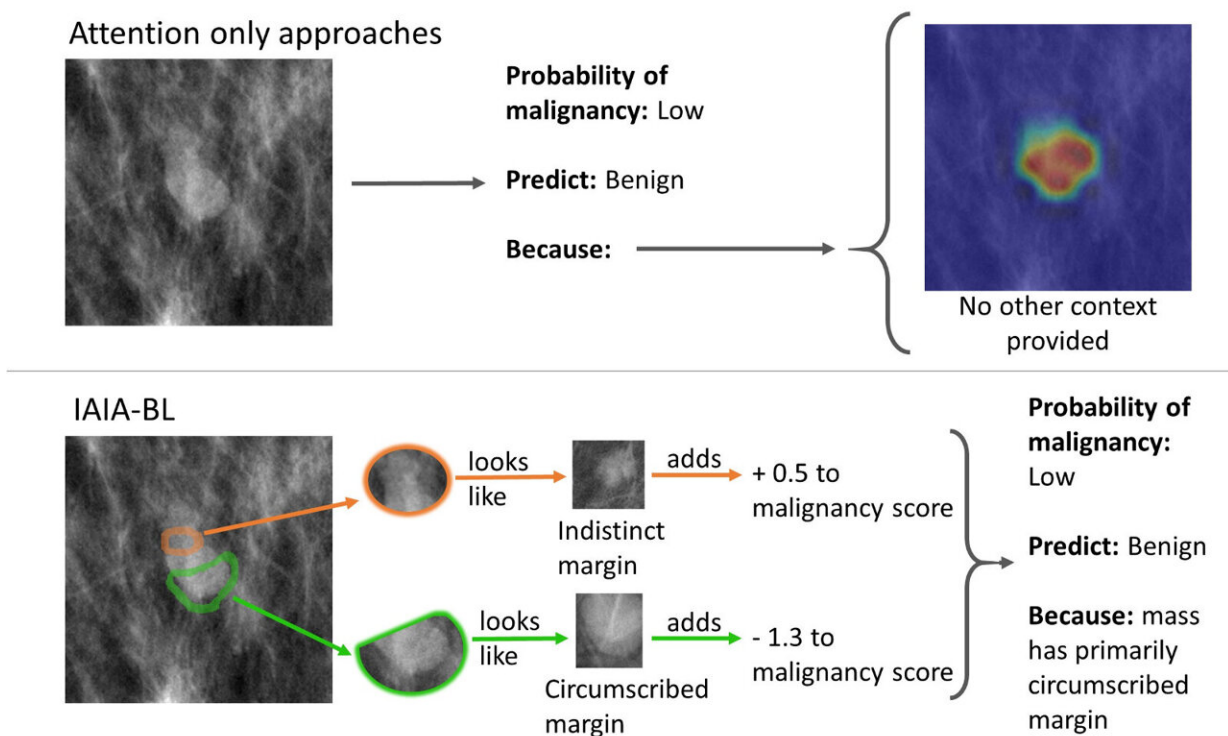


The first AI breast cancer sleuth that shows its work

January 14 2022, by Ken Kingery



Most AI for spotting pre-cancerous lesions in mammography scans don't reveal any of their decision-making process (top). If they do, it's often a saliency map (middle) that only tells doctors where they're looking. A new AI platform (bottom) not only tells doctors where it's looking, but which past experiences it's using to draw its conclusions. Credit: Alina Barnett, Duke University

Computer engineers and radiologists at Duke University have developed

an artificial intelligence platform to analyze potentially cancerous lesions in mammography scans to determine if a patient should receive an invasive biopsy. But unlike its many predecessors, this algorithm is interpretable, meaning it shows physicians exactly how it came to its conclusions.

The researchers trained the AI to locate and evaluate lesions just like an actual radiologist would be trained, rather than allowing it to freely develop its own procedures, giving it several advantages over its "black box" counterparts. It could make for a useful training platform to teach students how to read mammography images. It could also help physicians in sparsely populated regions around the world who do not regularly read mammography scans make better health care decisions.

The results appeared online December 15 in the journal *Nature Machine Intelligence*.

