

SA-ResNet50: A Four-Class Classification Model for Lung Diseases in Chest X-rays

Qi Deng

*School of Artificial Intelligence, Southeast University, Nanjing, China
213233438@seu.edu.cn*

Abstract. Deep learning in medical image analysis is becoming more important for diagnosing disease, especially when it comes to spotting stuff like COVID-19 and pneumonia in CXR pictures of someone's chest. But it is inevitable that there will be some limitations for CNN in terms of CXR images because it always has blurred and unclear edges, strange shape and low contrast with other tissues around the lesion area. And these usually cause wrong localizations and some important informations will be lost when extracting features. So it is hard for classification models to do well. In order to solve above problems, I propose the SA-ResNet50 architecture in which a spatial attention block is introduced inside residual bottleneck blocks of ResNet-50. This work contributes in some ways. First is to improve classification accuracy by improving the localization and representation of important lesion area. Second, in a complete pneumonia data set, the Accuracy is 0.9847, the Precision is 0.9826, the Recall is 0.9827, the F1-Score is 0.9826, and the Specificity is 0.9939. By doing full-comparison experiment about every baseline model and detailed ablation studies, I will strictly demonstrate it is effective and better than others. The experimental results not only validate the efficacy of this model, but also propose new ways and references to deal with similar issues in other fields related to recognizing pictures of medicine.

Keywords: Deep Learning, Spatial Attention Mechanism, Chest X-ray, Lung Disease Classification, ResNet-50

1. Introduction

March 2020: The World Health Organisation (WHO) declared that COVID-19 was a global pandemic [1]. During the pandemic, when there were respiratory symptoms, doctors often used CXR for patient diagnosis of whether it was corona virus or not. CXR does not only have a fast examination time, it is also far cheaper to do on average for every examination compared to other imaging modalities such as CT and MRI, therefore CXR remains one of the most common medical imaging techniques today [2,3]. Therefore, there is a huge amount of public dataset about CXR examinations, which provides a lot and useful data for training the model.

In the field of medical image analysis, Convolutional Neural Networks(CNNs) have made great achievements [4]. They can auto learn image's features, surpassing old fashioned hand crafted feature models. All kinds of cnn architectures have been used in medical picture recognition. For example, AlexNet uses dermoscopic pictures to find out if there is an illness, such as melanoma; VGG checks the picture of a person's eye using a camera that has blood sugar problems, it does this so the person

can get better fast; ResNet looks at the sick parts of a patient when they take their tummy out for surgery, and guesses what kind of sickness might be there [5-7].

But using the same kind of general purpose models on CXR images is proving hard. The pulmonary lesions like GGO and consolidation, the edge is always blurry and not regular. There is no definite boundary between the transition zone and normal lung tissue. This is imprecisely meaning to generalized model it can easily seen features as noise by doing convolutions, so there will be lots of problem such as bad feature extracting, wrong localizing therefore poor classifying. The challenge is very clinical; GGOs could be anything between benign inflammation and early stage adenocarcinoma so it's really hard but important to get them right.

In order to solve those problems, I use an attention method in ResNet-50. The attention mechanism in computer vision is inspired by human visual attention. It allows a neural network to "focus" on the more important parts of an input image of data and give them more weight towards the end of providing useful info for making decisions. The most commonly seen attention methods are Channel Attention (CA), Spatial Attention (SA), and the Convolutional Block Attention Module (CBAM) [8-10]. Channel attention is over the channel dimension of a feature map to learn which channels are more important on their own, then adjust by these weights. In a CNN, different channels would correspond to different feature detectors, e.g. for texture, color, edge etc. For Channel attention, it makes the model decide which features are currently more important to do what task and that can be described as "What you should focus on". In contrast, it is spatial attention which decides what parts of the feature map deserve to be paid attention, i.e., where should we look at? It creates a 2D weight mask and use this weight to modify the old feature map.

Considering that the main problem with recognizing lesions in CXR images is their spatial location, our approach is based on the idea that improving how the model perceives space is more important and beneficial than improving its understanding of channels. And thus our choice of a spatial attention mechanism is not arbitrary; rather, it is a deliberate and hypotheses driven engineering decision designed specifically to tackle the fundamental problem of spatial ambiguity in pulmonary lesions.

