

SPATIALLY-PRESERVING FLATTENING FOR LOCATION-AWARE CLASSIFICATION OF FINDINGS IN CHEST X-RAYS

Neha Srivathsa¹, Razi Mahmood², Tanveer Syeda-Mahmood³

¹Stanford University, ²University of California, Berkeley, ³IBM Almaden Research Center

ABSTRACT

Chest X-rays have become the focus of vigorous deep learning research in recent years due to the availability of large labeled datasets. While classification of anomalous findings is now possible, ensuring that they are correctly localized still remains challenging, as this requires recognition of anomalies within anatomical regions. Existing deep learning networks for fine-grained anomaly classification learn location-specific findings using architectures where the location and spatial contiguity information is lost during the flattening step before classification. In this paper, we present a new spatially preserving deep learning network that preserves location and shape information through auto-encoding of feature maps during flattening. The feature maps, auto-encoder and classifier are then trained in an end-to-end fashion to enable location aware classification of findings in chest X-rays. Results are shown on a large multi-hospital chest X-ray dataset indicating a significant improvement in the quality of finding classification over state-of-the-art methods.

Index Terms— Chest X-rays, deep learning network design, spatially-aware flattening.

1. INTRODUCTION

Chest X-rays are the most common imaging exams done in emergency rooms and intensive care units in hospitals. With the availability of large open source datasets labeled for selected anomalies[1, 2], they have become the focus of vigorous deep learning model research in recent years[3, 2, 4, 5]. The goal of this research is to produce automated radiology reports, which requires recognition of fine-grained findings that describe the laterality as well as location of findings (e.g. "left basal atelectasis"). While it is possible to do detection and classification with modern deep learning networks in computer vision[6, 7], fine-grained finding recognition is different as it requires the anomaly to be located within the correct anatomical reference (e.g. pneumothorax in left apex). Further, certain anomalies only appear in specific anatomical regions, requiring the modeling of these constraints (e.g. opacities are seen in lungs). Similarly, capsule networks[8] that are meant to capture relative spatial relationships of parts of an object in multiple poses are not very applicable for chest

X-rays since anomalies need to be localized within relevant anatomical areas that in turn, also need to be identified.


Large-scale fine-grained finding recognition in chest X-rays is a relatively recently posed problem. The current methods to handle such fine-grained findings use 3 main approaches: (a) incorporating an anatomical atlas for localization prior to classification[9], (b) using soft attention with gradCAM within classifiers themselves[10], and (c) direct classification using detailed finding labels that already reflect the location and laterality in the label names[5]. In a recent work, a fixed number of anatomical zones in a chest X-ray were localized through bounding boxes within atlas-guided anatomical segmentation[11]. The alignment of these bounding boxes for new chest X-rays required ad hoc registration. Later work used the training bounding boxes per anatomical region as supervision data to build a faster RCNN network[12] for detection. The inter-relationship features between regions was then extracted using a graph convolutional network and a multi-class classifier was applied to each region to classify anomaly labels within each region[9]. This two-stage approach, while reasonable, can lead to potential incorrect combinations of anomalies with anatomical locations as all regions are examined for all anomalies. In the second class of approaches, soft attention using heat maps was generated through gradCAM operators in classifier networks[10], or obtained using occlusion sensitivity as a measure of localization[13]. Others have adopted hard attention applied through regions of interest obtained by a prior region segmentation algorithm based on U-net[14] for lungs[11]. Soft attention approaches are based on the hope that the classifier is looking in the relevant region for the label which can only be verified in a post-hoc way through operators like GradCAM. The hard attention models, on the other hand, rely on the accuracy of deep learning-driven region segmentation, which is difficult even for large regions such as lungs, in the presence of severe anomalies. Both atlas-based and attention-based approaches have only been attempted on a limited class of findings. Finally, the third class of approaches uses a pure classifier but increases the granularity of the labels to cover location and laterality information within anomaly labels. Specifically, companion radiology reports were used to derive detailed fine-grained labels which were then directly learned using a custom deep

learning model designed for multi-label classification[5, 15]. As the model used whole image inputs and a flattening layer for classification that lost spatial contiguity information, the fine-grained classification performance was low (weighted average AUC of 0.73). Due to the bundled approach, an error in class label could point to gross errors in localization as well. Nevertheless, the classifier approach has been the only one so far to cover a large spectrum of findings suitable for the ultimate application of automatic report generation.

