

Is Medical Chest X-ray Data Anonymous?

Kai Packhäuser^{1,*,+}, Sebastian Gündel^{1,*,+}, Nicolas Münster¹, Christopher Syben¹, Vincent Christlein¹, and Andreas Maier¹

¹Pattern Recognition Lab, Department of Computer Science, Friedrich-Alexander University Erlangen-Nürnberg, 91058 Erlangen, Germany

*{kai.packhaeuser, sebastian.guendel}@fau.de

+these authors contributed equally to this work

ABSTRACT

With the rise and ever-increasing potential of deep learning techniques in recent years, publicly available medical datasets became a key factor to enable reproducible development of diagnostic algorithms in the medical domain. Medical data contains sensitive patient-related information and is therefore usually anonymized by removing patient identifiers, e. g., patient names before publication. To the best of our knowledge, we are the first to show that a well-trained deep learning system is able to recover the patient identity from chest X-ray data. We demonstrate this using the publicly available large-scale ChestX-ray14 dataset, a collection of 112,120 frontal-view chest X-ray images from 30,805 unique patients. Our verification system is able to identify whether two frontal chest X-ray images are from the same person with an AUC of 0.9940 and a classification accuracy of 95.55%. We further highlight that the proposed system is able to reveal the same person even ten and more years after the initial scan. When pursuing a retrieval approach, we observe an mAP@R of 0.9748 and a precision@1 of 0.9963. Furthermore, we achieve an AUC of up to 0.9870 and a precision@1 of up to 0.9444 when evaluating our trained networks on CheXpert and the COVID-19 Image Data Collection. Based on this high identification rate, a potential attacker may leak patient-related information and additionally cross-reference images to obtain more information. Thus, there is a great risk of sensitive content falling into unauthorized hands or being disseminated against the will of the concerned patients. Especially during the COVID-19 pandemic, numerous chest X-ray datasets have been published to advance research. Therefore, such data may be vulnerable to potential attacks by deep learning-based re-identification algorithms.

Chest radiography (X-ray) is a modality that is routinely used for diagnostic procedures around the world¹. It became the most common medical imaging examination for pulmonary diseases and allows a clear investigation of the thorax². Chest X-ray imaging is therefore well-suited for diagnosing several pathologies including pulmonary nodules, masses, pleural effusions, pneumonia, COPD, and cardiac abnormalities³. It is also used for COVID-19⁴ screening, as abnormalities characteristic of those infected with the coronavirus can be detected in radiographs⁵. While chest radiography plays a crucial role in clinical care, discovering certain diseases and abnormalities in chest radiographs can be a challenging task for radiologists, which potentially results in undesirable misdiagnoses⁶. Therefore, computer-aided detection (CAD) systems based on deep learning (DL)⁷ techniques have been developed in recent years to facilitate radiology workflows. These systems, characterized by their enormous benefits, can be utilized for a wide range of applications, e. g., for the automatic recognition of abnormalities in chest radiographs^{3,8} and the detection of tumors in mammography⁹. Some techniques even show the potential to exceed human performance¹⁰. However, the CAD systems are only treated as an additional source to support the radiologists and to increase certainty in their reading decisions.

On the one hand, the large variety of medical applications allows DL to grow and tackle real-life problems that were previously not solvable or improving solutions offered by traditional machine learning methods⁷. On the other hand, DL is a data-driven approach and well-known for its need for big data to train the neural networks^{11,12}. For these reasons, a vast amount of medical datasets have been published in recent years that enable researchers to develop diagnostic algorithms in the medical field in a reproducible way¹³. These include several large-scale chest radiography datasets, e. g., the *CheXpert*¹⁴, the *PLCO*¹⁵ and the *ChestX-ray14*¹⁶ datasets. But especially during the COVID-19 pandemic^{17,18}, the number of publicly available chest radiography datasets increased rapidly. A few selected examples are the *COVID-19 Image Data Collection*¹⁹, the *Figure 1 COVID-19 Chest X-ray Dataset Initiative*²⁰, the *ActualMed COVID-19 Chest X-ray Dataset Initiative*²¹, and the *COVID-19 Radiography Database*²².

Chest radiography datasets typically consist of two parts: First, the image data itself, which provides clinical information about the anatomical structure of the thorax. Second, the associated metadata, which contains sensitive patient-related information that is either stored in a separate file or embedded directly in the images²³. Proper data anonymization constitutes an important step when preparing medical data for public usage to ensure that a patient's identity cannot be revealed in publicly

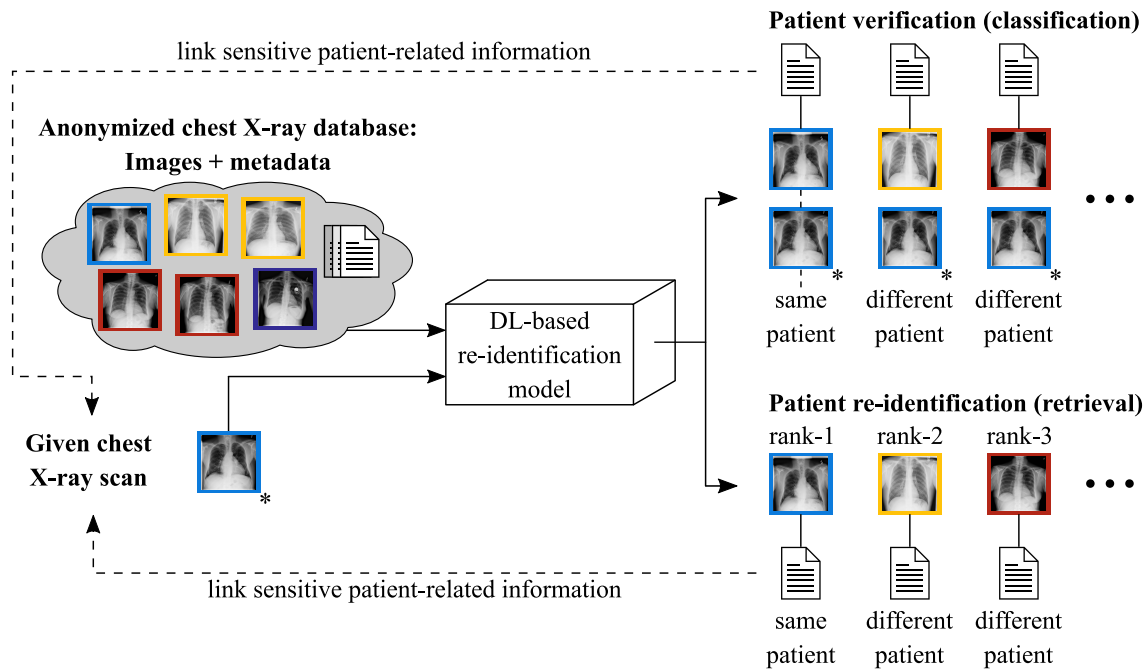


Figure 1. General problem scenario: Comparing a given chest radiograph to publicly available dataset images by means of DL techniques would either result in discrete labels indicating whether or not the dataset images belong to the same patient as the given radiograph (verification scenario) or yield a ranked list of the most similar radiographs related to the given scan (retrieval scenario). Images belonging to the same patient are highlighted with the same color. The given radiograph is marked with an asterisk. The shown cases would enable a potential attacker to link sensitive patient-related information contained in the dataset to the image of interest.

available datasets²³. In practice, any personally identifiable information is attempted to be removed from the data before it is shared. These objectives and requirements are specified, e. g., by the Health Insurance Portability and Accountability Act (HIPAA)²⁴ in the United States or the General Data Protection Regulation (GDPR)²⁵ in Europe.

In 2017, Google entered into a project with the National Institutes of Health (NIH) to publish a dataset containing 100,000 chest radiographs. However, the release was canceled two days before publication after Google was informed by the NIH that the radiographs still contained personal information which indicates that the data was incorrectly anonymized^{26,27}. This major incident highlights that many potential pitfalls can arise when clinical and technological institutions collect and share large medical datasets to revolutionize health-care.

In the past, various data de-identification techniques have been proposed, including commonly-used methods like pseudonymization²⁸ and k -anonymity²⁹. Pseudonymization describes a technique that replaces a true identifier, e. g., the name or the patient identification number by a pseudonym that is unique to the patient but has no relation to the person²⁸. However, pseudonymization is a rather weak anonymization technique as the patient's identity may still be revealed, e. g., by cross-referencing with other publicly available datasets. In contrast, k -anonymity modifies the data before sharing in such a way that every sample in the published dataset can be associated with at least k different subjects. In this way, the probability of performing identity disclosure is limited to at most $1/k$ ^{29,30}. Nevertheless, when background knowledge is available, k -anonymity is susceptible to many attacks.