2. Related work

In the last few years, many researchers have worked on various network architectures and optimization methods to improve the performance of the CNN models in automated classification of CXR images. Of which, attention strategy has attracted a lot of notice for its effectiveness in feature boosting.

Early work was mostly done on using transfer learning, where famous architectures that did good job recognizing natural images were used for diagnosing pulmonary disease such as COVID-19, They gave a way of using deep learning in combination with CXR images, they showed that this could be done for conditions such as pneumonia, pneumothorax and that it can outperform classical ML techniques that are feature-engineered [11-13].

The academics also tried to increase their classification score by building network which is custom made for CXR image and more specifically for detecting the COVID-19. Examples are COVID-Net, proposed by Wang et al. [14] and the modified DarkCovidNet, designed by D. K. Redie [15]. These models have achieved state-of-the-art performance on specific datasets through unique structural designs or special training strategies. But this does not mean that I attempt to build a completely novel network architecture in order to break previous records. On the contrary, my objective here is to show that we can put such a small target module inside any backbone network like ResNet-50 and solve the particular problem of blurry lesions in CXR images. This makes it a general enhancement strategy instead of being a purely task-specific architecture.

3. Materials and methods

3.1. Fundamental structure of the network

In order to solve the problem caused by blurry image edges for classification of CXR images, the SA-ResNet50 model uses a famous model ResNet-50 as the backbone and adds the spatial attention module into the main residual block of the network. Next we will go through the overall architecture of the model, how it works in the spatial attention module, and their integration.

ResNet-50 is composed of an initial convolutional layer (Stem), four stages of residual blocks, and a final classification head. As shown in Figure 1(A), the original CXR image is first fed into the Stem module, which includes a 7×7 convolutional layer, a Batch Normalization (BatchNorm) layer, a ReLU activation function, and a max-pooling layer. Subsequently, the four stages are composed of 3, 4, 6, and 3 bottleneck residual blocks, respectively, each integrated with the SAM. A typical bottleneck residual block, depicted in Figure 1(B), consists of three convolutional layers: a 1×1 convolution for dimensionality reduction, a 3×3 convolution for feature extraction, and another 1×1 convolution to restore the dimensionality.




Figure 1. The architecture of SA-ResNet50

The whole network figure

The proposed SA-bottle-neck block

The detailed structure of the Spatial Attention Module(SA-Module)

The input feature map x not only passes through this series of convolutional layers (the main path) but also directly connects to the output of the main path via a shortcut connection, forming a residual learning unit. Its mathematical expression can be simplified as:

$$y = \mathcal{F}(x, \{W_i\}) + x \quad (1)$$

where x and y are the input and output of the residual block, respectively, and $\mathcal{F}(x, \{W_i\})$ represents the residual mapping function of the main path containing multiple convolutional layers. This structure allows the network to easily learn an identity mapping, thereby preserving the flow of information in deep networks.

3.2. Spatial attention module

To localize lesions, I created a SAM which is Spatial Attention Model. It aims to learn a 2D spatial weight map which would make the model "pay attention" to important regions in the feature map but

"ignore" unimportant background information. In Fig. 1 (C), I show the detailed structure of SAM, which is named SA-Module.

Given an intermediate feature map $F \in \mathbb{R}^{C \times H \times W}$ as input, where C , H , and W represent the number of channels, height, and width, respectively, the computation process of the SAM is as follows (refer to Figure 1(C)).

3.2.1. Channel-wise information aggregation

First, to effectively aggregate channel information at each spatial location, two different types of channel pooling operations are performed in parallel on the input feature map F : Average Pooling, which calculates the average of features across all channels for each spatial position, generating a feature map $F_{\text{avg}} \in \mathbb{R}^{1 \times H \times W}$; and Max Pooling, which extracts the maximum feature value across all channels at each spatial position, generating a feature map $F_{\text{max}} \in \mathbb{R}^{1 \times H \times W}$. These two pooling operations capture global statistical information and the most salient feature responses, respectively.

