset [11, 12], the Kenya dataset, and the JSRT dataset [13]. As a result, the dataset used for experimental testing and method evaluation has large variance. It contains images obtained from different countries, different machines, and may have undergone with different ways of intensity conversion (for example, from HU to regular range). We also apply our method to the Indiana University dataset. Since, the Indiana University dataset does not have gender information, we cannot quantitatively evaluate the results on it. Below are brief introductions to each individual dataset.

Shenzhen Hospital Dataset: The chest radiographs in this dataset were collected by Shenzhen No.3 Hospital in Shenzhen, Guangdong province, China, as part of routine care. The dataset is mainly used for research on computer aided screening/diagnosis of tuberculosis (TB). It has been made publicly available by NLM [11, 12]. The set consists of a total of 662 frontal images of which 326 are normal x-rays, and 336 show various manifestations of TB. The images are in PNG format. In addition to the normal/abnormal reading, the dataset also provides information on the patient's gender and age for each image in a text file.

Montgomery County Dataset: The Montgomery county dataset [11, 12] is also made available to the public by the NLM. X-ray images in this set have been acquired from the tuberculosis control program of the Department of Health and Human Services of Montgomery County, MD, USA. This set contains 138 posterior-anterior x-rays, of which 80 x-rays are normal, and 58 x-rays are abnormal with manifestations of TB. All images are de-identified and available in DICOM format. Similar to the Shenzhen Hospital dataset, this dataset includes radiology readings as well as patient's sex and age in the text files. The information on patient's sex and age is also available in the DICOM image header.

Kenya Dataset: In collaboration with Indiana University School of Medicine and AMPATH (Academic Model Providing Access to Healthcare), NLM has collected de-identified x-rays from western Kenya, Africa. At present, the dataset contains 729 normal and 290 abnormal frontal chest radiographs with detailed radiologists' readings. The images are in DICOM format and de-identified. The DICOM image header contains patient's sex information but no age information. This image set is not publicly available due to restrictions placed by the data provider.

JSRT Dataset: This dataset was created by the Japanese Society of Radiological Technology in cooperation with the Japanese Radiological Society [13]. It contains 247 frontal x-rays, among them, 154 images having lung nodules and 93 images having no nodules. The images are provided in raw data (12 bits unsigned) with no header. Information such as patient age, gender, and nodule location are given in text files.

Indiana University Dataset: The dataset consists of approximately 4000 radiology reports and 8300 images (include both frontal and lateral views). The data was collected at various hospitals across the state of Indiana, USA. These images, as well as corresponding radiology report information, have been integrated into NLM Open-i system. The images are available to the public (in both PNG and DICOM formats) [14]. There is no information on patient's gender in this dataset due to aggressive de-identification processes applied to the image headers.

We combine all the frontal CXR images that have a gender label in the first four datasets (Shenzhen Hospital, Montgomery County, Kenya, and JSRT) into one dataset. We quantitatively evaluate the performance of the method on this combined labeled dataset. Table 1 shows the number of images for each gender (female vs. male) for each individual dataset and the

combined dataset. The total number of females in the combined dataset is 1097 while the total number of males is 969. We also apply our method to the frontal images in the Indiana University dataset, as this work was partially motivated by adding gender tags on this dataset. We visually examine the classification results to get a qualitative assessment on the performance.

Table 1. Number of images in each gender category.

Dataset	Female	Male	Total
Shenzhen Hospital	213	449	662
Montgomery County	74	64	138
Kenya	682	337	1019
JSRT	128	119	247
Total	1097	969	2066

