Tuberculosis Chest X-Ray Image Retrieval System Using Deep Learning Based Biomarker Predictions

Bradley C. Lowekamp, Andrei Gabrielian, Darrell E. Hurt, Alex Rosenthal, and Ziv Yaniv

Office of Cyber Infrastructure and Computational Biology, National Institute of Allergy and Infectious Diseases, Bethesda, MD 20892, USA

ABSTRACT

The world health organization's global tuberculosis (TB) report for 2022 identifies TB, with an estimated 1.6 million, as a leading cause of death. The number of new cases has risen since 2020, particularly the number of new drug-resistant cases, estimated at 450,000 in 2021. This is concerning, as treatment of patients with drug resistant TB is complex and may not always be successful. The NIAID TB Portals program is an international consortium with a primary focus on patient centric data collection and analysis for drug resistant TB. The data includes images, their associated radiological findings, clinical records, and socioeconomic information. This work describes a TB Portals' Chest X-ray based image retrieval system which enables precision medicine. An input image is used to retrieve similar images and the associated patient specific information, thus facilitating inspection of outcomes and treatment regimens from comparable patients. Image similarity is defined using clinically relevant biomarkers: gender, age, body mass index (BMI), and the percentage of lung affected per sextant. The biomarkers are predicted using variations of the DenseNet169 convolutional neural network. A multi-task approach is used to predict gender, age and BMI incorporating transfer learning from an initial training on the NIH Clinical Center CXR dataset to the TB portals dataset. The resulting gender AUC, age and BMI mean absolute errors were 0.9854, 4.03years and $1.67\frac{kg}{m^2}$. For the percentage of sextant affected by lesions the mean absolute errors ranged between 7% to 12% with higher error values in the middle and upper sextants which exhibit more variability than the lower sextants. The retrieval system is currently available from https://rap.tbportals.niaid.nih.gov/find_similar_cxr.

Keywords: deep learning, tuberculosis, chest x-ray, content-based image retrieval, explainable biomarkers, precision medicine

1. INTRODUCTION

Treatment of drug resistant tuberculosis (TB) is complex, requiring the use of patient specific drug regimens and consistent progression monitoring. Use of some drug combinations can potentially lead to prolonged periods of treatment due to less effective regimens. This can potentially reduce patient adherence which, in turn, leads to less favorable outcomes. This paper presents an image retrieval system which facilitates precision medicine. Using a chest X-ray as input, cases for similar patients are returned and one can examine the individual cases to gain a better understanding of disease and treatment progression over time.

Medical content-based image retrieval has been studied extensively over the past several decades.[1] By using features derived from an image of a new patient, similar images, in terms of feature values, are returned with their accompanying meta-data. The meta-data, e.g. existing patient records, can then be used to inform medical decision making for the new patient. Early on, image features were manually defined, using domain knowledge to model low level image characteristics. More recently, features learned by deep learning networks have been shown as more effective image descriptors.[2] In both approaches, feature selection and the choice of a specific distance metric in feature space are designed so that images that belong to the same category are closer in the feature space than those that are in different categories. As the same image can belong to many categories,

Further author information:
B.C.L.: E-mail: blowekamp@mail.nih.gov
Z.Y.: E-mail: zivyaniv@nih.gov

Medical Imaging 2024: Imaging Informatics for Healthcare, Research, and Applications, edited by Hiroyuki Yoshida, Shandong Wu, Proc. of SPIE Vol. 12931, 129310X
© 2024 SPIE · 1605-7422 · doi: 10.1117/12.3006848

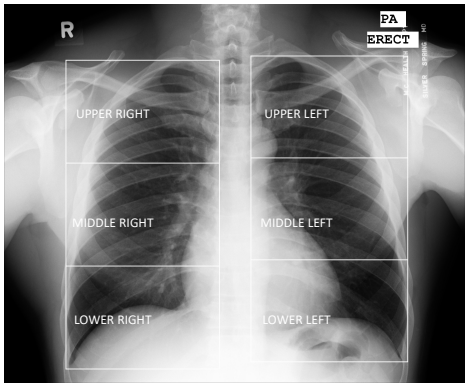

Figure 1. Definition of lung sextants used to estimate the percentage of sextant affected by lesions.

learning meaningful features and an appropriate distance metric is context specific. This is often referred to as the "semantic gap".