In this paper, therefore, we adopt the classifier approach for fine-grained labels, but address the limitation of existing classifier deep learning networks that lose spatial contiguity during the flattening step prior to classification. Specifically, we present a new spatially preserving deep learning network that preserves location and shape information through auto-encoding of feature maps during flattening. Results are shown on a large multi-hospital chest X-ray dataset indicating a significant improvement in the quality of finding classification over state-of-the-art methods without requiring detailed anatomy segmentation or large-scale region annotation.

2. PRESERVING SPATIAL INFORMATION DURING CLASSIFICATION

To understand the key idea behind our approach, consider the architecture of most deep learning networks for image classification, such as VGG-16[16]or Resnet101[17]. Typically, they have a feature extraction stage formed from layers of convolutional filtering and pooling operations followed by an intermix of fully connected and dropout layers for classification[16, 17]. Although the feature extraction operations using convolutional filters are applied at a pixel or super-pixel level, these are ultimately flattened into vectors before proceeding to the classification stage, resulting in the loss of spatial contiguity and layout information. This can be



(a) Result 1

Fig. 1. Illustration of loss of spatial contiguity during flattening - simple examples.

easily seen through a simple example in Figure 1, where the 2D layout makes it trivial to see the arrangement as representing the number 1 in Figure 1a, in comparison to its flattened

representation in Figure 1b. Similarly, in Figure 1c and d, the spatial adjacencies of the intensity values across a column are lost in the flattening, resulting in a loss of spatial information. *If the spatial layout information could be preserved during the flattening step, it would be possible to distinguish between location-specific anomalies better during classification.* This is the key idea proposed in this paper. Specifically, we preserve spatial information in two ways, (i) by capturing the shape information conveyed in the filtered representations, and (ii) by keeping the identities of the filters separate during the flattening. We now describe this in detail.

2.1. Spatially-preserving flattening

The convolutional filters in a deep learning network learn to detect image characteristics at specific locations in either the direct image (at the input layer) or its successive abstractions obtained through pooling operations capturing more global characteristics. The result of applying these filters to an image generates an activation map, also known as a feature map, highlighting the relevant features detected or preserved in the input by the filter. While feature maps close to the input detect small or fine-grained detail, and feature maps close to the output of the model capture more general features, each feature map can be seen as defining a shape. Figure 2 illustrates this for a Resnet50[17] architecture. Here the feature maps produced by applying the filters in Layer 2 (*conv1conv* layer) consist of 64 filters of 16x16 each. The 64 feature maps produced using an instance of Resnet50 pretrained on Imagenet from the input image of Figure 2a are shown in Figure 2b. *Since auto-encoders are known to form a low-dimensional*

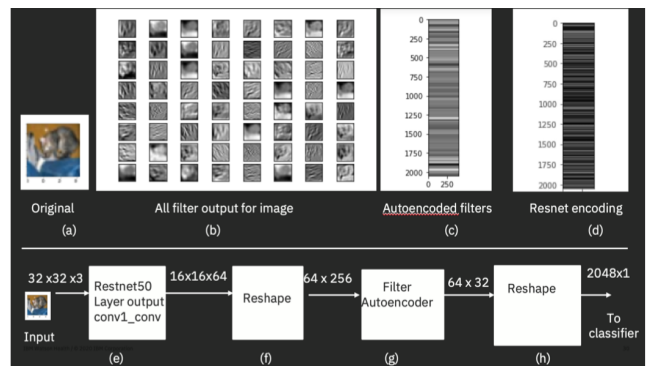


Fig. 2. Illustration of spatially preserving flattening in ResNet50v2 using the filters from the second layer.

representation or encoding of shapes, we can represent each such feature map using an auto-encoder and inherently preserve its spatial information. Further, by concatenating all the encodings of the feature maps, we form a new flattening vector that separately retains the shape information from the individual feature maps. A fixed ordering can be used among the feature maps to define the layout of their encod-

ings. Figure 2c shows the resulting encoding vector generated from the auto-encoding of feature maps of Figure 2b using a left to right, top to bottom ordering. A 32 bit auto-encoding was used for each of the feature maps resulting in a flattened vector of size $64 \times 32 = 2048 \times 1$. The traditional flattening obtained by the pre-trained ResNet50 prior to classification, which also generates a 2048×1 vector, is shown in Figure 2d. As can be seen by their intensity renderings, the spatially preserving flattening shows more discriminable intensity values, intuitively supplying the basis for improvement in subsequent classification.

