May 2020

Smarter, More Accurate, and Less Expensive

How Artificial Intelligence is Revolutionizing Cancer Diagnostics

Written by: Ethan Gauvin

Table of Contents

I. Introduction

From image analysis to data management to biomarker discovery, artificial intelligence ("AI") is increasingly being used by researchers and healthcare professionals to improve patient outcomes. Nowhere is this truer than in the field of cancer diagnostics, where advances in AI have led to the development of tests that can catch cancer earlier, with greater accuracy, and for less money than existing solutions. Not only will these AI-enabled diagnostic tools prevent countless deaths, but their widespread adoption will save healthcare systems trillions of dollars and transform the field of preventative medicine.

II. Why Now?

The concept of AI—computers capable of performing tasks that normally require human intelligence—has been around since the late 1950s. Moreover, the medical benefits of early-diagnosis have been known for centuries. So why are we just now seeing groundbreaking advancements in AI-enabled diagnostics? The answer can be primarily attributed to two emerging trends: an explosion in clinical data and the advent of machine learning.

1. The Rise of Big Data and Machine Learning

The past two decades have witnessed an extraordinary growth in the volume, variety, and velocity of data across all industries, but especially in healthcare. According to a report by the International Data Corporation, the volume of healthcare data —such as lab reports, clinical notes, wearable device data, and radiology scans—is expected to grow faster than any other sector between 2018 and 2025. [1] Doctors have gone from not having enough data on a patient to having far too much data to comprehend.

Fortunately, recent advances in AI have now made it possible to quickly make sense of large quantities of data. A subcategory of AI known as machine learning ("ML") is particularly well-suited for big-data analysis, and is responsible for the bulk of today's AI research, investment, and product development. Whereas traditional computing (and early forms of AI) followed a "top-down" approach requiring computer systems to adhere to a rigid set of instructions (if x input, then y output), ML takes a "bottom-up" approach allowing software to train itself on large amounts of data through trial-and-error—not unlike the development of the human brain. [2]

---

[1] Reinsel, D., Gantz, J., & Ryding, J. "The Digitization of the World." International Data Corporation, Nov. 2018, https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf
[2] Vigliarolo, Brandon. "Machine Learning: A Cheat Sheet." TechRepublic, 9 Apr. 2018, techrepublic.com/article/machine-learning-the-smart-persons-guide/

By the time IBM's Deep Blue supercomputer—which mostly relied on traditional computing—bested world chess champion Gary Kasporav in the late 1990s, researchers were still convinced top-down approaches were superior to ML. Yet they were also increasingly aware of the limitations of traditional methods of computing. The reason Deep Blue won was because it was programmed to understand the rules of a well-defined game and counter every potential move an opponent could make, however, this "brute force" approach required an enormous amount of computing power and could not be realistically applied to games or problems that had millions or billions of potential outcomes. Top-down approaches faced a scaling problem: it was neither feasible nor practical to program a list of responses to every conceivable event or input. [3]

2. How Machine Learning Works

Since then, computing power has grown exponentially while becoming cheaper and more accessible. The average smartphone in 2020 is millions of times more powerful than the most advanced computer in the 1990s. [4] These developments, combined with an explosion in data, have not only made ML useful, but positioned it as the dominant method of AI being studied and deployed around the world.

---

[3] Greenmeier, Larry. "20 Years after Deep Blue: How AI Has Advanced Since Conquering Chess." Scientific American, 2 June 2017, https://www.scientificamerican.com/article/20-years-after-deep-blue
[4] Roser, Max, and Hannah Ritchie. "Technological Progress." Our World in Data, 2020, ourworldindata.org/technological-progress#exponential-progress-moore-s-law

Rather than code a response to countless potential outcomes, today researchers can feed data (e.g. images, soundbites, text, or games of chess) through an ML algorithm and teach it to make accurate decisions over time through millions of rounds of trial-and-error. The more data used to train an ML network, the more accurate it becomes.