In the context of the current work, TB frontal Chest X-Rays (CXRs), we define image categories using the following clinically relevant biomarkers: gender, age, Body Mass Index (BMI) and percentage of lung sextant affected by lesions. Gender and age are generic biomarkers which have been previously shown to significantly contribute to retrieval performance of frontal CXRs.[3] In the context of TB in general and drug resistant TB in particular, gender and age are correlated with higher disease rates in males and higher treatment failure rates for older patients. The former is associated with both biological differences, more effective immune system in females, and socio-cultural behaviors, e.g. higher prevalence of smoking in males.[4] The latter is due to reduced intestinal function and motility, associated with older age, resulting in lower drug absorption.[5] BMI is an overall health indicator, with lower BMI known to be a risk factor for both acquiring TB and for unfavorable TB outcomes.[6,7] Finally, percentage of lung sextant affected by lesions indicates the spatial location and severity of the disease. Sextants are defined by dividing each lung into three equal regions as shown in Figure 1.

To circumvent the problem of a semantic gap between derived image features, distance metric and image categories, in this work the category defining biomarkers are directly predicted and the Manhattan distance between them is used as the similarity metric. In this way, the order of retrieved images is readily explainable, and errors can be directly traced to specific biomarker prediction discrepancies. Consequentially, there is no gap between the user's intent and the system behavior.

2. MATERIALS AND METHODS

2.1 Data

Two publicly available frontal CXR datasets are utilized in this work, one from the NIH Clinical Center (CC)[8] and the other from the TB Portals (TBP) program,[9] August 2023 data release. The CC dataset was used to train models for predicting gender and age. The TBP dataset was used to train models for predicting gender, age, BMI and percentage of lung sextant affected by lesions.

TBP – This dataset contains about 11,000 CXRs from 8800 TB patients. The dataset is heterogeneous in terms of image intensities and sizes. It includes both high dynamic range images with intensities greater than 255, and low dynamic range images with intensities in [0,255]. Image sizes vary significantly, ranging between 128×123 to 6512×5687. For each patient, the CXR acquired closest to the patient registration date is associated with the following metadata, gender, patient age and BMI. In addition, the following lesion types and the percentage of sextant affected by them are listed: cavities (<3cm, 3-5cm, >5cm), ground glass opacity, and nodules (<3mm, 3-8mm, 8-30mm, >30mm). These quantities were manually obtained by a single clinician reading of each image. Note that it is not the same clinician reading all images. The total percentage of lung sextant affected by lesions is computed by summing the percentages for all lesions types per sextant and capping the result at 1.0. The need for capping the total value is due to the manual reading. The clinician estimated the percentage of the

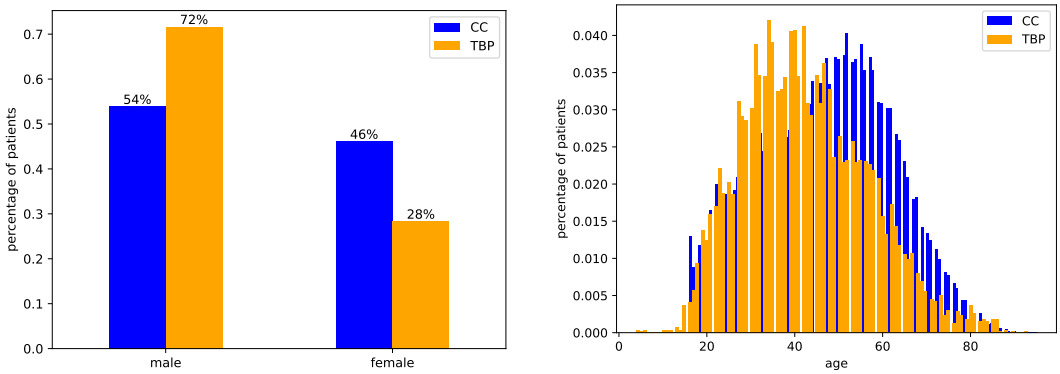

Figure 2. Data distributions for the TBP and CC datasets illustrating the domain shift between the two datasets. Note that the gender imbalance in the TBP dataset reflects disease characteristics.[4]

sextant affected independently per lesion type and in some cases this results in a value that is greater than the maximal possible one.

CC – This dataset contains about 112,000 images from about 31,000 patients. The majority of images (about 84,000) have no clinical findings and are classified as normal. The remaining images contain one or more of eight clinical findings (atelectasis, cardiomegaly, effusion, infiltration, mass, nodule, pneumonia, and pneumothorax). The dataset is homogeneous in terms of image intensities and sizes. All images have a low dynamic range and were resized to 1024×1024. For all images the gender and age of the patient is available.

2.2 Predicting Gender, Age and BMI

To predict gender, age in years and BMI in $\frac{kg}{m^2}$, the high dynamic range images from the TBP dataset, about 1100 images, and the images from the CC dataset were used. As the goal is to use these biomarkers in conjunction with the total percentage of lung sextant affected by lesions in the computation of the Manhattan distance, their values are mapped to the interval [0,1]. The values for age and BMI are linearly scaled from their original bounded possible ranges of [0,120][10] and [0,80].[11] In addition, gender is treated as a continuous variable and not a categorical one. Male is mapped to 0.0 and female to 1.0, with the predicted continuous gender value in [0,1].

