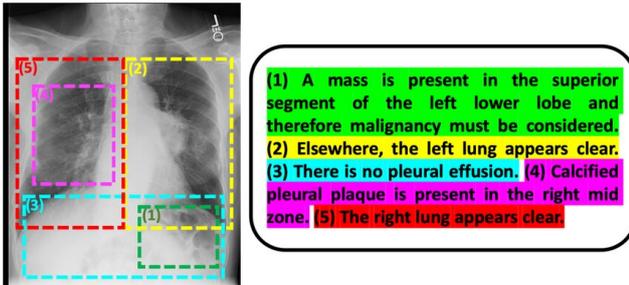# Using AI and old reports to understand new medical images

27 September 2021, by Rachel Gordon

An example image-text pair of a chest radiograph and its associated radiology report. Credit: Massachusetts Institute of Technology

Getting a quick and accurate reading of an X-ray or some other medical images can be vital to a patient's health and might even save a life. Obtaining such an assessment depends on the availability of a skilled radiologist and, consequently, a rapid response is not always possible. For that reason, says Ruizhi "Ray" Liao, a postdoc and a recent PhD graduate at MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL), "we want to train machines that are capable of reproducing what radiologists do every day." Liao is first author of a new paper, written with other researchers at MIT and Boston-area hospitals, that is being presented this fall at MICCAI 2021, an international conference on medical image computing.

Although the idea of utilizing computers to interpret images is not new, the MIT-led group is drawing on an underused resource—the vast body of radiology reports that accompany medical images, written by radiologists in routine clinical practice—to improve the interpretive abilities of machine learning algorithms. The team is also utilizing a concept from information theory called mutual information—a statistical measure of the interdependence of two different variables—in order to boost the effectiveness of their approach.

Here's how it works: First, a neural network is trained to determine the extent of a disease, such as pulmonary edema, by being presented with numerous X-ray images of patients' lungs, along with a doctor's rating of the severity of each case. That information is encapsulated within a collection of numbers. A separate neural network does the same for text, representing its information in a different collection of numbers. A third neural network then integrates the information between images and text in a coordinated way that maximizes the mutual information between the two datasets. "When the mutual information between images and text is high, that means that images are highly predictive of the text and the text is highly predictive of the images," explains MIT Professor Polina Golland, a principal investigator at CSAIL.

Liao, Golland, and their colleagues have introduced another innovation that confers several advantages: Rather than working from entire images and radiology reports, they break the reports down to individual sentences and the portions of those images that the sentences pertain to. Doing things this way, Golland says, "estimates the severity of the disease more accurately than if you view the whole image and whole report. And because the model is examining smaller pieces of data, it can learn more readily and has more samples to train on."

While Liao finds the computer science aspects of this project fascinating, a primary motivation for him is "to develop technology that is clinically meaningful and applicable to the real world."

The model could have very broad applicability, according to Golland. "It could be used for any kind of imagery and associated text—inside or outside the medical realm. This general approach,

moreover, could be applied beyond images and text, which is exciting to think about."

More information: Ruizhi Liao et al, Multimodal Representation Learning via Maximization of Local Mutual Information, arXiv:2103.04537v3 [cs.CV] arxiv.org/abs/2103.04537

Provided by Massachusetts Institute of Technology

APA citation: Using AI and old reports to understand new medical images (2021, September 27) retrieved 28 June 2022 from https://techxplore.com/news/2021-09-ai-medical-images.html