To date, little attention has been paid to the possibility of re-identifying patients in large medical datasets by means of DL techniques. However in theory, medical data disclosure, as illustrated in Figure 1, could be facilitated for potential attackers by using suitable DL approaches. Consider a publicly available dataset that is supposedly anonymized but contains further sensitive patient-related information, e. g., diagnosis, treatment history, and clinical institution. If a radiograph of known identity is accessible to a potential attacker and a properly working verification or re-identification model exists, then the model could be used to compare the given radiograph to each image in the dataset which would essentially result in a set of images belonging to the same patient (patient verification) or yield a ranked list of the most similar images to the given radiograph (patient re-identification). In this way, the patient's identity may be linked to sensitive data contained in the dataset. As a result, more patient-related information may have been leaked, highlighting the enormous data security and data privacy issues involved.

Table 1. Overview of the obtained verification results for our experiments using varying training set sizes N_s at different learning rates η . Moreover, different data handling techniques were used (FTS: Fixed training set; RNP: Randomized negative pairs). For each experiment, the training sets were balanced with respect to the amount of positive and negative image pairs. In this table, we present the area under the curve (AUC) (together with the lower and upper bounds of the 95 % confidence intervals from 10,000 bootstrap runs), the accuracy, the specificity, the recall, the precision, and the F1-score. Bold text emphasizes the overall highest AUC value.

Data handling	N_s	η	AUC + 95 % CI	Accuracy	Specificity	Recall	Precision	F1-score
FTS	100,000	10^{-3}	0.8610 <small>0.8632 0.8588</small>	0.7782 <small>0.7808 0.7755</small>	0.7710 <small>0.7746 0.7673</small>	0.7853 <small>0.7889 0.7818</small>	0.7742 <small>0.7779 0.7706</small>	0.7797 <small>0.7826 0.7769</small>
	200,000	10^{-3}	0.9448 <small>0.9461 0.9435</small>	0.8743 <small>0.8763 0.8722</small>	0.8685 <small>0.8715 0.8656</small>	0.8800 <small>0.8829 0.8771</small>	0.8700 <small>0.8729 0.8671</small>	0.8750 <small>0.8772 0.8728</small>
	400,000	10^{-4}	0.9587 <small>0.9599 0.9575</small>	0.8755 <small>0.8775 0.8734</small>	0.9290 <small>0.9313 0.9268</small>	0.8219 <small>0.8252 0.8186</small>	0.9205 <small>0.9231 0.9180</small>	0.8684 <small>0.8707 0.8661</small>
	800,000	10^{-4}	0.9896 <small>0.9901 0.9891</small>	0.9537 <small>0.9550 0.9524</small>	0.9541 <small>0.9559 0.9522</small>	0.9532 <small>0.9551 0.9513</small>	0.9541 <small>0.9558 0.9522</small>	0.9536 <small>0.9550 0.9523</small>
RNP	800,000	10^{-4}	0.9940 <small>0.9944 0.9937</small>	0.9555 <small>0.9567 0.9542</small>	0.9822 <small>0.9834 0.9811</small>	0.9287 <small>0.9309 0.9264</small>	0.9812 <small>0.9824 0.9800</small>	0.9542 <small>0.9556 0.9529</small>

In our work, we investigated whether conventional anonymization techniques are secure enough and whether it is possible to re-identify and de-anonymize individuals from their medical data using DL-based methods. Therefore, we considered the public ChestX-ray14 dataset¹⁶, which is one of the most widely used research datasets in the field of radiographic problems. Our algorithms are trained to determine whether two arbitrary chest radiographs can be recognized to belong to the same patient or not. Moreover, we showed that our proposed methods are able to perform a successful linkage attack on publicly available chest radiography datasets. Furthermore, this work aims to draw attention to the massive problem of releasing medical data without considering that DL systems can easily be used to reveal a patient’s identity. Therefore, we call for reconsidering conventional anonymization techniques and developing more secure methods that resist potential attacks by DL algorithms.

Patient Verification

First, we trained a siamese neural network (SNN) architecture on the ChestX-ray14 dataset to determine whether two individual chest radiographs correspond to the same patient or not. Our model was designed to process the two input images in two identical network branches, which are then combined by a merging layer. The fused information is fed through further network layers resulting in a single output score indicating the identity similarity.

Table 1 summarizes the outcomes of our evaluation. We analyzed a multitude of different experimental setups with varying learning rates η and differing balanced training set sizes N_s . Moreover, we investigated the effect of using epoch-wise randomized negative pairs (RNP) versus fixed training sets (FTS) for the entire learning procedure. When using RNP as the data handling technique, the negative image pairs were randomly constructed in each epoch, meaning that much more negative pairs could be utilized in a complete training run compared to FTS where the generated image pairs remain the same for the entire learning procedure. For all experiments on the ChestX-ray14 dataset, we used the same balanced validation and testing set with 50,000 and 100,000 image pairs, respectively, without patient overlap between any split. To assess the performance of the trained models, we performed a receiver operating characteristic (ROC) analysis by computing the AUC value together with the 95 % confidence intervals from 10,000 bootstrap runs. Moreover, we calculated the accuracy, specificity, recall, precision, and F1-score.

The results indicate that the amount of training data plays a crucial role in the patient verification task. We observe a significant performance increase as the training set size grows. For instance, when using a subset of 100,000 image pairs for training, we obtain an AUC value of 0.8610. In contrast, by enlarging the training set size to 800,000 image pairs (i. e. the total number of 400,000 positive image pairs combined with 400,000 negative pairs), we receive an AUC score of 0.9896. These findings have been visualized in the ROC curves shown in Figure 2 which illustrates the effect of the training set size on the verification performance when using fixed training sets. Note that Table 1 only shows the best experiments per training set size N_s . Additional experiments were conducted to investigate the effect of the learning rate (LR). The corresponding results are provided in a separate table in the appendix (see Table 5).

We also observed that randomly constructing the negative image pairs in each epoch led to further improvements in the final model performance. By using this data handling technique, we achieved our overall best results. The respective outcomes are reported in Table 1. When training our network architecture with a total of 800,000 training samples with epoch-wise randomly constructed negative pairs, the AUC score improved from 0.9896 to 0.9940. Besides, the other reported evaluation metrics apart from the recall also increased compared to the results achieved by the model trained with the fixed set. Figure 3 depicts the confusion matrix resulting from our best-trained model listed in Table 1 (last row), thus giving clear insights into the patient verification performance.




Figure 2. ROC curves for different training set sizes N_s and a fixed LR of $\eta = 10^{-4}$. During training, the fixed data handling technique was employed.

		Predicted		Total
		neg	pos	
Actual	neg	49,111 (TN)	889 (FP)	50,000
	pos	3,566 (FN)	46,434 (TP)	50,000
Total		52,677	47,323	100,000

Figure 3. Confusion matrix corresponding to the best experiment shown in Table 1 (last row) giving clear insights into the performance of our trained model.

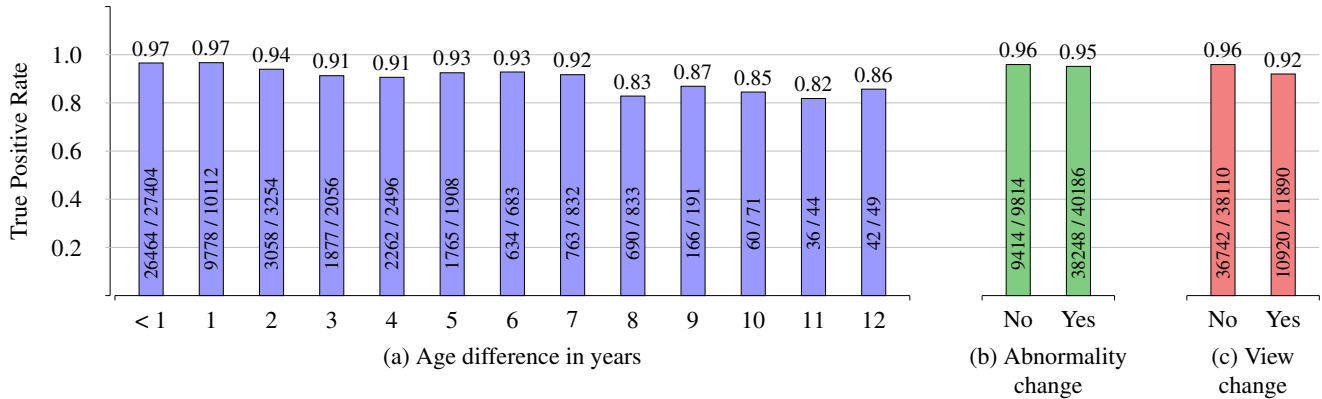


Figure 4. True positive rates (TPRs) for image pairs with (a) age differences, (b) with changes in the disease pattern, and (c) with changes in the projection view. The absolute numbers of true positives and overall positives are given for each bin. Note that the number of image pairs with age differences of more than 12 years is comparatively small, which is why the corresponding TPRs are neglected in this figure.