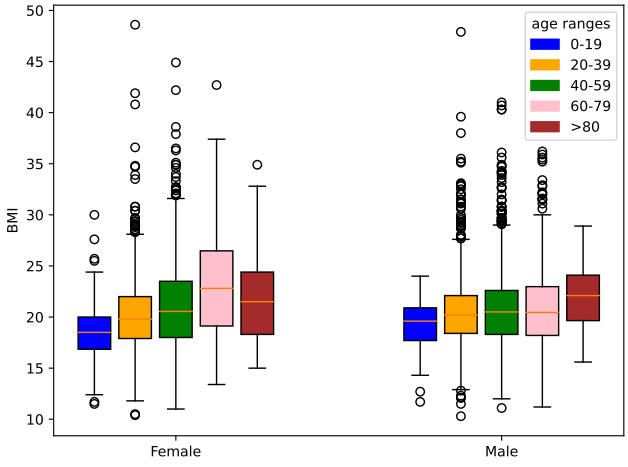

Figure 3. BMI as a function of gender and age group in the TBP dataset. Median BMI per age group female(male): 18.5(19.6), 19.8(20.2), 20.55(20.5), 22.8(20.45), 21.5(22.1). Age ranges are in years.

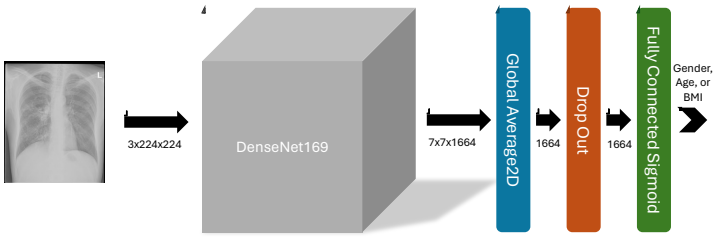

Figure 4. Network architecture used by the single task formulation. Each prediction task (gender, age and BMI) has its own network weights.

Figure 2 shows the gender and age distributions for the two datasets, highlighting the domain shift between them, which informed our modeling approach. Using the CC dataset for training to predict gender and age followed by transfer learning to the TBP dataset.

Initially a separate Convolutional Neural Network (CNN) model was trained for each of the prediction tasks, resulting in three models independently trained and tested on the TBP dataset. We then observed that in the TBP dataset gender, age and BMI do not appear to be independent, as illustrated in Figure 3. Consequentially, the problem was re-formulated using multi-task optimization, training a single CNN model to predict all three biomarkers. Finally, to utilize the much larger CC dataset, the problem was formulated as a two stage multi-task optimization. Initially training with the CC dataset to predict gender and age, and then performing transfer learning and adding the BMI prediction task using the TBP dataset. Detailed descriptions of each of the approaches follow.

2.2.1 Separate Task Approach

To predict gender, age and BMI the original DenseNet169[12] CNN architecture was modified. The classification layers used for the imagenet task were replaced by a 2D global average pooling layer, followed by a dropout layer (20% dropout rate), and a fully connected layer with a single output and sigmoid activation. Note that the dropout layer was only utilized during training to reduce the likelihood of over-fitting to the training data. The resulting network architecture is shown in Figure 4. The total number of network parameters was about 12.5 million.

All three models were trained with Tensorflow 2.8.3 and Keras 2.8.0 [13] using the following two-step approach. First, model weights were initialized using the imagenet DenseNet169 weights. These weights were fixed and only the weights of the new regression layers were optimized (about 1500 parameters). Second, the weights for all but the first 100 layers of the network were optimized (about 11.5 million parameters).

Both steps used the Adam optimizer with learning rates of 0.001 and 0.0001 for the first and second steps respectively. In both steps, the Mean Absolute Error (MAE) loss function was optimized until the validation loss did not decrease for five consecutive epochs. The weights for the model with the minimal validation loss were

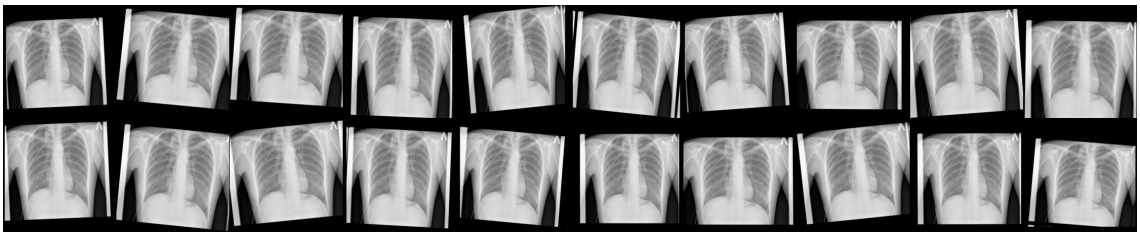

Figure 5. Results of applying the data augmentation strategy used in this work to a single CXR. Increasing the data variability during model training.