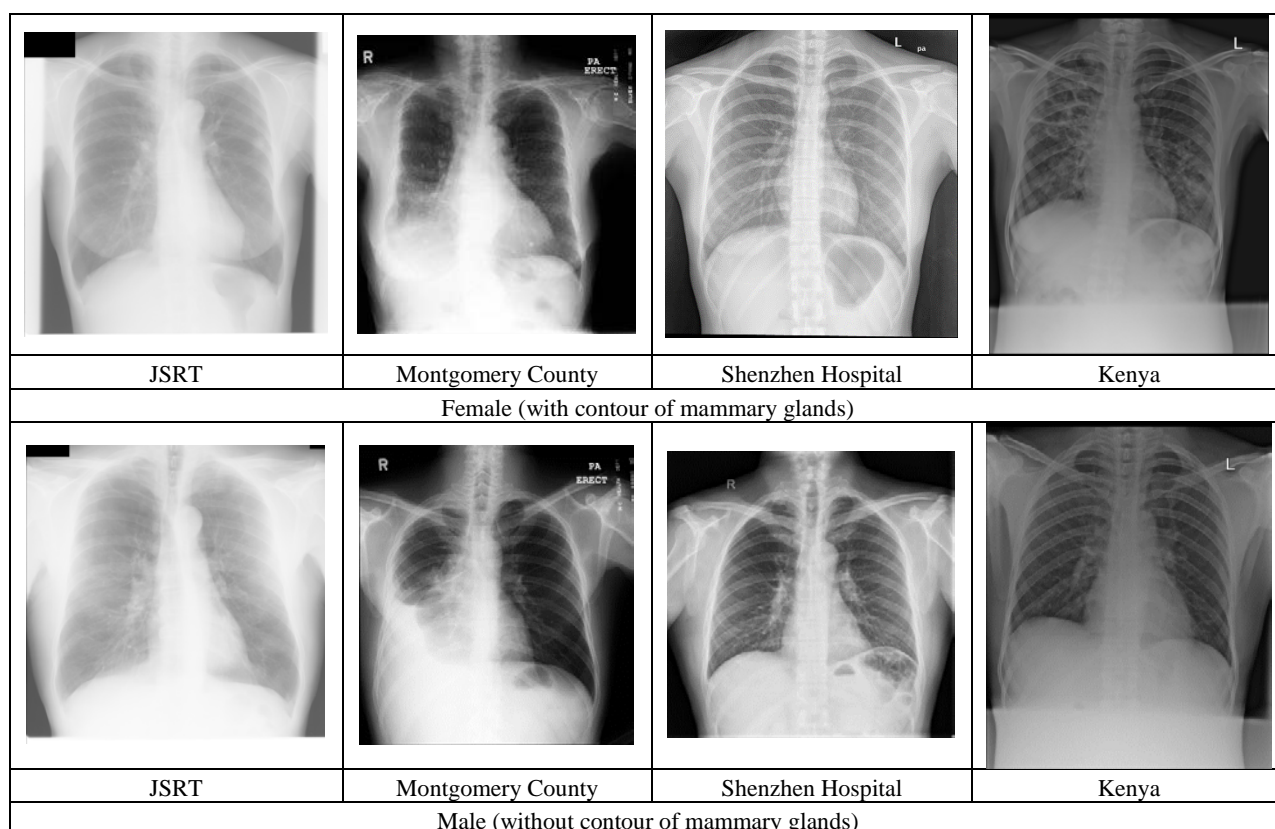


Figure 1. Chest x-rays (Female vs Male)

3. METHOD

There are a few visual difference hints that may be exhibited in the chest x-rays between adult females and males. For example, as shown in Figure 1, one significant characteristic is the presence of the breast contour in chest x-rays. The study on sex differences in thoracic dimensions and configuration in [15] also has shown, in normal adults, “females had smaller radial rib cage dimensions in relationship to height than males, a greater inclination of ribs, a comparable diaphragm dome position relative to the spine, and a shorter diaphragm length” [15]. It has also been suggested that chest radiographs findings may differ among females and males for lung diseases such as TB [16] and cystic fibrosis [17]. Finding good features that can capture these differentiating visual characteristics between sexes is crucial. The process of feature

handcrafting may require significant skills and time for trial and error. A data-driven approach which can automatically learn good features is appealing. The overall diagram of the method is shown in Figure 2. It consists of four major steps: pre-processing, CNN feature extractor, feature selection, and classifier. In the following, we will describe each step in detail.