3.2.2. Feature fusion and attention map generation

Next, these two complementary feature maps, F_{avg} and F_{max} , are concatenated along the channel dimension to form a descriptor of size $\mathbb{R}^{2 \times H \times W}$. This descriptor is then passed through a 7×7 convolutional layer (denoted as $f^{7 \times 7}$) with a single filter, which learns to fuse the two types of pooled information and effectively capture long-range spatial dependencies. Finally, a Sigmoid activation function σ is applied to normalize the output values to the range $[0, 1]$, generating the final 2D spatial attention map $M_s(F) \in \mathbb{R}^{1 \times H \times W}$. This process is precisely described by the following formula:

$$M_s(F) = \sigma(f^{7 \times 7}([\text{AvgPool}(F); \text{MaxPool}(F)])) \quad (2)$$

where $[;]$ denotes the concatenation operation. Each value in $M_s(F)$ represents the importance score of the corresponding spatial location in the original feature map.

3.2.3. Feature refinement

Finally, the generated spatial attention map $M_s(F)$ is multiplied element-wise with the original input feature map F via broadcasting to obtain the spatially refined output feature map F' . This step can be viewed as dynamically re-weighting the original features based on their spatial importance.

$$F' = M_s(F) \otimes F \quad (3)$$

where $\otimes$ denotes element-wise multiplication. The resulting feature map F' has enhanced feature responses in critical regions (such as lesions) and suppressed responses in non-critical background areas.

3.3. Integration into ResNet-50 bottleneck

In order to take advantage of spatial attention and utilize ResNet's strong capabilities, I replaced the regular ResNet-50 bottleneck block with a SA-Bottleneck block. The detailed structure is illustrated by the SA-Bottleneck part of Figure 1(B).

I embedded the Spatial Attention Module (SA-Module) at the end of the main path (the residual branch) of the bottleneck block, specifically after the final 1×1 convolution and its subsequent BatchNorm layer, but before the addition with the output of the shortcut connection. Specifically, after the input feature x passes through the three convolutional layers and corresponding BN and ReLU operations of the main path, a residual feature F_{res} is obtained. At this point, instead of directly

adding F_{res} to x from the shortcut connection, it is first passed through the SA-Module for processing:

$$F'_{\text{res}} = \text{SAM}(F_{\text{res}}) \quad (4)$$

Then, the spatially refined residual feature F'_{res} is added to the output of the shortcut connection, followed by the final ReLU activation function:

$$y = \text{ReLU}(F'_{\text{res}} + x) \quad (5)$$

It is refining this residual feature on the spatial side right before it gets used for the update, without disturbing the original residual learning. This allows the model to see where it is more important to look while learning about the rest. And this method is more straightforward and practical compared with applying it after adding, which will make each update more focused, so it can draw the outline of a blurred lesion better.

4. Experimental structure

4.1. Dataset description

The data set adopted in this article is the COVID-19 radiography data set formed by research groups at Qatar University, the University of Dhaka in Bangladesh, and other institutions. It has pictures of COVID-19, Viral pneumonia, Bacterial Pneumonia and Normal Lungs. There are 3616 COVID-19 images, 1345 Viral Pneumonia images, 2772 Bacterial Pneumonia images and 10192 Normal lung images. In this experiment, I resize all CXR pictures into the shape of 224x224 to fit my model's input size then convert the original gray scale image into a 3-channel RGB images. The processed data set is divided by 8:2 into the training set and the test set.

4.2. Experimental environment and parameter setting

To ensure the fairness and reproducibility of all experimental results, all model training and testing in this study were conducted on a unified hardware and software platform. The computing equipment was equipped with an NVIDIA GeForce RTX 4060 Laptop GPU, and all experiments were performed on the Windows 11 operating system. The primary deep learning framework was PyTorch (2.8.0), used in conjunction with the torchvision (0.23.0) library for data processing and model construction. The underlying computational acceleration was supported by the CUDA (12.1) toolkit. The Adam optimizer was chosen for updating model parameters, with a learning rate set to 1×10^{-4} . For the final residual stage (layer4), which was fine-tuned based on a pre-trained model, a smaller learning rate of 1×10^{-5} was set to avoid disrupting the general features learned during pre-training. The StepLR strategy was used to dynamically adjust the learning rate; every 7 epochs, the learning rate was multiplied by a decay factor of 0.1. This strategy helps the model converge more stably to an optimal solution in the later stages of training. The batch size was set to 32, and the training was conducted for 50 epochs.