We also analyzed how our best model behaves when comparing images of the same patient where the acquisition dates are several years apart. The results are illustrated in Figure 4a. We received a TPR of 0.97 for image pairs that had small age differences of one year or less. As the age variation between the follow-up images and the initial scan increases, we observe a slight decrease in the TPR values. Nevertheless, our model still shows competitive results even if the patient’s age in two images differs by several years. Even for an age difference of twelve years, we can verify that two images belong to the same patient by 86%. At this point, we want to mention that the NIH ChestX-ray14 dataset contains a few images with unrealistic age information about the patient, e. g., an age of 414 years. With only 16 such cases, the proportion of images with unreasonable age information is very low, accounting for only 0.01% of the total amount of images. Since incorrect information about a patient’s age does not have a negative influence on our training strategy, we have not excluded such images for our experiments, meaning that there could potentially occur a few positive pairs with large follow-up intervals during training. However, the test set on which Figure 4 is based does not contain any positive image pairs with unreasonable follow-up intervals. We only report the TPRs for image pairs with follow-up intervals of up to 12 years in Figure 4a as the number of pairs with larger intervals is relatively small.

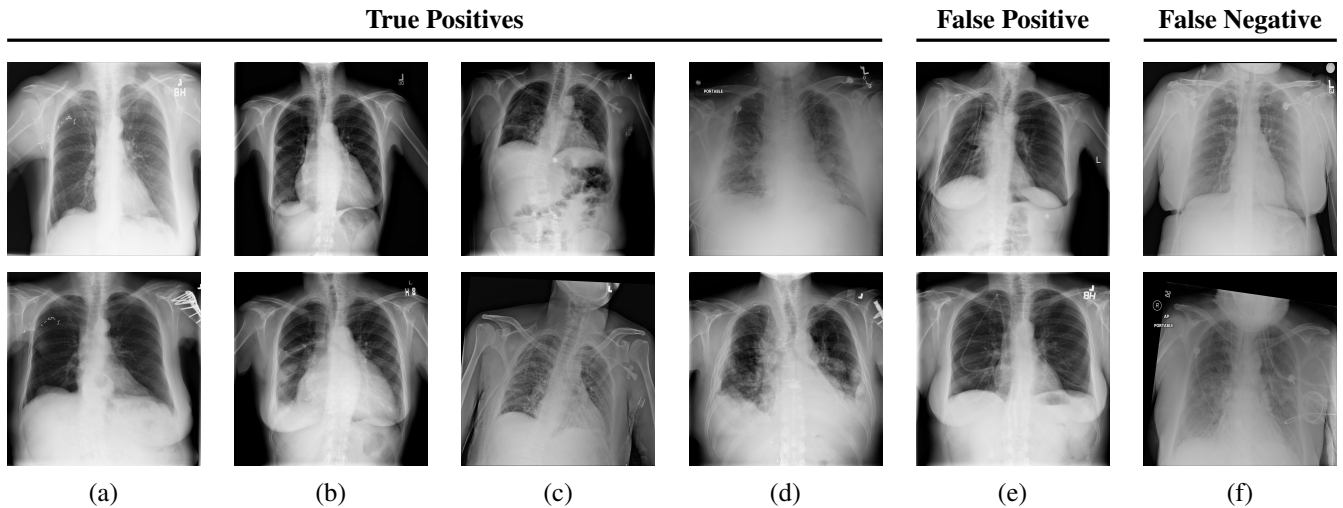


Figure 5. Exemplary image pairs are classified by our best performing verification model. Each column represents one image pair. The first four columns (a)–(d) show true positive classifications. The last two columns (e) and (f) depict a false positive and a false negative classification, respectively.

Table 2. Comparison of the verification performance on two different subsets of the ChestX-ray14 dataset that either contain foreign material or not (first two rows). Furthermore, we show the verification results for the CheXpert dataset and the COVID-19 Image Data Collection (last two rows). We present the AUC (together with the lower and upper bounds of the 95 % confidence intervals from 10,000 bootstrap runs), the accuracy, the specificity, the recall, the precision, and the F1-score.

Dataset	Subset	AUC + 95 % CI		Accuracy		Specificity		Recall		Precision		F1-score	
ChestX-ray14	w/ foreign material	0.9970	0.9993 0.9938	0.9796	0.9898 0.9679	0.9854	0.9971 0.9708	0.9738	0.9886 0.9556	0.9853	0.9970 0.9708	0.9795	0.9896 0.9678
	w/o foreign material	0.9972	0.9999 0.9909	0.9862	0.9954 0.9748	0.9908	0.9999 0.9771	0.9817	0.9957 0.9624	0.9907	0.9999 0.9763	0.9862	0.9956 0.9741
CheXpert	-	0.9870	0.9884 0.9855	0.9440	0.9475 0.9405	0.9629	0.9669 0.9586	0.9250	0.9308 0.9193	0.9614	0.9656 0.9572	0.9429	0.9466 0.9391
COVID-19	-	0.9763	0.9825 0.9696	0.9180	0.9315 0.9037	0.9780	0.9884 0.9677	0.8579	0.8820 0.8327	0.9750	0.9858 0.9625	0.9127	0.9275 0.8971

Additionally, we investigated the model’s verification capability in the case of new abnormality patterns appearing in follow-up scans that did not occur in previously acquired chest radiographs. Figure 4b shows that regardless of the abnormality, we nearly observe no decline in the TPR values, emphasizing the robustness of our trained SNN architecture. Furthermore, Figure 4c illustrates that changes in the projection view (e. g., one image taken in the anterior-posterior position and the other image acquired using the posterior-anterior view) hardly lead to any deterioration in the performance.

Moreover, we perform a qualitative evaluation where we visually inspect some exemplary image pairs evaluated using our best-performing verification model. In Figure 5, we show four true positive (TP) classifications (a)–(d), one pair that has been classified as a false positive (FP) (e), and one example for a false negative (FN) image pair (f). The shown images clearly illustrate the high technical variance present in the ChestX-ray14 dataset. The first image pair (a) shows two images belonging to the same patient with a difference of seven years. Clear differences in pixel intensities and lung shape are observed. However, both images belong to the same person, cf. the small vascular clips in the area of the upper right lung. Also, image pairs with large difference in scaling (b) or rotation (c) are verified correctly. Our model is also robust to the patients’ pathology: While the upper image of (b) shows characteristics of pneumothorax, the patient suffered from cardiomegaly, effusion, and masses in the lower image, according to the provided annotations. Similarly in (c), where the upper image indicates the presence of infiltration and pneumothorax, whereas the lower scan shows signs of infiltration and nodules. Figure 5e shows an exemplary image pair that has falsely been classified as positive. Conversely, (f) depicts a positive image pair that has been incorrectly classified as negative. To visually demonstrate which parts of the images are responsible for the verification task, we applied a siamese attention mechanism³¹ to our network architecture which utilizes the Grad-CAM algorithm³². The obtained attention maps can be seen in Figures 7 and 8 in the appendix. They clearly indicate that the human anatomy, especially the shape of the lungs and ribs, is the driving factor for the network decisions.

To investigate how foreign material (see Figure 5a) affects the verification performance, we evaluated our trained network on two small manually created subsets of around 200 images. The first one consisted only of images in which foreign material is visible, whereas the second one solely contained images without foreign material. When constructing the subsets, we selected

Table 3. Overview of the obtained results for our image retrieval experiments. In this table, we report the $mAP@R$, the R -Precision, and the Precision@1. The first 4 rows show the results on the ChestX-ray14 dataset for different image resolutions used for evaluation. The fifth row shows the outcomes on the CheXpert dataset. The last row indicates the results on the COVID-19 Image Data Collection. Bold text represents the overall highest performance metrics.