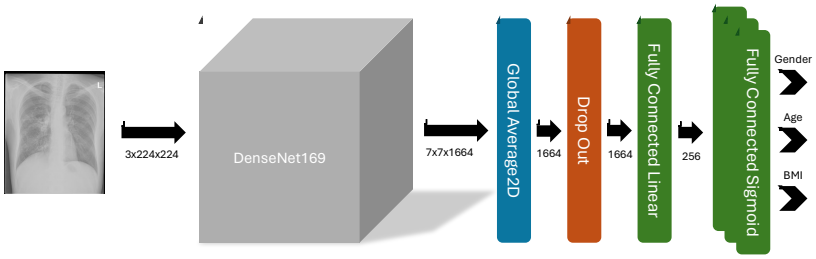

Figure 6. Network architecture used in the multi-task (gender, age, BMI) formulation.

used as the final model weights. The batch size was set to 64 and the dropout rates was 20%. Keras dataset input pipelines were used with SimpleITK[14] to dynamically read DICOM images, perform data augmentation, and then resize the images to (224×224). Data augmentation included random translation ($\pm 7\%$ of image size), rotation ($\pm 10^o$) and zooming in (0% to 10% of image size). Each training image was dynamically permuted 4-times per epoch for the TBP data only. Figure 5 illustrates the effects of these data augmentation settings on a single CXR.

2.2.2 Multi-Task Approach

Similar to the single task approach, in the multi-task approach a variation of the DenseNet169 architecture was utilized. The DenseNet169 classification layers were replaced by a 2D global average pooling layer, followed by a dropout layer (10% dropout rate), then a fully connected layer with 256 outputs and ReLU activation, and lastly an additional fully connected layer with sigmoid activation and three outputs. The resulting network architecture is shown in Figure 6. The total number of network parameters was about 13 million.

Training followed the same two step approach as described above with the same parameter settings and data augmentation settings. The main difference is that the loss function was defined as a weighted combination of the MAE losses for gender ($w_2 = 0.7$), age ($w_1 = 0.15$) and BMI ($w_3 = 0.15$). For this network, the number of optimized parameters in the first training step was about 400,000 and in the second step about 12 million.

2.2.3 Transfer Multi-Task Approach

As the size of the subset of the TBP dataset used in this work is not large, we leveraged the much larger CC dataset to improve the final model's performance. Additionally, it has been shown that transfer learning from a dataset with a different distribution can reduce the probability of the final model learning shortcuts (usage of data attributes that improve performance during training but are unrelated to the task at hand).[15]

The transfer multi-task approach used the same network architecture as the multi-task approach. The network was first trained on the CC data set for gender and age with the BMI output removed. The same two step optimization approach described above was performed. The only differences were that the dropout rate was 20% and the loss function was defined as a weighted combination of the MAE losses for gender ($w_2 = 0.3$) and age ($w_1 = 0.7$).

Next, a fully connected output layer with sigmoid activation for BMI was added to the network and initialized with random weights. The remaining layers were initialized using the weights from the CC trained network. Training was then performed with the TB Portals dataset. The two step optimization approach described above was performed with a dropout rate of 10% and the loss function was defined as a weighted combination of the MAE losses for gender ($w_2 = 0.1$), age ($w_1 = 0.2$) and BMI ($w_3 = 0.7$).

2.3 Predicting Percentage of Lung Sextant Affected by Lesions

To predict the percentage of lung sextant affected, the classification layers of the DenseNet169 architecture were replaced by a dropout layer, a fully connected layer with 256 outputs and linear activation, and a final, fully connected layer with six outputs and sigmoid activation. The resulting network architecture is shown in

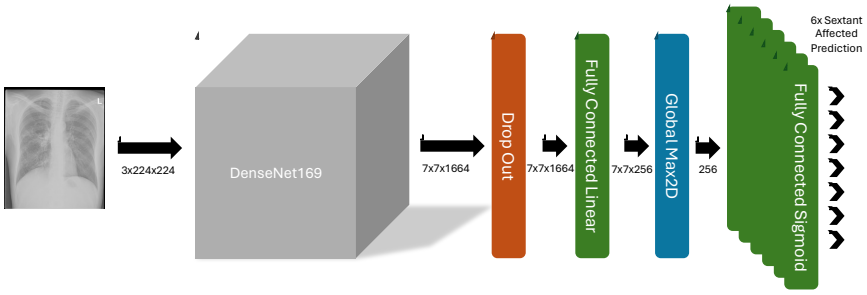

Figure 7. Network architecture for predicting the percentage of lung sextant affected by lesions.