Figure 2. Schematic diagram showing the processes in the method




Figure 3. Black border removal

Figure 4. Image cropping

Figure 5. Contrast enhancement

There is a large variance in visual appearance between the five individual datasets, as demonstrated in Figure 1. We first apply several pre-processing steps to the images before inputting them to a pre-trained CNN to extract features. The steps include: intensity conversion/normalization, black border removal, image cropping, contrast enhancement, and image resizing. Each individual image dataset may require different pre-processing steps due to the distinctive characteristics of the images in that dataset. For images that are in DICOM format (such as the images in the Kenya and Montgomery datasets), we convert their intensity from Hounsfield units to the range of [0, 1] and save the images to PNG format (8 bits) based on the windowing information and the photometric interpretation field in the DICOM header if they are available [18]. The raw image data (12 bits) in the JSRT dataset is converted to PNG format (8 bits). As shown in Figure 3, some images in the Montgomery and Kenya datasets contain black borders (intensity = 0) therefore those borders are removed [2]. For the images in the Kenya dataset, a brighter band (which is the protection pad being put on the patient’s stomach) frequently appears in the bottom of the images (as shown in Figure 4). As a result, the aspect ratio (height/width) of images in the Kenya dataset is usually larger than that of images in other datasets due to the extra inclusion of the stomach. At present, we crop those images coarsely based on a simple rule. Specifically, the 1/6 height on the bottom of the image is removed if the height/width ratio of the image is larger than a threshold (for example, 1.2). A more accurate segmentation method may be developed as future work. The contrast of the images in the JSRT dataset and Kenya dataset is enhanced using the method of histogram equalization. One example is shown in Figure 5. The low contrast images in

the Indiana dataset are identified and enhanced using the method presented in [18]. Finally, all the images are resized to a specific resolution required by the employed CNN architectures.

To extract features, we compare four CNN architectures. They are AlexNet, VggNet, GoogLeNet, and ResNet. AlexNet [19] was the first CNN method entered in the annual ImageNet Large Scale Visual Recognition Challenge (ILSVRC) competitions, and it outperformed the conventional methods significantly. Since then (ILSVRC 2012), the CNNs have become the leading methods for the following ILSVRC competitions and for applications in computer vision and speech recognition. VggNet [20] and GoogLeNet [21] were introduced in ILSVRC 2014. They both have substantially more layers (around 20) than previous architectures. VggNet is characterized by using only 3x3 sized filters in all convolutional layers. GoogLeNet is designed to efficiently utilize the computing resources inside the network. Different from previous architectures that layers are always stacked up sequentially, it consists of modules named “Inception” which contains parallel layers. ResNet [22] is the winner of ILSVRC 2015. The hallmark of this architecture is its ultra-deep characteristic (for example, with 152 layers). A good introduction to the characteristics of these CNN architectures is given in [23]. Implementations of above CNNs are included in several publicly available software, such as Caffe [24], TensorFlow [25], Torch [26], DeepLearning4j [27], and MatConvNet [28]. All these software implementations provide models of above CNNs that have been trained using the ImageNet dataset. For our application, we use the software MatConvNet and Matlab. For all four CNN architectures, we use the output of the layer right before the classification layer as the feature descriptor for the input image. Table 2 lists the input image size and output feature length for each CNN architecture.

Table 2. Image size and feature length for each CNN architecture

CNNs	Image Size	Feature Length
AlexNet	227 x 227 x 3	4096
VggNet	224 x 224 x 3	4096
GoogLeNet	224 x 224 x 3	1024
ResNet	224 x 224 x 3	2048

Comparing the feature length in Table 2 and the number of images in the dataset in Table 1, it shows the length of features obtained from AlexNet or VggNet (= 4096) is quite larger than the number of images (=2066). Therefore, we apply the technique of feature selection to reduce the feature length in order to circumvent the curse of dimensionality by removing redundant information. There are many methods for feature selection. We use the correlation-based feature selection (CFS) [29] method that is implemented in Weka [30]. It evaluates the value of a subset of features by considering the individual predictive ability of each feature as well as the degree of redundancy between them. The Best First method [31] (which searches the space of feature subsets by greedy hill-climbing augmented with a backtracking facility) is used as the strategy for searching the feature subset space. After getting the features, we apply and compare two conventional classifiers: support vector machines (SVM) and random forest (RF). Both classifiers are used frequently and are reported to show good performance in the literature.

4. EXPERIMENTAL RESULTS AND DISCUSSION

All the CNN architectures (AlexNet, VggNet, GoogLeNet, ResNet) were pre-trained using the ILSVRC ImageNet dataset. For VggNet, two versions are tested: VggNet-16 and VggNet-19 (having 16 and 19 layers respectively). For ResNet, we also tested two versions: ResNet50 and ResNet152 (having 50 layers and 152 layers respectively). We use the open source machine learning software Weka for the step of feature selection and classification. Specifically, the classifier is set to be the “AttributeSelectedClassifier” in which attributes are selected and the classifier is trained based on the training set only [32]. For the attribute selection method used in the “AttributeSelectedClassifier”, we select the feature evaluator as “CfsSubsetEval” and the search method as “BestFirst”. For the classifier used in the “AttributeSelectedClassifier”, we select either the sequential minimal optimization (SMO) algorithm (for SVM) or random forest. Table 3 lists the length of feature vector after feature selection for each CNN architecture. For all CNN architectures, the feature length was significantly reduced, from an original length of thousand(s) to less than two hundred. We evaluated the classification performance using five-fold cross-validation with the following performance measures: accuracy, TP (true positive) rate, FP (false positive) rate, precision, recall, F-score, and ROC area. Table 4 and Table 5 list the values of all these measures for each CNN architecture for SMO and RF classifiers, respectively. For all CNN features, the performance of SMO is better than that of RF. For both SMO and RF classifiers, VggNet-16 feature and AlexNet feature are the top two performers,