Dataset	Input dimensions	$mAP@R$	R -Precision	Precision@1
ChestX-ray14	1024×1024	0.9748	0.9763	0.9963
	800×800	0.9709	0.9726	0.9958
	512×512	0.9572	0.9601	0.9945
	224×224	0.7730	0.7979	0.9756
CheXpert	original	0.8001	0.8148	0.9444
COVID-19	original	0.8569	0.8707	0.8821

Table 4. Comparison of the re-identification performance on two different subsets (ChestX-ray14) that either contain foreign material or not. We report the $mAP@R$, the R -Precision, and the Precision@1.

Subset	$mAP@R$	R -Precision	Precision@1
w/ foreign material	0.9925	0.9925	> 0.9999
w/o foreign material	> 0.9999	> 0.9999	> 0.9999

the patients at random and then assigned the corresponding patient images to the respective subset after visual assessment. Furthermore, we ensured that no more than 5 images were used per patient. Table 2 summarizes the results indicating that the patient verification works with high performance regardless of the occurrence of foreign material. We even observe a slight improvement in performance for the subset where no foreign material is visible in the images.

Finally, to analyze whether our trained model is able to generalize to other datasets which have not been used during training, we evaluated our network on the CheXpert dataset and the COVID-19 Image Data Collection. The results are summarized in the last two rows of Table 2. It can be seen that our network still yields high AUC values of 0.9870 (CheXpert) and 0.9763 (COVID-19) for the verification task although the model has not been fine-tuned on the respective datasets. Also the other presented evaluation metrics show competitive values without deteriorating too much.

Patient Re-identification

For our patient re-identification experiments, we trained another SNN architecture on the ChestX-ray14 dataset. In contrast to the verification model, we omitted all the layers from the merging layer onwards. The main objective was to learn appropriate feature representations instead of directly determining whether the inputs belong to the same patient or not. After training the network, we used the ResNet-50 backbone as a feature extractor for the actual image retrieval task. By computing the Euclidean distance between the embeddings of the query image and each other image, we obtained for each query image a ranked list of its most similar images in terms of identity.

The results of the corresponding image retrieval experiments are summarized in Table 3. When using the original image size of 1024×1024 pixels for evaluation, we obtain a precision@1 of more than 99 % showing that the closest match nearly always is the same patient. The high mean average precision at R ($mAP@R$) of about 97 % further depicts that most of the most similar images are correctly identified. We observe a slight decrease in performance as the image size was reduced. Nevertheless, when the images were downsampled to a resolution of 512×512 pixels, we still obtain high performance values. When the image size was reduced too aggressively, e. g., to 224×224 pixels, the $mAP@R$ and the R -Precision rates drop. Yet, we still observed a high Precision@1 of more than 97 %.

Similar to the experiments in the patient verification section, we evaluated our best-trained re-identification model on two small subsets, one of which only contained images with visible foreign material and the other consisted exclusively of images without the presence of foreign material. The obtained results are presented in Table 4. As can be seen, we achieve high performance values for both subsets. Thus, we hypothesize that our outcomes are independent of foreign material which may occur only for specific patients.

Lastly, we analyzed the re-identification performance on the CheXpert dataset and the COVID-19 Image Data Collection. As can be seen in the last two rows of Table 3, we also obtain high retrieval values although we haven't performed any fine-tuning on both datasets, which demonstrates the feasibility of the trained re-identification network on previously unseen datasets.

Discussion and Conclusion

In this paper, we investigated the patient verification and re-identification capabilities of DL techniques on chest radiographs. We have shown that well-trained SNN architectures are able to compare two individual frontal chest radiographs and reliably predict whether these images belong to the same patient or not. Moreover, we have shown that DL models have the potential to accurately retrieve relevant images in a ranked list. Our models have been evaluated on the publicly available ChestX-ray14 dataset and showed competitive results with an AUC of up to 0.9940 and classification accuracy of more than 95 % in the verification scenario and an $mAP@R$ of 97 % and a precision@1 of about 99 % in the image retrieval scenario. Especially the fact that basic SNNs have the capability to re-identify patients despite potential age differences, disease changes or differing projection views demonstrated the effectiveness of DL techniques for this task. However, note that the shown results were obtained empirically, i. e. they do not necessarily reflect true measures of certainty.

As shown in Figure 5, the used dataset suffers from a high technical variance which may occur due to various windowing techniques applied to the images. In a real-life scenario, the resulting variations in image contrast and brightness could be significantly mitigated by using dynamic normalization approaches³³. Furthermore, we believe that variations in rotation and scaling can be counteracted by appropriate alignment algorithms. Nevertheless, even without such pre-processing steps, we were able to show that patient matching for chest radiographs is possible with a high performance by using DL techniques.

Moreover, we hypothesize that special noise patterns characteristic for unique patients appear in the images which might unintentionally improve the re-identification performance. For example, the initial anonymization strategy may be biased towards the clinical institution and, therefore, also towards follow-up images. To get a better impression of the re-identification capability of our SNN architecture, we also intend to investigate other datasets which show less or ideally no correlation between potential noise patterns and the patient identity. Therefore, further research on multiple datasets should ideally be considered. For our experiments, we already evaluated our models on two completely different datasets, the CheXpert dataset and the COVID-19 Image Data Collection. While the evaluation metrics are lower, we still obtain AUC scores of over 97 % (COVID-19) and 98 % (CheXpert) and precision@1 values of more than 88 % (COVID-19) and 94 % (CheXpert) without fine-tuning on these datasets. This indicates that patient verification and re-identification is also applicable for data that was acquired in various hospitals around the world where other pre-processing steps may be taken before data publication compared to the ChestX-ray14 dataset.

The COVID-19 Image Data Collection is very heterogeneous, containing, e. g., images of different sizes, both gray-scale and color images, and images with visible markers, arrows or date displays. For our experiments, only those images in the COVID-19 Image Data Collection were used that were acquired using the anterior-posterior or the posterior-anterior view, while images taken in the lateral position and CT scans were discarded. Apart from this, no further steps were taken to ensure the quality of the dataset. Although some of the factors mentioned above (e. g., brandings such as markers, arrows or dates) may facilitate the patient re-identification, we hypothesize that the COVID-19 Image Data Collection poses a realistic example of a public medical dataset and we therefore consider the conducted experiment as an authentic real-life application scenario.

Furthermore, we want to accentuate that our trained network architectures are able to handle non-rigid transformations that may appear between two images of the same person in the ChestX-ray14 dataset. Such deformations can occur due to various breath states in follow-up scans or due to different positioning during X-ray acquisition. Hence, the shape of the heart and lungs, or the contours of the ribs may appear deformed compared to an initial scan. The obtained results lead to the assumption that our trained SNN architectures can withstand such deformations and can therefore be used for reliable patient re-identification on chest radiographs.

We conclude that publicly available medical chest X-ray data is not entirely anonymous. Using a DL-based re-identification network enables an attacker to compare a given radiograph with public datasets and to associate accessible metadata with the image of interest. Thus, sensitive patient data is exposed to a high risk of falling into the unauthorized hands of an attacker who may disseminate the gained information against the will of the concerned patient. At this point, we want to emphasize that data leakage of this kind requires that the attacker has previously gained access to an image of a known person. This could happen, for example, through a stolen CD containing raw medical data of a specific patient, or by bribing corrupt medical staff at a radiological facility. Furthermore, data breaches due to inadequate data security measures at, e. g., healthcare institutions or health insurance companies, represent a possibility for attackers to obtain images of known patients, which could subsequently be utilized for a linkage attack as presented in our work. However, even if the attacker owns an image of an unknown identity, a re-identification model can be used to find the same patient across various datasets. Assuming multiple datasets contain the same patient but different metadata, an attacker would be able to obtain a more complete picture of the respective patient. We hypothesize that collecting patient information by this means could significantly help an attacker infer the true identity of the patient. We therefore urge that conventional anonymization techniques be reconsidered and that more secure methods be developed to resist the potential attacks by DL-based algorithms.