Figure 7. The total number of network parameters was about 13 million. For this network, the number of optimized parameters in the first training step was about 400,000 and in the second step about 12 million.

Training followed the two step approach used to train the single task networks described above. Model weights were initialized using the imagenet DenseNet169 weights and only the new regression layer weights were optimized, then the weights for all but the first 100 layers were optimized. All training parameter settings were the same as listed in section 2.2.1.

3. EXPERIMENTAL EVALUATION

All models were evaluated using a five fold cross validation approach. To ensure that there is no information leakage between training and testing, the data was first split into 80% for training-validation and 20% for testing. The split was done using patient IDs, so that all images belonging to the same patient are either in the training-validation set or test set. The training-validation data was then further split into 90% training and 10% validation. All testing was done on the TBP images.

Approach	Gender AUC	Age Mean Absolute Error(stddev)	BMI Mean Absolute Error(stddev)
Separate Task	0.9592	6.69(5.26)	2.07(1.89)
	0.9722	6.52(4.57)	2.02(2.17)
	0.9625	6.76(5.15)	2.01(1.53)
	0.9900	7.15(5.81)	2.29(2.29)
	0.9847	6.11(5.11)	2.11(2.01)
Multi-Task	0.9671	8.17(6.46)	2.39(1.95)
	0.9714	6.82(5.35)	2.05(2.04)
	0.9548	7.59(5.53)	2.09(1.63)
	0.9739	7.09(5.44)	2.65(2.58)
	0.9772	6.91(5.96)	2.27(1.96)
Transfer Multi-Task	0.9692	4.38(4.06)	1.51(1.58)
	0.9849	3.84(3.38)	1.41(1.68)
	0.9836	3.60(3.00)	1.60(1.43)
	0.9934	4.14(3.63)	2.08(2.41)
	0.9963	4.18(3.95)	1.74(1.76)

Table 1. Results for five fold cross validation using the three different modeling approaches. Mean absolute errors for age are in years and for BMI in $\frac{kg}{m^2}$.

Lower Left	Lower Right	Middle Left	Middle Right	Upper Left	Upper Right
8.32(18.37)	7.28(15.88)	12.36(21.22)	10.16(17.07)	14.36(17.44)	11.49(14.92)
7.52(18.60)	7.07(17.29)	10.88(20.53)	9.58(15.42)	9.13(17.44)	11.59(14.10)
8.66(18.87)	6.38(14.83)	11.84(20.47)	10.85(18.69)	11.01(15.58)	11.92(15.99)
8.26(17.73)	6.00(13.71)	8.10(16.04)	9.14(14.87)	10.90(16.31)	13.22(15.32)
8.30(18.03)	8.14(16.23)	9.56(17.42)	12.39(19.50)	12.17(17.37)	12.63(15.84)

Table 2. MAE results for five fold cross validation using the modified DenseNet169 architecture to predict the percentage of lung sextant affected by lesions. Errors are in percent, with the maximal possible error value being 100.

4. RESULTS

The results for predicting the gender, age and BMI for each of the three approaches were as follows. Using the separate task approach, the mean values for the five fold testing for gender AUC, age MAE and BMI MAE were 0.9737, 6.65years and 2.10$\frac{kg}{m^2}$. Using the multi-task approach, the mean values for gender AUC, age MAE and BMI MAE were 0.9689, 7.32years and 2.29$\frac{kg}{m^2}$. Finally, using the transfer multi-task approach, the mean values