For instance, if a researcher wanted to train an ML algorithm to distinguish dogs from cats, she would run it on thousands of (accurately labeled) dog and cat images. For each image, the algorithm makes a guess as to whether the creature in the photo is a dog or cat. After guessing, the algorithm is provided with the correct answer, and then records whether its guess was right or wrong. After repeating this thousands or millions of times, the algorithm will have "learned" how to identify a dog versus a cat, and will be able to do so with a high degree of accuracy.

While this process sounds straightforward, it can be successfully applied to extremely complex problems. In 2016, Google subsidiary DeepMind was able to develop an ML algorithm that beat the world champion of the notoriously complex game of Go, shocking the global AI research community.

Because there are more potential Go configurations than atoms in the known universe, it is impossible to program software to counter every conceivable move. Instead, Deepmind researchers trained their program on tens of millions of games, and over time it developed Go strategies that the best human players could not match, much less imagine. [5]

ML is particularly well-suited for medical diagnoses for the same reasons it became useful in the first place: an abundance of clinical data and cheap computing power has allowed researchers to train ML to instantly recognize patterns and pathologies that humans are incapable of detecting quickly (or at all). One can imagine how, after being trained on hundreds of thousands of x-ray images, an ML algorithm could become a quick and effective diagnostic aid. This is exactly what we are starting to see today, and ML experts believe this is only the beginning.

[5] "AlphaGo: The Story so Far." Deepmind, 2020, deepmind.com/research/case-studies/alphago-the-story-so-far.

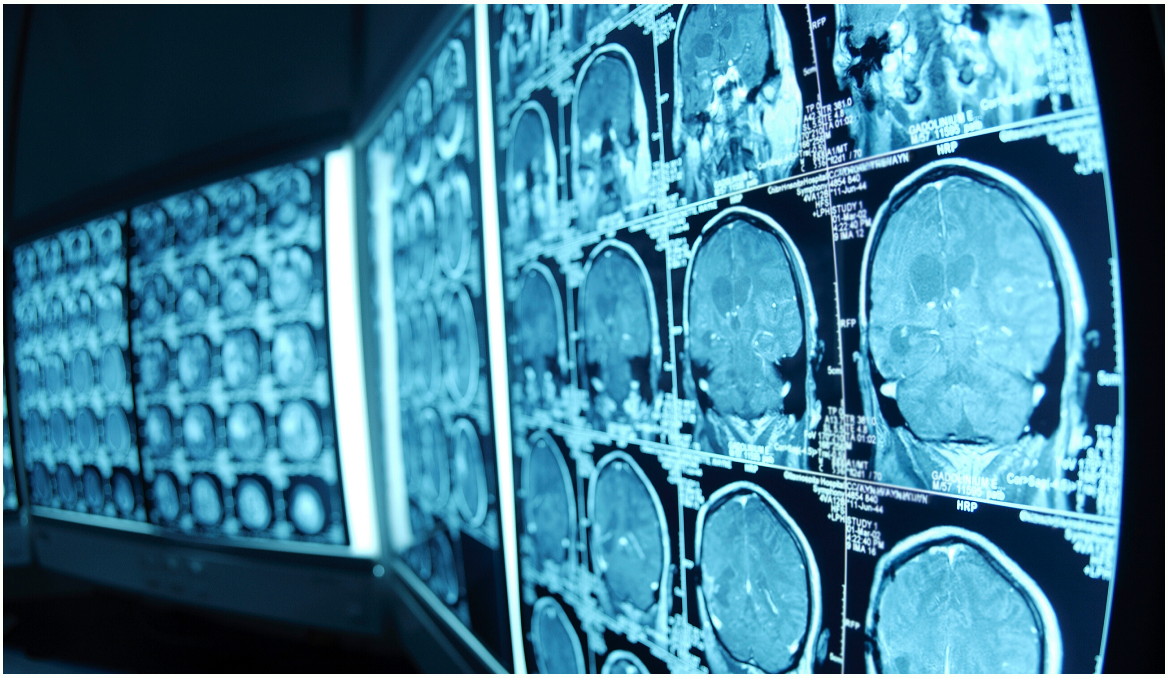

III. Improving Cancer Diagnosis

While the global cancer death rate has declined 15% since 1990, cancer continues to kill around ten million people per year and is the second leading cause of death behind cardiovascular diseases. [6] Today, most cancer treatments are expensive, invasive, non-targeted, and too late to be effective (after cancer has already spread), which is why medical professionals believe early detection and diagnosis remains the single best way to combat the disease. A rapidly growing number of companies, research labs, and medical institutions are developing ML techniques to catch cancer faster and with greater accuracy.