Potential solutions to the problems addressed in our work may be found in privacy-preserving approaches such as differential privacy (DP)^{34,35} which guarantees that the global statistical distribution of a dataset is retained while individually recognizable

information is reduced³⁶. This means that an outside observer is unable to draw any conclusions about the presence or absence of a particular individual. Differentially private datasets can therefore withstand linkage attacks attempting to reveal a patient’s identity. One commonly-used technique to achieve DP is to modify the input by adding noise to the dataset^{36,37}. This however may degrade and harmfully affect the quality of the data, especially when applied to medical images. In this context, we consider differentially private synthetic medical image generation³⁸ to be a promising research area. Nevertheless, further exploration on these topics is required before general conclusions can be made.

Aside from data-perturbation-based privacy approaches, we want to mention that the use of collaborative decentralized learning protocols such as federated learning (FL)³⁹ can significantly contribute to a safer use of medical data. By training a machine learning model collaboratively without centralizing the data, the need of raw data sharing or dataset release is eliminated⁴⁰. Thus, the medical data is able to reside with its owner, e. g., the healthcare institution where the data was acquired, which resolves data governance and ownership issues³⁶. However, FL itself does not provide full data security and privacy, meaning that some risks remain unless combined with other privacy-preserving methods.

Methods

Siamese neural networks

To re-identify patients from their chest radiographs, we employ SNN architectures for both the classification and the retrieval tasks. A SNN receives two input images which are processed by two identical feature extraction blocks sharing the same set of network parameters. The resulting feature representations can then be used to compare the inputs. The concept of a SNN was initially introduced by Bromley et al.⁴¹ for handwritten signature verification. Taigman et al.⁴² applied this idea in the field of face verification and proposed the *DeepFace* system. Moreover, Koch et al.⁴³ presented an approach for one-shot learning on the Omniglot⁴⁴ and MNIST⁴⁵ datasets.

NIH ChestX-ray14 dataset

With a total of 112,120 frontal-view chest radiographs from 30,805 unique patients, the NIH ChestX-ray14¹⁶ dataset counts to one of the largest publicly available chest radiography datasets in the scientific community. Due to follow-up scans, the image collection provides an average of 3-4 images per patient. The originally acquired radiographs were published as 8-bit gray-scale PNG images with a size of 1024 × 1024 pixels. Associated metadata is available for all images in the dataset. The additional data comprises information about the underlying disease patterns (either no finding or a combination of up to 14 common thoracic pathologies), the number of follow-up images taken, the patients’ age and gender, and the projection view (anterior-posterior or posterior-anterior) used for radiography acquisition. According to the publisher, the dataset was carefully screened to remove all personally identifiable information before release⁴⁶. Therefore, the patient names were replaced by integer IDs. Moreover, personal data in the image domain itself has been made unrecognizable by placing black boxes over the corresponding image areas.

CheXpert dataset

The CheXpert¹⁴ dataset contains 224,316 frontal and lateral chest radiographs of 65,240 patients, who underwent a radiographic examination from Stanford University Medical Center between October 2002 and July 2017. The originally acquired radiographs were published as 8-bit gray-scale JPG images with varying image resolutions. Note that only frontal chest radiographs were used in our work, whereas lateral images were excluded.

COVID-19 Image Data Collection

The COVID-19 Image Data Collection¹⁹ is a dataset that was created and published as an initiative to provide COVID-19 related chest radiographs and CT scans for machine learning tasks. It comprises data of 448 unique patients and a total of around 950 images with different image resolutions. In this work, only the available frontal radiographs were utilized, whereas the lateral images and CT scans were discarded.

Data preparation

Since SNN architectures require pairs of images for training and evaluation, we construct both positive and negative image pairs from the images contained in the ChestX-ray14 dataset. In this context, a positive pair consists of two images belonging to the same patient, whereas a negative pair comprises two images that belong to different patients. Mathematically, the constructed dataset can be described according to

$$\mathcal{S} = \{(\mathbf{x}_{11}, \mathbf{x}_{12}, y_1), \dots, (\mathbf{x}_{m1}, \mathbf{x}_{m2}, y_m), \dots, (\mathbf{x}_{M1}, \mathbf{x}_{M2}, y_M)\} , \text{ with } y_m \in \{0, 1\} , \quad (1)$$

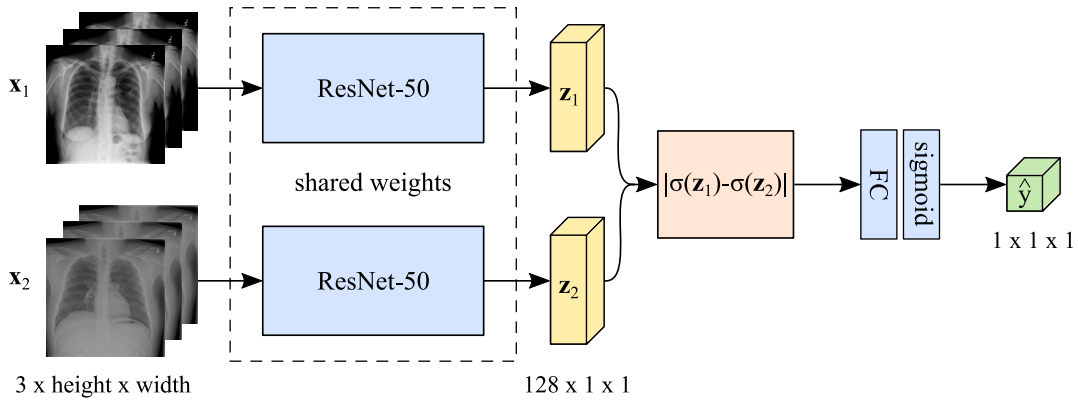


Figure 6. SNN architecture used for patient verification on the ChestX-ray14¹⁶ dataset. The feature extraction blocks (light blue) share the same set of network parameters and produce the feature representations \mathbf{z}_1 and \mathbf{z}_2 (yellow). After merging (orange) and an additional FC and sigmoid layer σ (light blue), the network yields the final output score \hat{y} (green). For our patient re-identification experiments, we used the same architecture but ejected all the layers from the merging layer onwards.

where the triplet $(\mathbf{x}_{m1}, \mathbf{x}_{m2}, y_m)$ represents one sample consisting of two images \mathbf{x}_{m1} and \mathbf{x}_{m2} , and the corresponding label y_m . M refers to the total number of samples and m denotes an iterator variable in the range of $m \in [1, M]$. The class label y_m symbolizes a binary variable that takes the value 0 for negative image pairs and 1 for positive image pairs.

To ensure that images from one patient only appear either in the training, validation, or testing set, we use the patient-wise splitting strategy. According to the official split, the data is roughly divided into 70 % training, 10 % validation, and 20 % testing. Based on this split, we construct the actual image pairs for each subset.

Offline mining

For patient verification, we follow an offline mining approach, meaning that the positive and negative image pairs are generated once before conducting the experiments. First, the positive pairs are generated by only considering the patients for whom multiple images exist in the respective subset. For each patient with follow-up images, we produce all possible tuple combinations assuming the images to be unique. By following this approach, we are able to construct a total of around 400,000 positive image pairs for our training set. The negative pairs in each subset are randomly generated and concatenated with the respective positive pairs afterwards.

Online mining

For the patient re-identification experiments, we choose an online mining approach, meaning that image pairs are formed in each batch during the training procedure. This means that the embeddings of all batch images are first computed and then subsequently used in all possible combinations as input for the loss function. Moreover, all patients with only one available image were discarded from the training set.

Patient verification

Deep learning architecture

For patient verification, the used SNN architecture (see Figure 6) receives two images \mathbf{x}_1 and \mathbf{x}_2 of size $3 \times 256 \times 256$. Both inputs are processed by a pre-trained ResNet-50 incorporated in each network branch. In its original version, the ResNet-50 was designed to classify images into 1,000 object categories trained on the ImageNet⁴⁷ dataset. To adapt the ResNet-50 to our specific needs, we replace its classification layer with a layer consisting of 128 output neurons producing the feature representations \mathbf{z}_1 and \mathbf{z}_2 , respectively. To merge both network branches, the absolute difference of the sigmoid activations of the two feature vectors is computed. We add a fully-connected (FC) layer to reduce the dimensionality to one neuron, followed by another sigmoid activation function σ which yields the final output score $\hat{y} \in [0, 1]$.

