LEARNING IN MACHINES & BRAINS

DEEP LEARNING IDENTIFIES PNEUMONIA BETTER THAN EXPERTS

Researchers created a computer algorithm that can detect pneumonia in chest X-ray images more accurately than experienced radiologists and enhanced detection over a previous algorithm in 13 other conditions.

KEYWORDS

X-ray, deep learning, pneumonia, neural network, medical diagnosis

FOCUS OF STUDY

This study tested the capabilities of a deep-learning computer algorithm called CheXNet, which the researchers trained to detect 14 different chest diseases by analyzing X-ray images. The goal: to automate the detection of pneumonia and other thoracic pathologies.

BACKGROUND

Neural network algorithms that analyze medical imaging have surpassed the judgments of medical experts at diagnosing several life-threatening conditions, including skin cancer and hemorrhages. New work led by Andrew Ng, associate fellow in CIFAR's Learning in Machines & Brains program, has developed a technology that adds pneumonia and 13 other diseases of the chest to the list.

About 50,000 people die from pneumonia every year in the United States alone. Currently, this lung infection can only be diagnosed by radiologists, who detect the disease by looking at X-ray images of patients' chests. Yet, two thirds of people in the world don't have access to proper diagnosis tools or personnel, according to an estimate by the World Health Organization. It's a barrier that prevents many people from getting the diagnosis they need to treat their disease.

Even when radiologists are available, diagnosing some conditions can be particularly tricky. For example, pneumonia often occurs alongside other pathologies, and can look similar in an X-ray to certain harmless conditions.

This study improves upon previously created neural networks to create an algorithm with the ability to identify thoracic pathologies to a level of accuracy that surpasses that of expert radiologists.

STUDY DESIGN AND METHODS

The researchers created a computer algorithm called CheXNet, a 121-layer convolutional neural network, to analyze the input of frontal chest X-rays from patients.

The output of CheXNet is a positive or negative result for each of these diseases in question, with a probability value for the likelihood that diagnosis was positive.

To train, validate and test their model, they used a recent database of 112,120 frontal-view chest X-ray images from 30,805 patients with 14 different thoracic diseases, including pneumonia.

For each thoracic pathology, the researchers randomly divided the dataset images into three groups: one group of images they used to train the network, one group for validation, and one for testing. Every patient in the data set belonged to just one of these groups, even if they had contributed multiple images.

The CheXNet neural network algorithm is a dense convolutional neural network, where layers within a dense block are connected to every other layer in a feed-forward configuration. These dense connections improved the flow of information through the network. The algorithm also uses batch normalization, a process that has been shown to increase speed and accuracy.

After training CheXNet and validating their results, the researchers tested their network on 420 images

and compared its results to those of four practicing academic radiologists from Stanford University. These radiologists provided their independent analyses of the same images with no prior knowledge of the patients or the data set. The researchers coded a diagnosis by a human or CheXNet as correct for a particular image if it agreed with the majority of the radiologists.

The researchers also modified CheXNet to identify the other 13 diseases, and found that its diagnoses on the X-ray images in the same large data set outperformed existing algorithms.
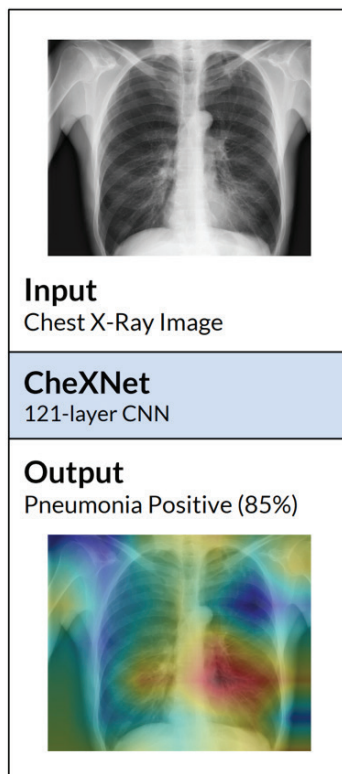
The researchers also visualized CheXNet's results by creating a heatmap on the X-ray images that highlighted the areas most associated with the diagnosis classification. The colour-coded image of each X-ray allows the viewer to see the areas most indicative of the pathology.

KEY FINDINGS

The CheXNet model surpassed the performance of human radiologists in detecting pneumonia, both in sensitivity and specificity. That is, the model's diagnoses were more frequently correct on both positive and negative diagnoses.

When tuned to detect the 13 other thoracic pathologies, CheXNet likewise outperformed previous state-of-the-art algorithms.

FIGURES

Input
Chest X-Ray Image

CheXNet
121-layer CNN

Output
Pneumonia Positive (85%)

CONCLUSION AND IMPLICATIONS

The CheXNet model could help deliver accurate diagnoses quickly and accurately without the need for an expert radiologist. Automating this process would help expedite diagnosis in clinical settings, helping patients get early diagnosis and treatment. For instance, in the case of pneumonia, CheXNet could provide early diagnosis leading to early treatment, which is critical for preventing complications and death.

CheXNet could also help fill the need for expert diagnosis in the many places around the world with inadequate access to radiologists. In this way, it could help deliver healthcare to vulnerable populations.

This study not only shows the potential for an automated diagnosis for pneumonia, but also for 13 other thoracic pathologies, including cardiomegaly, edema, emphysema and hernia.

And the technology may be tweaked to provide even better diagnosis. This study only allowed the radiologists and CheXNet to analyze the X-ray images with a frontal view. However, the algorithm may be modified to also interpret images that present a lateral view of the patient, potentially increasing the accuracy of diagnosis by 15 percent.

REFERENCES

CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning. Andrew Y. Ng et al., arXiv:1711.05225.

RESEARCH BRIEF AUTHORSHIP

Stephanie Orford, BSc (Honours) Behavioural Neuroscience, Simon Fraser University, for CIFAR.

ABOUT CIFAR'S PROGRAM IN LEARNING IN MACHINES & BRAINS

The Learning in Machines & Brains program is revolutionizing the field of artificial intelligence, and creating computers that think more like people – that can recognize faces, understand what is happening in a picture or video, and comprehend the actual meaning of language. The result will be computers that are not only powerful but intelligent, and that will be able to do everything from conduct a casual conversation to extract meaning from massive databases of information.

CIFAR