[6] Roser, Max, and Hannah Ritchie. “Cancer.” Our World in Data, 2020, ourworldindata.org/cancer#cancer.

1. Improving X-ray and CT Scan Diagnosis

The majority of cancers are diagnosed through medical imaging, such as x-rays and CT scans. The challenge with this approach is that it depends on human interpretation and does not always provide a reliable diagnosis—even for the most seasoned radiologists and oncologists. Healthcare professionals routinely see a false negative diagnosis rate of 20-30% (false positives are also common). [7] Moreover, the growing number and availability of diagnostic images is rapidly exceeding the capacity of specialists—especially in developing countries.

Even in the earliest stages of its development, ML proved it could rival the best human experts at accurately diagnosing a variety of diseases. A meta-analysis of 82 studies published between 2012 and 2018 compared the diagnostic accuracy of human professionals against ML, and found that on average ML models equaled their human counterparts in both sensitivity and specificity (sensitivity measures the proportion of positive cases accurately identified and specificity measures the proportion of negative cases accurately identified). [8]

---

[7] Bradley, Stephen, et al. "Sensitivity of Chest X-Ray for Lung Cancer: Systematic Review." British Journal of General Practice, 1 June 2018, bjgp.org/content/68/suppl_1/bjgp18X696905.
[8] Liu, Xiaoxuan, et al. "A Comparison of Deep Learning Performance against Health-Care Professionals in Detecting Diseases from Medical Imaging: a Systematic Review and Meta-Analysis." The Lancet, 1 Oct. 2019, thelancet.com/journals/landig/article/PIIS2589-7500(19)30123-2/fulltext.

Today, ML algorithms are beginning to outperform their human counterparts. A recent study conducted by an international team of researchers found that ML was superior to a group of 58 dermatologists at diagnosing skin cancer, catching 95% of melanomas compared to the human specialists' 86%. [9] In 2020, Google researchers and the National Cancer Institute collaborated to see if ML could predict lung cancer by analyzing 43,000 low-dose CT scans of the lungs from 15,000 different people. For patients with only one scan available, Google's model outperformed every radiologist who also examined the scans to assess risk of lung cancer. Google's ML algorithm also reduced false negatives by 5% and false positives by 11%. [10]

Of course, the future of diagnostics will require both the rigorous analysis of ML and careful interpretation by human specialists. A study in 2019 found that, when reading hundreds of breast cancer screening images, radiologists' sensitivity improved by 8% and specificity by 6% when they had an ML tool to assist them. Moreover, the time required to read and interpret the findings fell from a minute to less than thirty seconds with the help of ML. [11] Not only is ML improving the accuracy of diagnostics, but it is also improving the efficiency of human specialists, which will ultimately result in a greater number of diseases diagnosed and lives saved.

---

[9] "Man Against Machine: Artificial Intelligence Is Better than Dermatologists at Diagnosing Skin Cancer." ESMO, 29 May 2018, esmo.org/newsroom/press-office/artificial-intelligence-skin-cancer-diagnosis.
[10] Savage, Neil. "How AI Is Improving Cancer Diagnostics." Nature News, Nature Publishing Group, 25 Mar. 2020, nature.com/articles/d41586-020-00847-2.
[11] Sokol, Emily. "Deep Learning, AI Improve Accuracy of Breast Cancer Detection." HealthITAnalytics, 1 Aug. 2019, healthitanalytics.com/news/deep-learning-ai-improve-accuracy-of-breast-cancer-detection.