4.3. Experimental evaluation indicators

To thoroughly and fairly evaluate how well each model performed for this multi-class problem, especially given that there was an imbalance between classes for medical images, I used an overall score made up of different measures. The calculation of all these metrics relies on the four basic numbers from the Confusion matrix, True Positives (TP), False Positives (FP), True Negatives (TN) and False Negatives (FN) [16].

Accuracy is the ratio of how many correct samples we have to all samples. It is probably the best overall indication of how well the model did.

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad (6)$$

Precision is the ratio of true positive samples to all positive samples [17]. It is how correct the model's predictions were.

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (7)$$

Recall: It is the ratio of all actual positive samples that were correctly found by the model [17]. For medical diagnosis, it is very important for recall to be high because recall is the percentage of actual patients that are detected.

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (8)$$

F1-Score: It is an average of precision and recall [18]. Therefore, it provides a more balanced view of how the model performs when there is a trade-off between precision and recall.

$$\text{F1 - Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (9)$$

Specificity is the ratio of actual negative samples which are predicted to be negative by the model [19]. It measures how well it can recognize healthy samples or the absence of target disease.

$$\text{Specificity} = \frac{\text{TN}}{\text{TN} + \text{FP}} \quad (10)$$

4.4. Experimental design

4.4.1. Ablation study

In order to assess whether the pure spatial attention works well in my proposed SA-Resnet50, I also want to test my hypothesis that spatial attention is more important than channel attention for blurred lesions in CXR images, so I conducted some ablation experiments. In all these experiments, I used the CAM instead of our module, and used the combined CBAM.

Table 1. Performance comparison of different attention modules

CBAM	CAM	SAM	Accuracy	Precision	Recall	F1-Score	Specificity
✓			0.9822	0.9834	0.9789	0.9811	0.9919
	✓		0.9847	0.9794	0.9817	0.9805	0.9934
		✓	0.9847	0.9826	0.9827	0.9826	0.9939

Looking at the values given in Table 1 alongside the baseline ResNet-50 results provided in Section 4.4.2 (accuracy: 0.9777), we can see an improvement by adding any form of attention such as CBAM, CAM, SAM to the model. The above shows that attention methods are fairly good in general at improving the features of CXR images for classification.

Macro-level Accuracy: both SAM and CAM got a score of 0.9847, which is better than CBAM's result of 0.9822. But in terms of the F1-Score, which represents the comprehensive performance of

the model, it's obvious that SAM(0.9826) is higher than CAM(0.9805). Because the F1-Score is the harmonic mean of precision and recall, this advantage means that SAM is more effective at balancing precision and completeness.

Most importantly, our proposed pure spatial attention model(SAM) is also the top performer among four important metrics: Precision, Recall, F1-Score and Specificity. So it does support my hypothesis that for recognizing the CXR images where lesion edges are blurred and spatial localization is critical, a focused pure spatial attention is better than channel attention or simply applying CBAM. Consequently, by forcing the model to learn "where to look", SAM has achieved the most significant improvements in addressing this problem: inaccurate lesion localization.

4.4.2. Contrast experiment

In order to demonstrate the overall effectiveness and progress of the improved SA-ResNet50, I have chosen four classic convolutional neural network milestones, which are considered important landmarks in the development of image recognition. They include AlexNet, VGG16, ResNet-18 and standard ResNet-50. For every model, I used the same method of transfer learning, the same training parameters and tested every model on precisely the exact same test set. A detailed performance comparison can be seen in Table 2.

Table 2. Performance comparison of different baseline models

Models	Accuracy	Precision	Recall	F1-Score	Specificity
VGG16	0.9457	0.9162	0.9250	0.9203	0.9797
AlexNet	0.9602	0.9288	0.9396	0.9333	0.9857
ResNet-18	0.9498	0.9391	0.9290	0.9339	0.9789
ResNet-50	0.9777	0.9734	0.9744	0.9739	0.9910
SA-ResNet50	0.9847	0.9826	0.9827	0.9826	0.9939

As shown in the results, ResNets' (ResNet-18 & ResNet-50) performance is better than that of the previous VGG16 and AlexNet. Specifically, the deeper ResNet-50 (Accuracy: 0.9777) had considerable improvement over the shallower ResNet-18 (Accuracy: 0.9498); thus it can be deduced that having more depth along with powerful features is essential components towards getting a precise CXR image classification.