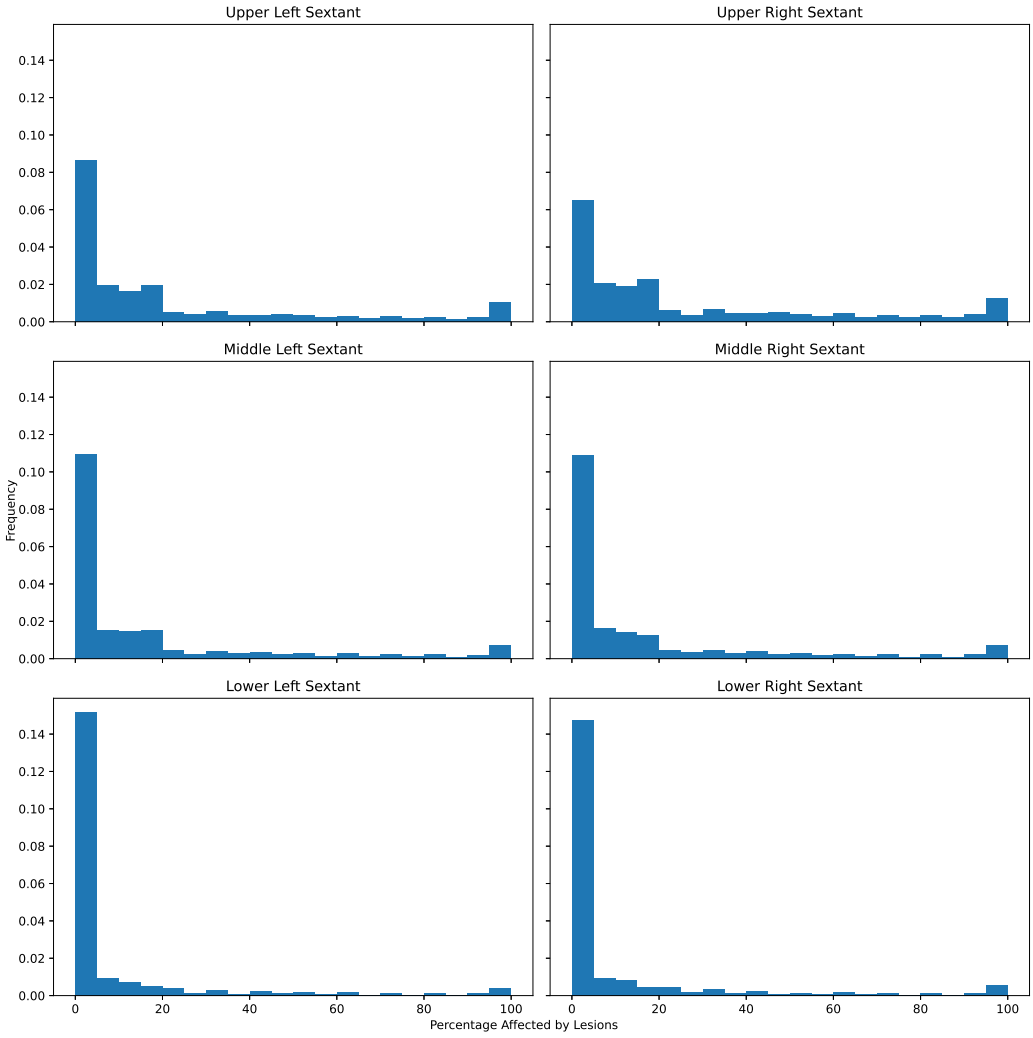

Figure 8. Reference data for the percentage of lung sextant affected by lesions. Notice the higher variability in the middle and apical sextants where lesions are more commonly found.[16]

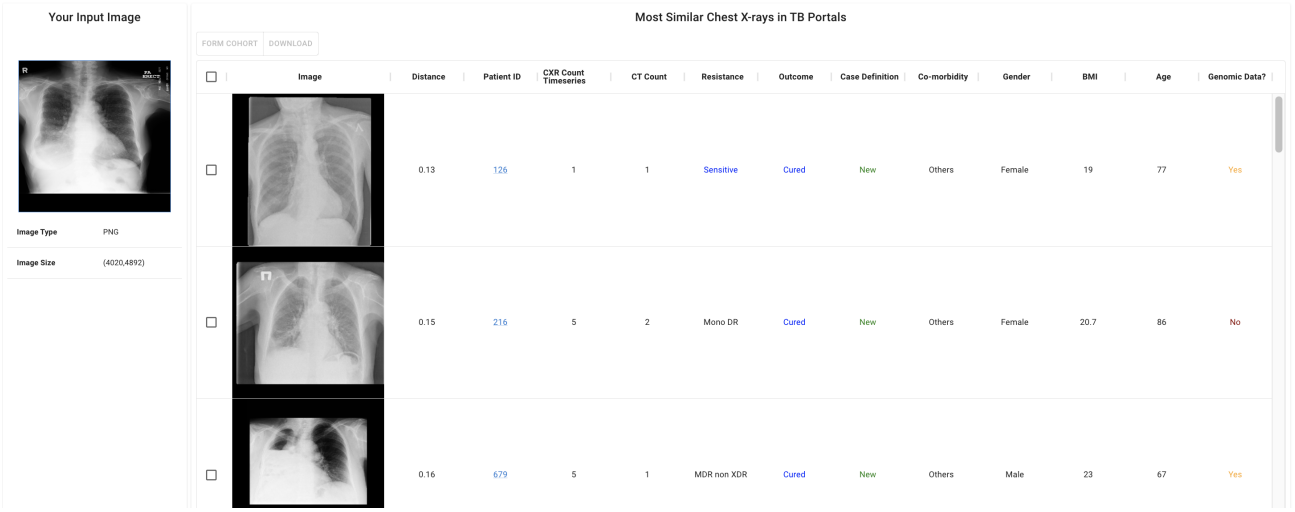

Figure 9. Image retrieval system interface, query image on the left and first three results with corresponding case information. Complete case details are available via the patient ID column hyperlinks.