2. Improving Tissue Sample and Biomarker Diagnosis

ML can be applied to more than just x-rays and CT scans. For example, ML has proven proficient at diagnosing images of tissue samples, as demonstrated earlier this year by researchers at the University of Waterloo. By training their algorithm on 30,000 previous biopsies of confirmed cancer cases, the researchers' algorithm was able to achieve near-100% sensitivity in diagnosing new biopsy images of 32 forms of cancer in 25 organs and body parts. [12]

However, perhaps the most exciting application of ML to cancer diagnosis is related to data optimization and genomics. For example, the startup 20/20 GeneSystems has developed a single blood test that can detect the presence of several different types of cancer in otherwise healthy people by measuring tumor protein biomarkers. When relying on biomarker measurements alone, the test yielded a sensitivity of 57%. But when combined with an ML algorithm trained on a database of 28,000 patients, the test's sensitivity increased to 82%. [13] Johns Hopkins researchers have also produced successful results by using a similar approach combining ML and molecular testing technology. [14]

---

[12] Kalra, Shivam, et al. "Pan-Cancer Diagnostic Consensus through Searching Archival Histopathology Images Using Artificial Intelligence." Nature , 10 Mar. 2020, www.nature.com/articles/s41746-020-0238-2.
[13] "About OneTest." One Test for Cancer, 2020, onetestforcancer.com/learn.
[14] Mone, Amy. "Early Cancer Detection Technology Receives Record Venture Investment." Johns Hopkins University , 3 June 2019, hub.jhu.edu/2019/06/03/cancerseek-blood-test/.

IV. Economic and Social Benefits

Consider a healthy, middle-aged woman in the year 2030. She visits her primary care doctor for a routine check-up and within a couple hours is informed that a protein biomarker indicating the presence of breast cancer was detected in her saliva. While the news is unsettling, she is reassured by her doctor that tests nowadays are capable of catching cancer at the earliest stage, and that she should come in for additional screening. Her mammogram is processed by an algorithm and, although invisible to the human eye, the computer identifies a group of cells that will one day form a tumor if left untreated. The cancerous cells are confirmed by a quick biopsy, removed, and she is left with instructions to conduct a monthly breast cancer screening, which can now be done in the comfort of her own home by taking a swab of saliva and mailing it to a lab.

While it may sound like science-fiction, these types of diagnostic solutions are actually in development today. A future where routine preventative care includes non-invasive ML-enabled tests for multiple types of cancer—just as patients are routinely tested for early stages of heart disease today—could become a reality in this decade.

1. Tangible Savings

Because of their inexpensive and non-invasive nature, AI-enabled diagnostic tests will greatly expand access to cancer screening. This will result in a dramatic increase in the number of people screened each year, especially within developing countries and low-income demographics, as well as the number of cancer cases diagnosed and treated early.

The economic and social benefits of these developments will be transformational. In 2019, the cost of cancer treatment in the U.S. reached $150 billion annually, with 1.8 million new diagnoses and 600,000 deaths. In addition, it is estimated that cancer-related mortality and productivity loss cost the American economy an additional $180 billion, with the overall economic burden of cancer totaling $330 billion or 1.8% of gross domestic product. [15]

---

[15] "The Economic Burden of Cancer." American Cancer Society, 2020, cancer.org/taking-action/economic-burden/.

While it is difficult to estimate what the savings might look like once AI-diagnostic technologies reach full maturity, studies conducted by the World Health Organization have shown that treating patients with early-diagnosed cancer is consistently two to four times less expensive than treating those with advanced-stage cancer. [16] By applying the same methodology, and assuming the majority of America's population has access to routine cancer screening, annual spending on cancer care in the U.S. could one day be reduced by anywhere between 67-80%. [17] We can expect these savings to be similar in other high-income countries with advanced healthcare systems, which translates to trillions of dollars saved globally.

2. Intangible Rewards

Crucially, these estimates cannot account for the intangible benefits facilitated by widespread cancer screening and early diagnosis. This includes, at minimum, a massive reduction in the mental health toll on families and loved ones, a decreased burden on hospital systems, and a redirection of funds to deserving healthcare needs that would otherwise be spent on late-stage cancer treatment.