2.2. Spatially-preserving deep learning network

Although all feature maps from the second layer were used in the above example, it is preferred to use the last layer of the feature extraction step to benefit from the local and global image characteristics captured in the filters. Specifically, we introduce a spatially preserving flattening module in between the feature extraction and classification layers to make a new end-to-end deep learning network as shown in Figure 3. The proposed network leverages the advantages of the feature extractor and classifier portion of a previously developed network[15]. Referring to [15] for details, we describe these stages only briefly here. The feature extractor uses a feature pyramid network to allow multiresolution analysis, and a cascade of dilated blocks with skip connections to improve convergence, while spatial dropout was used to reduce overfitting. Dilated blocks with different feature channels were cascaded with maxpooling to learn more abstract features. Finally the classification was done through a dense layer with sigmoid activation to allow multilabel classification. The parameters and dimensions are detailed in Figure 3. The last stage of the feature extractor generates 128 filters of size 64×64 . In the original architecture described in [15], global square pooling was used to flatten these filter weights into a feature vector of size $128 \times 128 = 16,384$. The new spa-

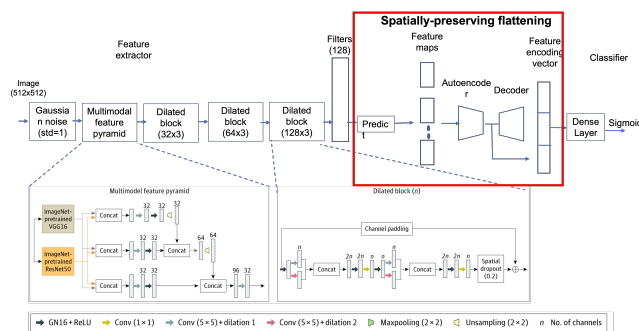


Fig. 3. Illustration of spatially preserving flattening deep learning network architecture.

tially preserving flattening (SPF) module aims to replace this 16,384 length vector with one generated from auto-encoded

Label	Images	Macro AUC		
		Ours	Network1	Detectron
Elevated right hemidiaphragm	1596	0.892	0.797	0.654
Bilateral pleural effusions	6696	0.823	0.793	0.71
Pneumothorax in the left lower lobe	4442	0.882	0.784	0.679
Right lower lobe pleural effusion	44770	0.812	0.766	0.71

Table 1. Illustration of frequently seen location-specific findings and the performance of various learning networks.

feature maps. Specifically, the new flattening module consists of (a) a predictor to predict activation maps using the 128 filters from the last stage of feature extraction, (b) an autoencoder to train on the feature maps for capturing their spatial information, and (c) a flattener that concatenates the encodings of the feature maps. Specifically, we used a 128-bit encoder for each of the 128 incoming 512×512 feature maps produced by the predictor. The resulting flattened vector was assembled by arranging the 128 encoded vectors of all features maps into a 16,384 length vector which was then given as input to the DenseNet classifier.

The overall network was trained end-to-end using a Nadam optimizer for fast convergence, with the learning rate as 2×10^{-6} . Two NVIDIA Tesla V100 graphics processing units with 16 GB memory were used for multi-graphics processing unit training with a batch size of 12 over 30 epochs. All development was done in Python using TensorFlow and Keras libraries.