Training strategy

The verification model is trained using the binary cross-entropy (BCE) loss. The network parameters are optimized by combining mini-batch stochastic gradient descent (SGD)^{48,49} with the adaptive moment estimation (Adam)⁵⁰ method. The batch size N_b is set to 32 in all our experiments. We use different LR to investigate their effect on the model's performance. Furthermore, we include an early stopping criterion with a patience $p = 5$, which means that the training procedure stops as soon as the network does not improve for 5 epochs. We train the architecture using input dimensions of $3 \times 256 \times 256$.

Evaluation techniques

We utilize ROC curves to visualize the trained verification models based on their performance. A ROC curve represents a two-dimensional graph in which the TPR is plotted against the false positive rate (FPR) at various threshold settings⁵¹, thus indicating how many true positive classifications can be gained as an increasing number of false positive classifications is allowed. Additionally, we calculate the AUC which reflects a proportion of the area of the unit square and will always range from 0 to 1⁵¹. The higher the AUC score, the better the model’s average performance. Nevertheless, it has to be mentioned that a classifier with a high AUC might perform worse in a specific region of ROC space than a classifier with a low AUC value. Moreover, we evaluate the performance by computing the accuracy, specificity, recall, precision and F1-score. Therefore, the threshold at the output neuron is set to $t = 0.5$.

Patient re-identification

Deep learning architecture

For patient re-identification, we train a SNN architecture which receives two images \mathbf{x}_1 and \mathbf{x}_2 of size $3 \times 1024 \times 1024$. Both inputs are processed by a pre-trained ResNet-50 incorporated in each network branch. However, the network head of the used ResNet-50 is slightly modified. The average pooling layer is replaced by an adaptive average pooling layer producing feature maps of size 5×5 . In addition to the adaptive average pooling layer, an adaptive max-pooling layer is applied which also yields feature maps of size 5×5 . The outputs of the pooling layers are concatenated and processed by a 1×1 convolutional layer reducing the number of feature maps from 2048 to 100. The feature maps are then flattened, followed by two successive FC layers resulting in 128-dimensional feature representations \mathbf{z}_1 and \mathbf{z}_2 for the first and the second network branch.

Training strategy

The re-identification model is trained using the contrastive loss function⁵² which is typically utilized to achieve a meaningful mapping F from high to low dimensional space. By using the contrastive loss, the network learns to map similar inputs to nearby points on the output manifold while dissimilar inputs are mapped to distant points. Negative pairs contribute to the loss only if their distance is smaller than a certain margin m . In this work, the margin is set to $m = 1$.

For our image retrieval experiments, the SNN architecture is optimized using the SGD algorithm in combination with the 1cycle learning policy^{53,54}. When using the 1cycle LR schedule, the LR η steadily increases until it reaches a chosen maximum value and gradually decreases again thereafter. This schedule changes the LR after every single batch and is pursued a pre-defined number of epochs. The upper bound is chosen at 0.1584 with the help of a LR finder. The lower bound is set to 0.0063. The L2 regularization technique is used with a decay factor of 10^{-5} . Moreover, the batch size is adjusted to 32. We optimize the SNN architecture by first training the adapted network head of the incorporated ResNet-50 for 30 epochs with all other parameters being frozen. Then, the complete architecture is trained for another cycle, this time consisting of 50 epochs.

Since the batch size limits the task of constructing informative positive and negative pairs in the online mining strategy, the concept of cross-batch memory⁵⁵ is utilized to generate sufficient pairs across multiple mini-batches. This concept is based upon the observation that the embedding features generally tend to change slowly over time. This “slow drift” phenomenon allows the use of embeddings of previous iterations that would normally be considered out-dated and discarded. For our experiments, a memory size of 128 is chosen, meaning that the last 4 batches are considered for mining.

Evaluation techniques

To evaluate the re-identification performance of our trained model, several metrics are computed. R -Precision represents the precision at R , where R denotes the number of relevant images for a given query image. In other words, if the top- R retrieved images show r relevant images, then R -Precision can be calculated from Equation (2). Note that this value is then averaged over all query samples. Precision@1 constitutes a special case and evaluates how many times the top-1 images in the retrieved lists are relevant.

$$R\text{-Precision} = \frac{r}{R} \quad (2)$$

To further consider the order of the relevant images within the retrieved list, the mean average precision at R ($mAP@R$) is computed according to Equation (3). The $mAP@R$ denotes the mean of the average precision scores at R ($AP@R$) over all Q query images. The $AP@R$ (see Equation (4)) is the average of the precision values over all R relevant samples, where $P@i$ refers to the precision at rank i and $rel@i$ is an indicator function which equals 1 if the sample is relevant at rank i and 0 if it is not relevant.

$$mAP@R = \frac{1}{Q} \sum_{i=1}^Q AP_i@R \quad (3) \quad AP@R = \frac{1}{R} \sum_{i=1}^R P@i \times rel@i \quad (4)$$

Data availability

The NIH ChestX-ray14 dataset used throughout the current study is available via Box at <https://nihcc.app.box.com/v/ChestXray-NIHCC>. The COVID-19 Image Data Collection is available on GitHub at <https://github.com/ieee8023/covid-chestxray-dataset>. The CheXpert dataset can be requested at <https://stanfordmlgroup.github.io/competitions/chexpert>.

Code availability

The code used to train and evaluate both the patient verification and the patient re-identification models will be made available after acceptance. Correspondence and requests for materials should be addressed to K.P. or S.G.

Acknowledgements

The research leading to these results has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (ERC grant no. 810316).

Author contributions

A.M., V.C., S.G., and C.S. conceived the main idea. K.P. and N.M. performed the experiments and the evaluation. K.P. wrote the main part of the manuscript. S.G. offered continuous support during the experiments and the writing process. S.G., C.S., V.C., and A.M. provided expertise through intense discussions. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

References

1. Maier, A., Steidl, S., Christlein, V. & Hornegger, J. *Medical Imaging Systems: An Introductory Guide*, vol. 11111 (Springer, 2018).
2. Raoof, S. *et al.* Interpretation of Plain Chest Roentgenogram. *Chest* **141**, 545–558 (2012).
3. Gündel, S. *et al.* Learning to recognize Abnormalities in Chest X-Rays with Location-Aware Dense Networks. In *Iberoamerican Congress on Pattern Recognition*, 757–765 (Springer, 2018).
4. World Health Organization (WHO). Coronavirus. <https://www.who.int/health-topics/coronavirus> (2020). [Online; accessed: 21.12.2020].
5. Wang, L., Lin, Z. Q. & Wong, A. COVID-Net: a tailored deep convolutional neural network design for detection of COVID-19 cases from chest X-ray images. *Sci. Reports* **10**, 1–12 (2020).
6. Lee, C. S., Nagy, P. G., Weaver, S. J. & Newman-Toker, D. E. Cognitive and System Factors Contributing to Diagnostic Errors in Radiology. *Am. J. Roentgenol.* **201**, 611–617 (2013).
7. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
8. Gündel, S. *et al.* Multi-task Learning for Chest X-ray Abnormality Classification on Noisy Labels. *arXiv: 1905.06362* (2019).
9. Akselrod-Ballin, A. *et al.* A region based convolutional network for tumor detection and classification in breast mammography. In *Deep learning and data labeling for medical applications*, 197–205 (Springer, 2016).
10. Rajpurkar, P. *et al.* CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning. *arXiv: 1711.05225* (2017).
11. Roh, Y., Heo, G. & Whang, S. E. A Survey on Data Collection for Machine Learning: A Big Data - AI Integration Perspective. *IEEE Transactions on Knowl. Data Eng.* (2019).
12. Maier, A., Syben, C., Lasser, T. & Riess, C. A gentle introduction to deep learning in medical image processing. *Zeitschrift für Medizinische Physik* **29**, 86–101 (2019).
13. Oakden-Rayner, L. Exploring Large-scale Public Medical Image Datasets. *Acad. Radiol.* **27**, 106–112 (2020).
14. Irvin, J. *et al.* CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 590–597 (2019).
15. Gohagan, J. K., Prorok, P. C., Hayes, R. B. & Kramer, B.-S. The Prostate, Lung, Colorectal and Ovarian (PLCO) Cancer Screening Trial of the National Cancer Institute: History, organization, and status. In *Controlled Clinical Trials*, vol. 21, 251S–272S (2000).