---

[16] "Early Cancer Diagnosis Saves Lives, Cuts Treatment Costs." World Health Organization, 3 Feb. 2017, who.int/news-room/detail/03-02-2017-early-cancer-diagnosis-saves-lives-cuts-treatment-costs.

[17] If the cost of treating late-stage cancer is 2-4x the cost of treating early-stage cancer, then by eliminating late-stage cancer cases you would theoretically reduce total cancer spending by 67-80%. Example 1: 100B total cancer spending, of which 33.3B is early treatment and late treatment is 2x the cost (66.6B). Eliminating late treatment leaves 33.3B in early treatment. 100B to 33.3B is a decrease of 67%. Example 2: 100B total cancer spending, of which 20B is early treatment and late treatment is 4x the cost (80B). Eliminating late treatment leaves 20B in early treatment. 100B to 20B is a decrease of 80%.

Today, the five year survival rates of common types of cancer—such as breast, prostate, testicular, and thyroid cancer—are all close to a remarkable 99% when caught at the earliest stage. [18] It is not a stretch to say that, when early-diagnosis becomes the norm for the majority of cancers in the near future, millions of lives will be saved every year. While economists have attempted to assign a monetary value to an individual human life (the U.S. government puts it at $10 million), the cumulative total of future lives saved is unarguably priceless. [19]

---

[18] Kandola, Aaron. "Top 7 Most Curable Cancers Based on 5-Year Relative Survival Rate." Medical News Today, 7 Aug. 2018, medicalnewstoday.com/articles/322700.

[19] Thomson-DeVeaux, Amelia. "What Should The Government Spend To Save A Life?" FiveThirtyEight, 27 Mar. 2020, fivethirtyeight.com/features/what-should-the-government-spend-to-save-a-life/.

V. Conclusion

AI and ML have only just started to revolutionize the field of cancer diagnostics. In the years ahead, they will radically democratize cancer screening, facilitate widespread early diagnosis, and save millions of lives. The subsequent economic savings and social benefits are difficult to comprehend but extraordinary in their potential.

It is important to note that these technologies will face many challenges and setbacks throughout their development and implementation. No discussion of AI and healthcare is complete without acknowledging important privacy, data governance, and ethical concerns, which carry with them intense regulatory scrutiny and political obstacles—not to mention entrenched private interests that will oppose any attempt to make healthcare more cost-effective. These concerns are outside the scope of this paper, but there are many excellent resources that cover them in detail. [20]

Lastly, while AI and ML are exciting tools with incredible potential, they must be recognized for what they are: tools. No algorithm is capable of the careful judgement, communication, creativity, and other uniquely-human literacies that medical professionals utilize on a daily basis. The purpose of these technologies is to augment physicians' ability to solve complex problems, not to replace their jobs or decision-making.

---

[20] Rigby, Michael J. "Ethical Dimensions of Using Artificial Intelligence in Health Care." American Medical Association, 1 Feb. 2019, journalofethics.ama-assn.org/article/ethical-dimensions-using-artificial-intelligence.

While not oblivious to the risks, this report hopes to establish a recognition that, given the trajectory of AI-enabled diagnostic technology, we are on the cusp of stopping one of the world's greatest killers. Although important to implement these technologies in a safe and ethical manner, it is categorically unethical to tolerate the shocking levels of death and suffering caused by cancer in the 21st century—especially when AI has given us the means to make progress.

* * *

About the author: Ethan Gauvin is dual MBA/MPA at the MIT Sloan of Management and the Harvard Kennedy School of Government, and is currently an Academic Fellow at DigitalDx Ventures, an early-stage venture capital firm focused on digital diagnostics. Before graduate school, he spent five years on Google's Legal and Government Affairs teams where he helped shape the company's position on a range of technology and AI policy issues.

Acknowledgements: I am indebted to Michele Colucci (Managing Partner, DigitalDX Ventures) and Michael Solotke (MD/MBA candidate at Yale) for their excellent advice and feedback. This paper was written with support from DigitalDX Ventures, however, all views are my own.