"If a computer is going to help make important medical decisions, physicians need to trust that the AI is basing its conclusions on something that makes sense," said Joseph Lo, professor of radiology at Duke. "We need algorithms that not only work, but explain themselves and show examples of what they're basing their conclusions on. That way, whether a physician agrees with the outcome or not, the AI is helping to make better decisions."

Engineering AI that reads medical images is a huge industry. Thousands of independent algorithms already exist, and the FDA has approved more than 100 of them for clinical use. Whether reading MRI, CT or mammogram scans, however, very few of them use validation datasets with more than 1000 images or contain demographic information. This dearth of information, coupled with the recent failures of several notable examples, has led many physicians to question the use of AI in high-stakes medical decisions.

In one instance, an AI model failed even when researchers trained it with images taken from different facilities using different equipment. Rather than focusing exclusively on the lesions of interest, the AI learned to use subtle differences introduced by the equipment itself to recognize the images coming from the cancer ward and assigning those lesions a higher probability of being cancerous. As one would expect, the AI did not transfer well to other hospitals using different equipment. But because nobody knew what the algorithm was looking at when making decisions, nobody knew it was destined to fail in real-world applications.

"Our idea was to instead build a system to say that this specific part of a potential cancerous lesion looks a lot like this other one that I've seen before," said Alina Barnett, a computer science Ph.D. candidate at Duke and first author of the study. "Without these explicit details, [medical practitioners](#) will lose time and faith in the system if there's no way to understand why it sometimes makes mistakes."

Cynthia Rudin, professor of electrical and computer engineering and computer science at Duke, compares the new AI platform's process to that of a real-estate appraiser. In the black box models that dominate the field, an appraiser would provide a price for a home without any explanation at all. In a model that includes what is known as a 'saliency map,' the appraiser might point out that a home's roof and backyard were key factors in its pricing decision, but it would not provide any details beyond that.

"Our method would say that you have a unique copper roof and a backyard pool that are similar to these other houses in your neighborhood, which made their prices increase by this amount," Rudin said. "This is what transparency in medical imaging AI could look like and what those in the medical field should be demanding for any radiology challenge."

The researchers trained the new AI with 1,136 images taken from 484 patients at Duke University Health System.

They first taught the AI to find the suspicious lesions in question and ignore all of the healthy tissue and other irrelevant data. Then they hired radiologists to carefully label the images to teach the AI to focus on the edges of the lesions, where the potential tumors meet healthy surrounding tissue, and compare those edges to edges in images with known cancerous and benign outcomes.

Radiating lines or fuzzy edges, known medically as mass margins, are the best predictor of cancerous breast tumors and the first thing that radiologists look for. This is because cancerous cells replicate and expand so fast that not all of a developing tumor's edges are easy to see in mammograms.

"This is a unique way to train an AI how to look at medical imagery," Barnett said. "Other AIs are not trying to imitate radiologists; they're coming up with their own methods for answering the question that are often not helpful or, in some cases, depend on flawed reasoning processes."

After training was complete, the researches put the AI to the test. While it did not outperform human radiologists, it did just as well as other black box computer models. When the new AI is wrong, people working with it will be able to recognize that it is wrong and why it made the mistake.

Moving forward, the team is working to add other physical characteristics for the AI to consider when making its decisions, such as a lesion's shape, which is a second feature radiologists learn to look at. Rudin and Lo also recently received a Duke MEDx High-Risk High-Impact Award to continue developing the algorithm and conduct a

radiologist reader study to see if it helps clinical performance and/or confidence.

"There was a lot of excitement when researchers first started applying AI to [medical images](#), that maybe the computer will be able to see something or figure something out that people couldn't," said Fides Schwartz, research fellow at Duke Radiology. "In some rare instances that might be the case, but it's probably not the case in a majority of scenarios. So we are better off making sure we as humans understand what information the computer has used to base its decisions on."

More information: Alina Jade Barnett et al, A case-based interpretable deep learning model for classification of mass lesions in digital mammography, *Nature Machine Intelligence* (2021). [DOI: 10.1038/s42256-021-00423-x](https://doi.org/10.1038/s42256-021-00423-x)

Provided by Duke University

Citation: The first AI breast cancer sleuth that shows its work (2022, January 14) retrieved 20 March 2024 from <https://techxplore.com/news/2022-01-ai-breast-cancer-sleuth.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.