The proposed SA-ResNet50 got the highest scores across all five runs. Compared to the best baseline model (ResNet-50), SA-ResNet50 improves by 0.0070 (from 0.9777 to 0.9847) in accuracy and by 0.0087 (from 0.9739 to 0.9826) in F1-Score.

These comparisons validate my choice of ResNet-50 as a robust baseline, but much more importantly, they show that putting my proposed spatial attention on top of an already good baseline leads to consistent and broad improvements. This suggests that we do not need to have a deep network; rather, the SA-ResNet50 can have a more robust extraction and utilization of CXR image features with this spatial attention.

4.4.3. Confusion matrix analysis

In order to examine the classification performance of these models in greater detail, I have done some visualization of the confusion matrices for the baseline ResNet-50 and also my proposed SA-ResNet50 on the test set (shown in Figure 2).

The values on the diagonal of the SA-ResNet50 confusion matrix are all greater than or equal to those of the baseline model, and the off-diagonal misclassification values are usually lower. Visually, it is evident that the SA-ResNet50 has more accurate overall classification and less mistaken judgment.

The misclassification between Normal and COVID-19 is prominent, as it's very hard to tell them apart. The baseline model misclassified 36 Normal cases as COVID-19. For the SA-ResNet50 model, this critical number of misclassifications dropped sharply by 44.4% to only 20 cases. This implies that my spatial attention helped the model learn more discriminative information when it came to separating the normal lung texture from COVID-19 lesions. Also, the model got better at telling apart "Bacteria" from "Viral Pneumonia". The baseline model has wrongly labeled the "Viral Pneumonia" as "Bacteria", which is 7 cases, while SA-ResNet50 just has one. This means that SAM is not just working well on majority classes; it is also important for making features better when we might have fewer samples of something.




Figure 2. Comparison of model confusion matrices

4.4.4. Visualization of attention maps

From the interpretable side of things to see how the spatial attention makes changes for the model's decision-making process, I employed the Grad-CAM technique and show what model's class activation heatmaps look like. The areas highlighted in red in the heat maps are the image region that the model focused on most heavily during its decision making for classification. I selected typical examples from each class and compared the differences of the attended region between baseline ResNet-50 and SA-ResNet50 (shown in Figure 3).




Figure 3. Visualization comparison of model attention regions

From the visualizations of Bacteria, COVID and Viral Pneumonia, it is evident that the attention areas the baseline ResNet-50 are quite scattered and wide. The heatmaps generally cover most of the lung fields, often extending into surrounding normal tissue outside the lesion areas. In contrast, the attention focus of the SA-ResNet50 is much better; its highlighted parts are small and concentrated. It could also cover the real lesion area precisely (like pulmonary infiltrates and consolidations in lung). So we can see it is learning from the spatial attention module on how to accurately localize lesions within the image. The heatmaps generated by the SA-ResNet50 are very similar to areas which doctor pays attention while diagnosing. It increases understand ability of the model's decision-making process and also confirms its clinical validity. A model which “looks” at the right places, yields classifications that we can trust much more.

5. Conclusion

The primary issue addressed in this paper is the phenomenon where blurred boundaries of pulmonary lesions in CXR images prevent deep learning models from accurately locating lesions, resulting in decreased classification accuracy. I put forward and proved an improved residual network model based on a pure spatial attention module, namely SA-ResNet50. The key concept in my work is to integrate SAM into a powerful backbone, such as ResNet-50's bottleneck residual blocks and make it help the model pay attention to important parts in images for diagnosis. This paper has successfully proved that the introduction of spatial attention mechanism can improve the accuracy of CXR image classification. Furthermore, it is also a good way to solve other problems about analyzing medical pictures with similar indistinct boundaries and troublesome targets in the same way. As for my future work, I'll try to combine this kind of spatial attention modules into other modern networks, and evaluate its generalizability on larger and more various clinical data sets, so that we can advance its application in clinical computer-aided diagnosis (CAD) systems.