3. CLASSIFICATION OF LOCATION-SPECIFIC FINDINGS IN CHEST X-RAYS

We now apply the developed deep learning network for location-specific finding classification in chest X-rays. A set of location-specific labels were catalogued for a large collection of chest X-ray images in a previous work[5] which was leveraged in our work as well. Specifically, this was a multi-institutional dataset collected from 2 reference sources, namely, MIMIC-CXR[1], NIH[2]. These X-rays showed a wide range of clinical settings, including intensive care units, urgent care, in-patient care, and emergency departments and represented a wide variety of chest X-ray findings in AP, PA and portable modes of acquisition. All data used was de-identified and covered under the secondary use of health data and informed patient consent was either waived (NIH) or obtained (MIMIC) by the data providers. An earlier work had painstakingly catalogued a full list of 457 findings in chest X-rays using text analytics on the accompanying radiology reports associated with these images[4]. In our work, we selected a subset of 57 of the 457 labels from this label set

Method	Labels	Train	Test	Macro AUC	Weighted AUC
[7]	57	75,613	20,941	0.512±0.025	0.573±0.022
[5]	457	75,613	20,941	0.729±0.001	0.716±0.002
[5]	57	75,613	20,941	0.74±0.003	0.723 ±0.012
Ours	57	75,613	20,941	0.81±0.001	0.79±0.002

Table 2. Illustration of the performance of location and laterality classification models against the selected dataset using labels. The first two columns list the method and the number of labels.

that were found to refer to either laterality or anatomical location and were most frequently found in radiology reports. A subset of these labels are shown in Table 1. Since the earlier work[4] had already verified these labels, we used the labeled dataset to train, validate, and test our network for fine-grained finding classification. We used a 70-10-20% split for training, validation and testing. Although the full dataset consisted of 335,189 chest x-rays and their reports, the dataset available for the 57 finding labels was found to contain 107,169 chest X-ray images.

4. RESULTS

By using spatially sensitive encodings from feature maps for classification, our network is able to predict the location of specific labels reliably. Figure 4 shows the type of location-specific label prediction on sample images by our network. Since the location-specific details were already included in the label, the AUC measured the combined accuracies of anatomy and anomaly classification. The weighted AUC for the 57 labels was found to be 0.81.




Fig. 4. Illustration of location specific finding classification using SPF deep learning network.

Comparison of performance

As the literature on location-specific finding classification in chest X-rays is sparse, we compared the performance of our deep learning network against two other state-of-the-art approaches, namely, those using an anatomical atlas to separate anatomy detection from anomaly location, and those that use a whole image approach to directly recognize location-specific labels. Since their networks were trained on different sets of labels, we report our evaluation on the same subset of

Encoding size	AUC
32	0.67±0.002
64	0.716±0.003
128	0.81±0.012
256	0.80±0.001

Table 3. Illustration of ablation study on the size of the auto-encoding used for capturing the shapes of filters. The train and test splits of the network are the same as in Table 2.

57 anatomy-specific findings we chose for our implementation. Specifically, for the atlas approach, we implemented a detection transformer (DETR) with ResNet-50 as base model using the implementation provided from DETR[7]. This was trained on the automatically extracted bounding box regions obtained by applying the anatomical atlas to the training and testing images as reported in [11]. Since no open source implementation of [9] released this year was provided, the above implementation closely approximates what is possible with SOTA approaches based on anatomical atlases, although the reference[9] was demonstrated for only 9 core findings. Finally, the network reported in [5] gave us a comparison to conventional non-spatially preserving encoding-based fine-grained classification. All networks used the same splits for train and test as shown in Table 2, and the same set of 57 labels. As can be seen, the average AUC obtained by our method is at least 15% higher than the comparable methods. The performance for individual fine-grained labels is shown in Table 1.

Ablation studies

The main ablation study was to see the effect of the use of spatially preserving flattening over conventional flattening. Since the comparison network of [5] used this form of flattening, the results in Table 1 and Table 2 constitute the ablation study for the flattening. We also performed an ablation study by changing the size of the auto-encoding for the feature maps and selecting the filters from different layers of the feature extraction stage. The effect of auto-encoding sizes on the average AUC performance is shown in Table 3, indicating that the size of the encoding chosen for producing the feature maps was appropriate for the image sizes being handled in our network for chest X-rays.

5. CONCLUSION

In this paper, we have presented a new spatially preserving deep learning network for location-aware finding classification in chest X-rays. Results are shown on a large multi-hospital chest X-ray dataset indicating a significant improvement in the quality of finding classification over state-of-the-art methods.