for gender AUC, age MAE and BMI MAE were 0.9854, 4.03years and $1.67\frac{kg}{m^2}$. Note that for gender AUC, the higher the value the better. For age and BMI MAE values, the lower the value the better, as this corresponds to a lower error. Table 1 provides detailed results per fold for each of the three approaches.

With respect to the percentage of lung sextant affected by lesions, the mean absolute errors for the five fold testing per lung sextant were: lower left 8.21%, lower right 6.98%, middle left 10.55%, middle right 10.43%, upper left 11.51%, and upper right 12.17%. Table 2 provides detailed results per fold and sextant.

5. DISCUSSION AND CONCLUSIONS

To enable precision medicine in the context of drug resistant TB, we developed a content based image retrieval system using a frontal CXR as input to retrieve patient cases from the TBP database. To eliminate the semantic gap between the input and the class it belongs to, an explicit set of biomarkers was used to define the class and distances between images. The specific biomarkers utilized in the work are a combination of generic markers (gender, age, BMI) and TB specific ones (percentage of lung sextant affected by lesions).

While the original biomarker ranges are non-commensurate, gender is a categorical marker, age is in [0,120], BMI is in [0,80], and percentage of sextant affected by lesions is in [0,1], the prediction task was formulated such that all predictions were continuous values in the range [0,1] as described above.This allowed the use of the Manhattan distance as the distance between images.

Initially, gender, age and BMI were estimated using the straightforward approach. That is, a dedicated model was utilized to predict each of these biomarkers. Upon further exploration of the data it was observed that these three biomarkers are not necessarily independent from each other. Consequentially, a multi-task approach was evaluated. Unfortunately, the multi-task approach trained on TBP data yielded worse results than the single task. Taking advantage of a much larger frontal CXR dataset which included gender and age information, a two step transfer multi-task approach was used. That is, an initial gender and age model was trained on the larger dataset using a multi-task approach and those weights were then used to initialize the multi-task approach on the TBP data which predicted gender, age and BMI. This later approach was shown to yield the best results, with an average AUC of 0.9854 for gender, mean absolute error of 4.03years for age and a mean absolute error of $1.67\frac{kg}{m^2}$. In addition to improved performance, the multi-task model has a significantly smaller memory footprint than the three independent models, 13 million parameters vs. 3×12.5 million parameters, a reduction of 100MB.

With respect to the predictions of the percentage of lung sextant affected by lesions, the median error was 1.5% with the 75'th percentile being 12%. We observe that the errors in the lower lung sextants, mean absolute

error of 7%-8%, are lower than the errors in the middle, mean absolute error of 10%-11%, and upper sextants, mean absolute error of 11%-12%. This is likely associated with the characteristics of TB where lesions are predominantly found in the apical and upper lung zones.[16] Consequentially, the variability of the percentage of lung sextant affected by lesions is higher in these regions as compared to the lower sextants as shown in Figure 8. As a result it is easier to learn to correctly predict these values for the lower sextants, hence they exhibit lower errors.

The models described above were found to be sufficiently accurate for our purposes, and were incorporated into an online content based image retrieval system as part of the TBP Radiomics portal. The system is currently available at https://rap.tbportals.niaid.nih.gov/find_similar_cxr. Figure 9 shows the system's user interface displaying the result of a query image.

The resulting table is ordered according to similarity from most to least similar. While the definition of similarity is based on gender, age, BMI and percentage of sextants affected by lesions, in some cases this may be too coarse of a category. For example, the query does not separate between drug sensitive and drug resistant cases, a task that has been shown to be hard to do based solely on CXRs.[17]

To refine the set of retrieved patient records one can use the actual case information found in the table. This includes comorbidities, resistance type, availability of pathogen genomic information and other case specific details. Once the user selects a subset cohort of interest they can analyze it using the TBP DEPOT web-enabled analytics platform[18] or any other TBP tool which facilitates analysis of cohorts created on the TBP platform (e.g. the Genomic Analysis Portal: https://gap.tbportals.niaid.nih.gov/).