References

- [1] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri and R. M. Summers, Chestxray8: hospital - scaled CXR data base and benchmark for the weakly supervised classification and localization of thoracic diseases in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, US 3462-3471.
- [2] G. N. Hounsfield, "Computerized transverse axial scanning (tomography): Part 1. Description of system," British Journal of Radiology, vol. 46, no. 552, pp. 1016-1022, Dec. 1973.
- [3] P. C. Lauterbur, "Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance," Nature, vol. 242, no. 5394, pp. 190-191, Mar. 1973.
- [4] Y. LeCun et al., "Backpropagation Applied to Handwritten Zip Code Recognition," in Neural Computation, vol. 1, no. 4, pp. 541-551, Dec. 1989.
- [5] A. Krizhevsky, I. Sutskever and G. E. Hinton "ImageNet Classification with Deep Convolutional Neural Networks", in Advances in Neuro-Logical Information Processing Systems (NIPS 2012), pages: 1097 – 1105, 2012.
- [6] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks For Large-scale Image Recognition", In Proc Int'l Conf On Learning Representations (ICLR), 2015:
- [7] K. He, X. Zhang, S. Ren and J. Sun, Deep Residual Learning for Image Recognition, CVPR (CONF., Las Vegas, NV) IN Proceedings DeDD, Proceedings of CONFERENCE ON COMPEUTER RISC AND PATTER RECOGNITION: pp. 760-780
- [8] J. Hu, L. Shen, S. Albanie, G. Sun and E. Wu, "Squeeze-and-excitation networks", In: Proceedings Of The Ieee Conference On Computer Vision And Pattern Recognition (Cvpr), 2018: Pp. 7132-7141.
- [9] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, "spatial transformer network," Adv Nerv infor pro Sys Vol 28 No. nips, (2015), pp. 2017 – 2025.
- [10] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "CBAM: Convolutional Block Attention Module", In Proceedings of the European Conference on Computer Vision (ECCV), 2018, pp. 3-19.
- [11] T. Wang, Y. Zhao, L. Zhu, G. Liu, Z. Ma and J. Zheng, "Lung CT image aided detection of COVID - 19 using alexnet network", 2020 5th international conference on communication, image and signal processing (CCISP) chendu, china, 2020, 199-203.

- [12] K. K. Shibly, S. K. Dey, M. Tahzib-Ul-Islam, and M. M. Rahman, Covid Faster R – Cnn: A Novel Framework To Diagnose Novel Coronavirus Disease (Cov Id – 19) In X - Ray Images, Informatics In Medicine Un locked Vol.20 P.100405 2020.
- [13] K. Tahir, A. Iftikhar and M. Mumraiz, "Forecasting COVID - 19 Via Registration Slips of Patient Using ResNet - 101 And Performance Analysis & Comparison Of Prediction For Covid - 19 Using FASTER RCNN MASK RCNN AND RESNET - 50," 2021 international conference on advances in electrical, computing, communication and sustainable technologies (icaect), bhiLai, India, 2021, p. 1-6
- [14] Wang, L., Lin, Z.Q. Wong, "A." COVI D-Net: A Tailored Deep Convolutional Neural Network Design For Detection of COVID - 19 Cases from Chest X - ray Images, Sci rep, 10, 19549(2020). <https://doi.org/10.1038/s41598-020-76550-z>.
- [15] D. K. Redie "Diagnosis of COVID-19 with Chest X-rays According to the Altered DarkCovidNet Model" Evol. Intell., pp. 1–10, 2022.
- [16] J. T. Townsend, "Theoretical analysis of an alphabetic confusion matrix," Perception & Psychophysics, vol. 9, no. 1, pp. 40-50, 1971.
- [17] A. Kent, M. M. Berry, F. U. Luehrs, Jr., and J. W. Perry, "Machine literature searching VIII. Operational criteria for designing information retrieval systems," American Documentation, vol. 6, no. 2, pp. 93-101, 1955.
- [18] C. J. Van Rijsbergen, Information retrieval: 2nd ed. London, UK: Butterworth-Heinemann; 1979.
- [19] J. Yerushalmy, "Statistical problems in assessing methods of medical diagnosis, with special reference to X-ray techniques," Public Health Reports, vol. 62, no. 40, pp. 1432-1449, 1947.