16. Wang, X. *et al.* ChestX-ray8: Hospital-scale Chest X-ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2097–2106 (2017).
17. Bandyopadhyay, D. *et al.* Covid-19 pandemic: cardiovascular complications and future implications. *Am. J. Cardiovasc. Drugs* 1–14 (2020).
18. Spinelli, A. & Pellino, G. Covid-19 pandemic: perspectives on an unfolding crisis. *The Br. journal surgery* (2020).
19. Cohen, J. P. *et al.* COVID-19 Image Data Collection: Prospective Predictions are the Future. *arXiv: 2006.11988* (2020).
20. Chung, A. Figure 1 COVID-19 Chest X-ray Dataset Initiative. <https://github.com/agchung/Figure1-COVID-chestxray-dataset> (2020).
21. Chung, A. ActualMed COVID-19 Chest X-ray Dataset Initiative. <https://github.com/agchung/Actualmed-COVID-chestxray-dataset> (2020).
22. Rahman, T., Chowdhury, M. & Khandakar, A. COVID-19 Radiography Database. <https://www.kaggle.com/tawsifurrahman/covid19-radiography-database> (2020).
23. Willemink, M. J. *et al.* Preparing Medical Imaging Data for Machine Learning. *Radiology* **295**, 4–15 (2020).
24. Centers for Disease Control and Prevention. Health Insurance Portability and Accountability Act of 1996 (HIPAA). <https://www.cdc.gov/phlp/publications/topic/hipaa.html> (2018). [Online; accessed: 23.12.2020].
25. European Union. Complete guide to GDPR compliance. <https://gdpr.eu/> (2020). [Online; accessed: 23.12.2020].
26. O'Connor, M. Google axed release of vast x-ray dataset following NIH privacy concerns. <https://www.healthimaging.com/topics/imaging-informatics/google-axed-release-x-ray-dataset-nih-concerns> (2019). [Online; accessed: 17.12.2020].
27. Vincent, J. Google scrapped the publication of 100,000 chest x-rays due to last-minute privacy problems. <https://www.theverge.com/2019/11/15/20966460/google-scrapped-publication-100000-chest-x-rays-nih-project-2017> (2019). [Online; accessed: 17.12.2020].
28. Noumeir, R., Lemay, A. & Lina, J.-M. Pseudonymization of Radiology Data for Research Purposes. *J. Digit. Imaging* **20**, 284–295 (2007).
29. Sweeney, L. k-anonymity: A model for protecting privacy. *Int. J. Uncertainty, Fuzziness Knowledge-Based Syst.* **10**, 557–570 (2002).
30. Gkoulalas-Divanis, A. & Loukides, G. *Medical Data Privacy Handbook* (Springer, 2015).
31. Zheng, M., Karanam, S., Wu, Z. & Radke, R. J. Re-Identification with Consistent Attentive Siamese Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5735–5744 (2019).
32. Selvaraju, R. R. *et al.* Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. In *Proceedings of the IEEE International Conference on Computer Vision*, 618–626 (2017).
33. Gündel, S. *et al.* Robust Classification from Noisy Labels: Integrating Additional Knowledge for Chest Radiography Abnormality Assessment. *Med. Image Analysis* 102087 (2021).
34. Dwork, C. A Firm Foundation for Private Data Analysis. *Commun. ACM* **54**, 86–95 (2011).
35. Dwork, C., Roth, A. *et al.* The Algorithmic Foundations of Differential Privacy. *Foundations Trends Theor. Comput. Sci.* **9**, 211–407 (2014).
36. Kaissis, G. A., Makowski, M. R., Rückert, D. & Braren, R. F. Secure, privacy-preserving and federated machine learning in medical imaging. *Nat. Mach. Intell.* **2**, 305–311 (2020).
37. Sarwate, A. D. & Chaudhuri, K. Signal Processing and Machine Learning with Differential Privacy: Algorithms and challenges for continuous data. *IEEE signal processing magazine* **30**, 86–94 (2013).
38. Torfi, A., Fox, E. A. & Reddy, C. K. Differentially Private Synthetic Medical Data Generation using Convolutional GANs. *arXiv preprint arXiv:2012.11774* (2020).
39. Konečný, J., McMahan, H. B., Ramage, D. & Richtárik, P. Federated Optimization: Distributed Machine Learning for On-Device Intelligence. *arXiv preprint arXiv:1610.02527* (2016).
40. Rieke, N. *et al.* The future of digital health with federated learning. *NPJ digital medicine* **3**, 1–7 (2020).
41. Bromley, J. *et al.* Signature Verification using a “Siamese” Time Delay Neural Network. *Int. J. Pattern Recognit. Artif. Intell.* **7**, 669–688 (1993).

42. Taigman, Y., Yang, M., Ranzato, M. & Wolf, L. DeepFace: Closing the Gap to Human-Level Performance in Face Verification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1701–1708 (2014).
43. Koch, G., Zemel, R. & Salakhutdinov, R. Siamese Neural Networks for One-shot Image Recognition. In *ICML Deep Learning Workshop* (2015).
44. Lake, B. M., Salakhutdinov, R. & Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science* **350**, 1332–1338 (2015).
45. LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-Based Learning Applied to Document Recognition. *Proc. IEEE* **86**, 2278–2324 (1998).
46. National Institutes of Health (NIH). NIH Clinical Center provides one of the largest publicly available chest x-ray datasets to scientific community. <https://www.nih.gov/news-events/news-releases/nih-clinical-center-provides-one-largest-publicly-available-chest-x-ray-datasets-scientific-community> (2017). [Online; accessed: 05.01.2021].
47. Deng, J. *et al.* ImageNet: A Large-Scale Hierarchical Image Database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255 (2009).
48. LeCun, Y. A., Bottou, L., Orr, G. B. & Müller, K.-R. Efficient BackProp. In *Neural Networks: Tricks of the Trade*, 9–48 (Springer, 2012).
49. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* (MIT Press, 2016). URL: <http://www.deeplearningbook.org>.
50. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980* (2014).
51. Fawcett, T. An introduction to ROC analysis. *Pattern Recognit. Lett.* **27**, 861–874 (2006).
52. Hadsell, R., Chopra, S. & LeCun, Y. Dimensionality Reduction by Learning an Invariant Mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, vol. 2, 1735–1742 (IEEE, 2006).
53. Smith, L. N. & Topin, N. Super-Convergence: Very Fast Training of Neural Networks Using Large Learning Rates. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, vol. 11006, 1100612 (International Society for Optics and Photonics, 2019).
54. Smith, L. N. Cyclical Learning Rates for Training Neural Networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 464–472 (IEEE, 2017).
55. Wang, X., Zhang, H., Huang, W. & Scott, M. R. Cross-Batch Memory for Embedding Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6388–6397 (2020).

Supplementary Material

Table 5. Additional patient verification results: Comparison of different data handling techniques (FTS and RNP), training set sizes N_s and learning rates η . We present the AUC (together with the lower and upper bounds of the 95 % confidence intervals from 10,000 bootstrap runs), the accuracy, the specificity, the recall, the precision, and the F1-score.