outperforming the features extracted from the other deeper architectures. Specifically, for the SMO classifier, VggNet-16 feature gets the best results (with accuracy being 86.6%); and for the RF classifier, AlexNet feature gets the best results (with accuracy being 84.6%). We examine the mis-classified cases in order to get an idea on when/why the method fails and how to improve the method. Figure 6 and Figure 7 give several mis-classified examples which are classified using the VggNet-16 feature with SMO for females and males, respectively. The dataset contains images of pediatric patients (such as images in Figure 6(a) and Figure 7(a)) (age information is not available to all the images in the dataset). It is better to separate adult images from pediatric images as they have different characteristics regarding to gender discrepancy as shown in chest x-rays. We expect the performance will be increased if considering only adult patients as the gender discrepancy in pediatric images is much less differentiable than adult images. There are images like the one shown in Figure 7(b) which the labeled sex information is male but is predicted as female probably because of the noticeable presence of the mammary glands contours in the images (this could be because the ground truth labels are not correct and we need to check them with radiologists). The examples in Figure 6 and 7 also indicate that the visual gender differences exhibited in chest x-rays can be very subtle. In the future, we will work on the following aspects for improving the performance: removing pediatric images from the dataset (may develop a method for classifying pediatric images and adult images); using the CNNs that are pre-trained by a large medical image dataset instead of ImageNet; exploring a combination of CNN features and traditional features; getting more data (we are collecting more chest x-rays in Kenya which can add to our dataset); testing the same method on lateral chest x-rays. We apply the method to the Indiana University dataset. Since there are no ground truth labels for this dataset, we cannot quantitatively evaluate the performance. We visually examine the results and consider the results are encouraging (our preliminary subjective evaluation is mainly based on the assumption that female patients have mammary glands contours presented in the images).

Table 3. Feature length after feature selection for each CNN

	AlexNet	VggNet-16	VggNet-19	GoogLeNet	ResNet-50	ResNet-152
Feature Length	140	168	152	93	123	119

Table 4. 5-fold cross validation performance for each CNN (SMO)

CNNs	Accuracy	TP Rate	FP Rate	Precision	Recall	F-Measure	ROC Area
AlexNet	85.8%	0.858	0.142	0.858	0.858	0.858	0.925
VggNet-16	86.6%	0.866	0.135	0.866	0.866	0.866	0.932
VggNet-19	84.3%	0.843	0.159	0.843	0.843	0.843	0.918
GoogLeNet	82.5%	0.825	0.175	0.826	0.825	0.825	0.892
ResNet-50	79.0%	0.790	0.211	0.790	0.790	0.790	0.862
ResNet-152	79.0%	0.790	0.211	0.790	0.790	0.790	0.871

Table 5. 5-fold cross validation performance for each CNN (RF)

CNNs	Accuracy	TP Rate	FP Rate	Precision	Recall	F-Measure	ROC Area
AlexNet	84.6%	0.846	0.157	0.846	0.846	0.845	0.923
VggNet-16	84.4%	0.844	0.157	0.844	0.844	0.844	0.911
VggNet-19	82.9%	0.829	0.175	0.829	0.829	0.828	0.907
GoogLeNet	80.0%	0.799	0.204	0.799	0.799	0.799	0.888
ResNet-50	76.3%	0.763	0.243	0.763	0.763	0.762	0.844
ResNet-152	77.4%	0.774	0.234	0.775	0.774	0.773	0.856