6. REFERENCES

- [1] A. E. W. Johnson, T. J. Pollard, N. R. Greenbaum, M. P. Lungren, C-Y Deng, Y. Peng, Z. Lu, R. G. Mark, S. J. Berkowitz, and S. Horng, "Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs," <https://arxiv.org/abs/1901.07042v5>, 1 2019.
- [2] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald Summers, "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*, 2017, pp. 3462–3471.
- [3] J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, C. Chute, H. Marklund, B. Haghighi, R. Ball, K. Shpan-skaya, J. Seekins, D. A. Mong, S. S. Halabi, J. K. Sandberg, R. Jones, D. B. Larson, C. P. Langlotz, B. N. Patel, M. P. Lungren, and A. Y. Ng, "Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison," in *Proceedings of the AAAI Conference on Artificial Intelligence.*, 2019, pp. 590–597.
- [4] T. Syeda-Mahmood, K. C. L. Wong, Y. Gur, J. T. Wu, A. Jadhav, S. Kashyap, A. Karargyris, A. Pillai, A. Sharma, A. Bin Syed, O. Boyko, and M. Moradi, "Chest x-ray report generation through fine-grained label learning," in *Proc. Medical Image Computing and Computer-Assisted Interaction (MICCAI)*, 2020, vol. 12262 LNCS, pp. 561–571.
- [5] Joy T. Wu, Ken C. L. Wong, Yaniv Gur, Nadeem Ansari, Alexandros Karargyris, Arjun Sharma, Michael Morris, Babak Saboury, Hassan Ahmad, Orest Boyko, Ali Syed, Ashutosh Jadhav, Hongzhi Wang, Anup Pillai, Satyananda Kashyap, Mehdi Moradi, and Tanveer Syeda-Mahmood, "Comparison of chest radiograph interpretations by artificial intelligence algorithm vs radiology residents," *JAMA Network Open*, vol. 3, October 2020.
- [6] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár, "Focal loss for dense object detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, pp. 318–327, August 2018.
- [7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko, "End-to-end object detection with transformers," in *Proc. European Conference on Computer Vision*, 5 2020, vol. 12346 LNCS, pp. 213–229.
- [8] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton, "Dynamic routing between capsules," *Advances in Neural Information Processing Systems*, vol. 2017-December, pp. 3857–3867, 2017.
- [9] Nkechinyere N. Agu, Joy T. Wu, Hanqing Chao, Ismini Lourentzou, Arjun Sharma, Mehdi Moradi, Pingkun Yan, and James Hendler, "Anaxnet: Anatomy aware multi-label finding classification in chest x-ray," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2021, pp. 804–813.
- [10] Satyananda Kashyap, Alexandros Karargyris, Joy Wu, Yaniv Gur, Arjun Sharma, Ken C. L. Wong, Mehdi Moradi, and Tanveer Syeda-Mahmood, "Looking in the right place for anomalies: Explainable ai through automatic location learning," *Proceedings - International Symposium on Biomedical Imaging*, vol. 2020-April, pp. 1125–1129, April 2020.
- [11] Joy Wu, Yaniv Gur, Alexandros Karargyris, Ali Bin Syed, Orest Boyko, Mehdi Moradi, and Tanveer Syeda-Mahmood, "Automatic bounding box annotation of chest x-ray data for localization of abnormalities," *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pp. 799–803, 2020.
- [12] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, pp. 1137–1149, June 2015.
- [13] Mohammad Tariqul Islam, Md Abdul Aowal, Ahmed Tahseen Minhaz, and Khalid Ashraf, "Abnormality detection and localization in chest x-rays using deep convolutional neural networks," May 2017.
- [14] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015, vol. 9351, pp. 234–241.
- [15] Tanveer Syeda-Mahmood, K. C. L. Wong, Joy T. Wu, Ashutosh Jadhav, and Orest Boyko, "Extracting and learning fine-grained labels from chest radiographs," in *Proc. American Medical Informatics Association (AMIA) Annual Symposium*, November 2020.
- [16] Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition," <https://arxiv.org/abs/1409.1556>, 2015.
- [17] Sun J. He K, Zhang X, Ren S, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 248–255.