6. ACKNOWLEDGMENTS

This project has been funded in part with Federal funds from the National Institute of Allergy and Infectious Diseases (NIAID), National Institutes of Health, Department of Health and Human Services under BCBB Support Services Contract HHSN316201300006W/75N93022F00001 to Guidehouse, Inc.

REFERENCES

[1] Müller, H. and Unay, D., "Retrieval from and understanding of large-scale multi-modal medical datasets: A review," *IEEE Transactions on Multimedia* **19**(9), 2093–2104 (2017).

[2] Barz, B. and Denzler, J., "Content-based image retrieval and the semantic gap in the deep learning era," in [*Pattern Recognition. ICPR International Workshops and Challenges*], 245–260 (2021).

[3] Anavi, Y., Kogan, I., Gelbart, E., Geva, O., and Greenspan, H., "Visualizing and enhancing a deep learning framework using patients age and gender for chest x-ray image retrieval," in [*SPIE Medical Imaging: Computer-Aided Diagnosis*], 978510 (2016).

[4] Hertz, D. and Schneider, B., "Sex differences in tuberculosis," *Seminars in Immunopathology* **41**(2), 225–237 (2019).

[5] Soeroto, A. Y., Pratiwi, C., Santoso, P., and Lestari, B. W., "Factors affecting outcome of longer regimen multidrug-resistant tuberculosis treatment in West Java Indonesia: A retrospective cohort study," **16**(2), e0246284 (2021).

[6] Lönnroth, K., Williams, B. G., Cegielski, P., and Dye, C., "A consistent log-linear relationship between tuberculosis incidence and body mass index," *Int J Epidemiol.* **39**(1), 149–155 (2010).

[7] Yen, Y.-F., Chuang, P.-H., Yen, M.-Y., Lin, S.-Y., Chuang, P., Yuan, M.-J., Ho, B.-L., Chou, P., and Deng, C.-Y., "Association of body mass index with tuberculosis mortality: A population-based follow-up study," *Medicine (Baltimore)* **95**(1), e2300 (2016).

[8] Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., and Summers, R. M., "ChestX-Ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in [*IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*], 3462–3471 (2017).

[9] Rosenthal, A., Gabrielian, A., Engle, E., et al., "The TB Portals: an open-access, web-based platform for global drug-resistant-tuberculosis data sharing and analysis," *Journal of Clinical Microbiology* **55**(11), 3267–3282 (2017).

[10] Shikimoto, R., Arai, Y., Yuasa, S., Gondo, Y., Yasumoto, S., Abe, Y., and Hirose, N., "Clinical course of the longest-lived man in the world: A case report," *Experimental Gerontology* **159**, 111679 (2022).

[11] Yoshizawa, T., Ishikawa, K., Nagasawa, H., Takeuchi, I., Jitsuiki, K., Omori, K., Ohsaka, H., and Yanagawa, Y., "A fatal case of super-super obesity (BMI>80) in a patient with a necrotic soft tissue infection," *Internal Medicine* **57**(10), 1479–1481 (2018).

[12] Huang, G., Liu, Z., and Weinberger, K. Q., "Densely connected convolutional networks," *CoRR* **abs/1608.06993** (2016).

[13] Chollet, F. et al., "Keras." https://keras.io (2015).

[14] Lowekamp, B. C., Chen, D. T., Ibáñez, L., and Blezek, D., "The design of SimpleITK," *Frontiers in Neuroinformatics* **7** (2013).

[15] Jabbour, S., Fouhey, D., Kazerooni, E., Sjoding, M. W., and Wiens, J., "Deep learning applied to chest X-rays: Exploiting and preventing shortcuts," in [*Proceedings of the Machine Learning for Healthcare Conference, MLHC*], 750–782 (2020).

[16] Nachiappan, A. C., Rahbar, K., Shi, X., Guy, E. S., Mortani Barbosa, E. J., Shroff, G. S., Ocazionez, D., Schlesinger, A. E., Katz, S. I., and Hammer, M. M., "Pulmonary tuberculosis: Role of radiology in diagnosis and management," *RadioGraphics* **37**(1), 52–72 (2017).

[17] Karki, M., Kantipudi, K., Yang, F., Yu, H., Wang, Y. X. J., Yaniv, Z., and Jaeger, S., "Generalization challenges in drug-resistant tuberculosis detection from chest X-rays," *Diagnostics* **12**(1) (2022).

[18] Gabrielian, A., Engle, E., Harris, M., Wollenberg, K., Juarez-Espinosa, O., Glogowski, A., Long, A., Patti, L., Hurt, D. E., Rosenthal, A., and Tartakovsky, M., "TB DEPOT (data exploration portal): A multi-domain tuberculosis data analysis resource," *PLoS One* **14**(5), e0217410 (2019).