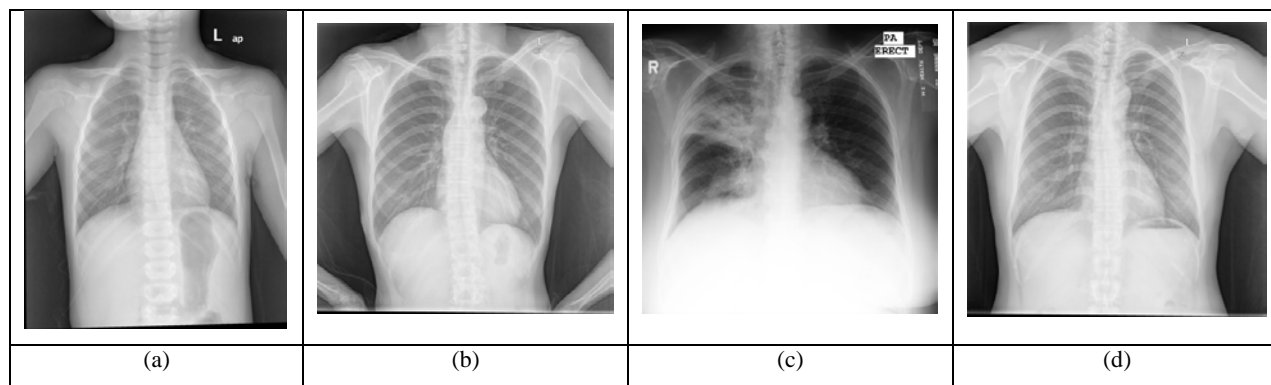


Figure 6. Misclassified images (Female predicted as Male)




Figure 7. Misclassified images (Male predicted as Female)

5. CONCLUSIONS

For chest x-ray images, sometimes, the gender information of patient is not available. In this paper, we present a method to automatically identify the gender of a patient based on the patient's chest x-ray only, as body differences between adult females and males that may be exhibited in the chest x-rays. The proposed method employs the technique of deep learning and transfer learning to deal with the challenge of feature handcrafting and the limit of data. Specifically, the method consists of four main steps: pre-processing, CNN feature extraction, feature selection and classification. The method is tested on a combined dataset obtained from several sources, which results in a final dataset with large variance. Different pre-processing steps are applied for each individual image dataset due to their source differences. For feature extraction, four CNN architectures are tested and compared. They are AlexNet, VggNet, GoogLeNet, and ResNet. A feature selection technique is applied, since the feature length is larger than the number of images. Two popular classifiers: SVM and Random Forest, are used and compared. The classification performance is evaluated by cross-validation and measured with several matrices. The best performer is the VggNet-16 feature extractor with the SVM classifier, with accuracy of 86.6% and ROC Area being 0.932 for 5-fold cross validation. We also discuss several misclassified cases and describe future work for performance improvement.

ACKNOWLEDGEMENT

This research was supported by the Intramural Research Program of the National Institutes of Health (NIH), National Library of Medicine (NLM), and Lister Hill National Center for Biomedical Communications (LHNCBC).