Data handling	N_s	η	AUC + 95 % CI	Accuracy	Specificity	Recall	Precision	F1-score
FTS	100,000	10^{-4}	0.8509 0.8532 0.8485	0.7461 0.7487 0.7433	0.8414 0.8446 0.8383	0.6507 0.6550 0.6465	0.8040 0.8079 0.8002	0.7193 0.7226 0.7160
		10^{-5}	0.7429 0.7459 0.7399	0.6322 0.6352 0.6293	0.7924 0.7960 0.7888	0.4721 0.4764 0.4677	0.6945 0.6994 0.6896	0.5621 0.5661 0.5581
		10^{-6}	0.7257 0.7289 0.7225	0.6702 0.6731 0.6672	0.6719 0.6761 0.6677	0.6685 0.6727 0.6643	0.6708 0.6749 0.6666	0.6696 0.6729 0.6662
		10^{-7}	0.7440 0.7470 0.7409	0.6741 0.6770 0.6712	0.7002 0.7042 0.6961	0.6480 0.6521 0.6438	0.6837 0.6879 0.6795	0.6654 0.6688 0.6620
	200,000	10^{-4}	0.8886 0.8906 0.8866	0.7920 0.7946 0.7895	0.8547 0.8578 0.8516	0.7293 0.7332 0.7254	0.8339 0.8373 0.8304	0.7781 0.7811 0.7752
		10^{-5}	0.8158 0.8185 0.8132	0.6758 0.6787 0.6729	0.8574 0.8604 0.8543	0.4943 0.4985 0.4898	0.7761 0.7805 0.7715	0.6039 0.6078 0.6000
		10^{-6}	0.7615 0.7645 0.7586	0.6755 0.6784 0.6726	0.7448 0.7487 0.7410	0.6062 0.6104 0.6018	0.7038 0.7080 0.6995	0.6514 0.6549 0.6477
	400,000	10^{-3}	0.9541 0.9552 0.9529	0.8852 0.8872 0.8832	0.8655 0.8685 0.8625	0.9049 0.9074 0.9023	0.8706 0.8734 0.8677	0.8874 0.8894 0.8854
		10^{-5}	0.8518 0.8542 0.8494	0.7450 0.7477 0.7423	0.8410 0.8442 0.8379	0.6490 0.6532 0.6448	0.8032 0.8070 0.7993	0.7179 0.7212 0.7145
		10^{-6}	0.7417 0.7449 0.7386	0.6572 0.6601 0.6542	0.7152 0.7191 0.7113	0.5991 0.6033 0.5948	0.6778 0.6822 0.6735	0.6360 0.6395 0.6324
		10^{-7}	0.7306 0.7337 0.7275	0.6567 0.6596 0.6537	0.6801 0.6842 0.6761	0.6333 0.6375 0.6289	0.6644 0.6686 0.6602	0.6485 0.6519 0.6449
	RNP	800,000	10^{-3}	0.9826 0.9833 0.9820	0.9324 0.9339 0.9308	0.9393 0.9414 0.9372	0.9254 0.9278 0.9231	0.9385 0.9406 0.9364
10^{-4}			0.9940 0.9944 0.9937	0.9555 0.9567 0.9542	0.9822 0.9834 0.9811	0.9287 0.9309 0.9264	0.9812 0.9824 0.9800	0.9542 0.9556 0.9529
FTS	800,000	10^{-5}	0.9278 0.9294 0.9262	0.8339 0.8362 0.8316	0.8946 0.8973 0.8919	0.7732 0.7769 0.7696	0.8801 0.8831 0.8771	0.8232 0.8258 0.8205
		10^{-5}	0.9200 0.9217 0.9182	0.8215 0.8240 0.8192	0.8888 0.8915 0.8861	0.7543 0.7580 0.7505	0.8715 0.8746 0.8684	0.8087 0.8114 0.8058
		10^{-6}	0.8669 0.8692 0.8646	0.7752 0.7778 0.7726	0.8165 0.8199 0.8131	0.7339 0.7378 0.7300	0.7999 0.8035 0.7963	0.7655 0.7685 0.7625
		10^{-7}	0.8126 0.8152 0.8099	0.7247 0.7275 0.7220	0.7502 0.7539 0.7463	0.6992 0.7031 0.6952	0.7368 0.7407 0.7328	0.7175 0.7206 0.7144

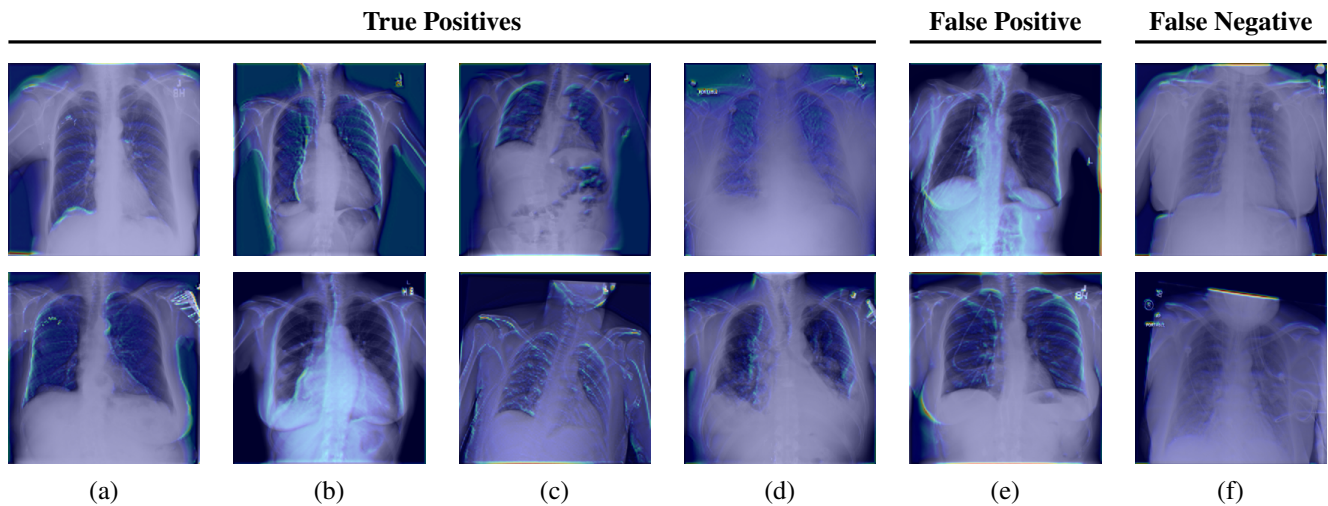


Figure 7. Grad-CAM visualizations for the first convolutional layer of the ResNet-50 incorporated in our SNN. Each column represents one image pair. The first four columns (a)–(d) show true positive classifications. The last two columns (e) and (f) depict a false positive and a false negative classification, respectively. The shown images illustrate that the anatomical structure of, e. g. the breast (cf. (a), (b), (e)), the lungs (cf. (a), (b), (c), (e)), and the heart (cf. (a), (b)) have a high impact on the final model prediction. Furthermore, it can be seen that our network focuses on the collarbones (cf. (a), (c), (d), (f)) and the ribs (cf. (b), (c), (e)). The upper images of (a) and (f) also highlight that our network pays attention to the contour of the diaphragm.

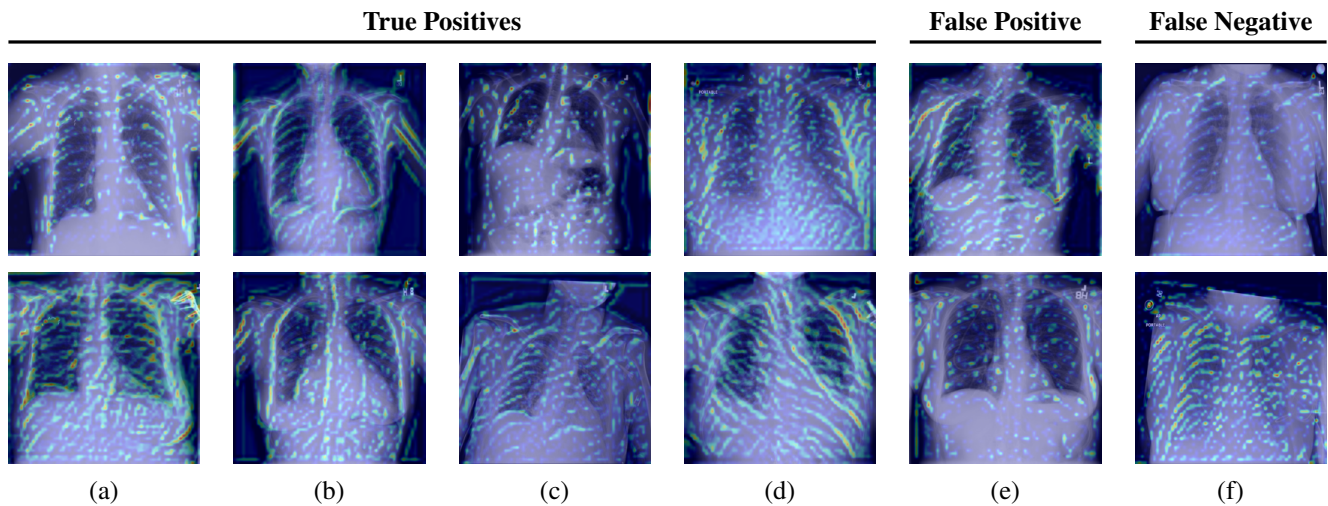


Figure 8. Grad-CAM visualizations for an intermediate convolutional layer of the ResNet-50 incorporated in our SNN. Each column represents one image pair. The first four columns (a)–(d) show true positive classifications. The last two columns (e) and (f) depict a false positive and a false negative classification, respectively. The obtained attention maps clearly illustrate that the selected network layer focuses on the ribs and the outline of the thorax.