REFERENCES

- [1] Stanescu, L., Burdescu, D.D., Brezovan, M., Mihai, C.J., "Creating new medical ontologies for image annotation: a case study," Springer Publishing Company, Incorporated (2012).
- [2] Candemir, S., Jaeger, S., Palaniappan, K., Musco, J.P., Singh, R.K., Xue, Z., Karargyris, A., Antani, S., Thoma, G. R., McDonald, C. J., "Lung segmentation in chest radiographs using anatomical atlases with nonrigid registration," *IEEE Transactions on Medical Imaging* 33 (2), 577-590 (2014).
- [3] Candemir, S., Jaeger, S., Lin, W., Xue, Z., Antani, S., Thoma, G.R., "Automatic heart localization and radiographic index computation in chest x-rays," *Proceedings of SPIE* 9785, 978517 (2016).
- [4] Candemir, S., Jaeger, S., Antani, S., Bagci, U., Folio, L., Xu, Z., Thoma, G.R., "Atlas-based rib-bone detection in chest X-rays," *Computerized Medical Imaging and Graphics* 51, 32-39 (2016)
- [5] Xue, Z., You, D., Candemir, S., Jaeger, S., Antani, S., Long, L. R., Thoma, G. R., "Chest X-ray image view classification," *Proceedings of the 28th IEEE International Symposium on Computer-Based Medical Systems*, 66-71 (2015).
- [6] Jaeger, S., Karargyris, A., Candemir, S., Folio, L., Siegelman, J., Callaghan, F., Xue, Z., Palaniappan, K., Singh, R., Antani, S., "Automatic tuberculosis screening using chest radiographs," *IEEE Transactions on Medical Imaging* 33 (2), 233-245 (2014).
- [7] Demner-Fushman, D., Kohli, M. D., Rosenman, M. B., Shooshan, S. E., Rodriguez, L., Antani, S., Thoma, G. R., McDonald, C. J., "Preparing a collection of radiology examinations for distribution and retrieval," *Journal of the American Medical Informatics Association* 23(2), 304-310 (2016).
- [8] Bar, Y., Diamant, I., Wolf, L., Greenspan, H., "Deep learning with non-medical training used for chest pathology identification," *Proceedings of SPIE Medical Imaging*, 94140V (2015).
- [9] Ginneken, B., Setio, A. A., Jacobs, C., Ciompi, F., "Off-the-shelf convolutional neural network features for pulmonary nodule detection in computed tomography scans," *Proceedings of the 12th IEEE International Symposium on Biomedical Imaging (ISBI)*, 286-289 (2015).
- [10] Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., "ImageNet: a large-scale hierarchical image database," *IEEE Computer Vision and Pattern Recognition (CVPR)*, 248-255 (2009).
- [11] Jaeger, S., Candemir, S., Antani, S., Wang, Y.-X. J., Lu, P.-X., Thoma, G., "Two public chest X-ray datasets for computer-aided screening of pulmonary diseases," *Quantitative Imaging in Medicine and Surgery*, 4(6), 475-477 (2014).
- [12] <https://ceb.nlm.nih.gov/repositories/tuberculosis-chest-x-ray-image-data-sets/>, last accessed 2017/05/01.
- [13] Shiraishi, J., Katsuragawa, S., Ikezoe, J., Matsumoto, T., Kobayashi, T., Komatsu, K., Matsui, M., Fujita, H., Kodera, Y., Doi, K., "Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists' detection of pulmonary nodules," *American Journal of Roentgenology* 174 (1), 71-74 (2000).
- [14] <https://openi.nlm.nih.gov/faq.php?it=xg>, last accessed 2017/05/01.
- [15] Bellemare, F., Jeanneret, A., Couture, J., "Sex differences in thoracic dimensions and configuration," *American Journal of Respiratory and Critical Care Medicine* 168(3), 305-312 (2003).
- [16] Thorson, A., Long, N. H., Larsson, L. O., "Chest X-ray findings in relation to gender and symptoms: A study of patients with smear positive tuberculosis in Vietnam," *Scandinavian Journal of Infectious Diseases* 39, 33-37 (2007).
- [17] Bellemare, F., Jeanneret, A., "Sex differences in thoracic adaptation to pulmonary hyperinflation in cystic fibrosis," *American Journal of Respiratory and Critical Care Medicine* 29, 98-107 (2007).
- [18] Xue, Z., Candemir, S., Antani, S., Long, L. R., Jaeger, S., Thoma, G.R., "Foreign object detection in chest X-rays," *Proceedings of the 2015 IEEE International Conference on Bioinformatics and Biomedicine*, 956-961 (2015).
- [19] Krizhevsky, A., Sutskever, I., Hinton, G.E., "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, 1097-1105 (2012).
- [20] Simonyan, K., Zisserman, A., "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556* (2014).
- [21] He, K., Zhang, X., Ren, S., Sun, J., "Deep residual learning for image recognition," *Proceedings of IEEE Computer Vision and Pattern Recognition (CVPR)*, 770-778 (2016).
- [22] Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., "Going deeper with convolutions," *Proceedings of IEEE Computer Vision and Pattern Recognition (CVPR)*, (2015).
- [23] <https://adeshpande3.github.io/adeshpande3.github.io/The-9-Deep-Learning-Papers-You-Need-To-Know-About.html>, last accessed 2017/07/06.
- [24] <http://caffe.berkeleyvision.org/>, last accessed 2017/07/06.
- [25] <https://www.tensorflow.org/>, last accessed 2017/07/06.
- [26] <http://torch.ch/>, last accessed 2017/07/06.
- [27] <https://deeplearning4j.org/>, last accessed 2017/07/06.
- [28] <http://www.vlfeat.org/matconvnet/>, last accessed 2017/07/06.
- [29] Hall, M. A., *Correlation-based Feature Selection for Machine Learning*, PhD dissertation, Department of Computer Science, University of Waikato, (1999)

- [30] <http://www.cs.waikato.ac.nz/ml/weka/>, last accessed 2017/07/06.
- [31] Rich, E., Knight, K., Artificial Intelligence, McGraw-Hill, (1991).
- [32] Domingos, P., "A few useful things to know about machine learning," Commun. ACM 55(10), 78-87 (2012).