

I, Dentist: Is artificial intelligence the future of oral healthcare?

by

Robert M. Davis

Submitted to the
Program in Comparative Media Studies/Writing
in Partial Fulfillment of the Requirements for the Degree of

Master of Science in Science Writing

at the

Massachusetts Institute of Technology

September 2021

© 2021 Robert M. Davis
All rights reserved

The author hereby grants to MIT permission to reproduce and to distribute
publicly paper and electronic copies of this thesis document in whole
or in part in any medium now known or hereafter created.

Signature of Author.....
Program in Comparative Media Studies/Writing
May 28, 2021

Certified by.....
Ashley Smart
Associate Director, Knight Science Journalism Program at MIT
Thesis Advisor

Accepted by.....
Seth Mnookin
Director, Graduate Program in Science Writing

I, Dentist

Is artificial intelligence the future of oral healthcare?

by

Robert M. Davis

Submitted to the Program in Comparative Media Studies/Writing
on May 28, 2021 in Partial Fulfillment of the Requirements
for the Degree of Master of Science in Science Writing

ABSTRACT

The field of dentistry has earned a reputation for being more prone to misdiagnosis and overtreatment than other medical subspecialties. This is driven, at least in part, by a professional culture that has traditionally been less scientific and evidence-based than that of general medicine. It is also partially driven by an array of economic pressures that have predisposed dentists towards more aggressive and expensive treatment options, as well as by the legitimate ambiguities of clinical decision making in oral healthcare.

In the past decade, at least half a dozen dental AI companies have begun selling software that they claim can help mitigate the problem of misdiagnosis and overtreatment in dentistry. Some of their systems attempt to do this by monitoring insurance claims, and flagging suspicious patterns in patient records. Other systems focus on automating the diagnosis process itself—scanning patient X-rays to identify simple issues like a cavity, for example, or a tooth abscess.

On paper, dental AI companies pitch their products as helpful tools that can assist dentists by automating busywork and providing a backstop against innocent human error. They suggest that computer vision technology can help ensure that every dentist has access to the latest best practices and evidence-based care recommendations. However, AI technology can do much more than merely play the role of digital assistant. In many ways, these systems also serve as a kind of hall monitor, providing tacit enforcement of the clinical best practices that inform their programming.

This project explores a simple question: What do we gain, and what do we lose, by bringing artificial intelligence to oral healthcare?

Thesis Advisor: Ashley Smart

Title: Associate Director, Knight Science Journalism Program at MIT

Skyler King was once the proud owner of 32 teeth—all of them planted firmly in his mouth. That changed, however, in the spring of 2009, after King visited his then-dentist, Dr. Mark D. Meyers of St. Louis, Missouri.

In a lawsuit, King claims he arrived at his dentist's office on May 7 to seek treatment for a dental abscess, a kind of pus-filled tooth pimple. It is not a pleasant condition. A tooth abscess can cause intense pain that radiates into the jawbone and the neck, or as far as the ear.

In most cases, the condition is relatively easy to treat. But an abscessed tooth can be deadly. Without proper care, the bacterial infection that causes the abscess can spread to nearly any part of the head and neck—including the brain. It can also spread to the bloodstream, where the infection may send the body into sepsis, a dangerous immune response that can lead to organ failure and death.

It is likely that even an inexperienced dentist could have diagnosed the problem without much difficulty, but King says Meyers was lightning quick. After performing a cursory oral examination, no X-rays required, Dr. Meyers declared that King would need to have all of his teeth, every single one of them, removed immediately. A failure to act, Meyers warned, would mean risking potentially fatal blood poisoning. King was just 23 at the time.

As King recalls it, Meyers instructed his patient to come back to his office the following day. Meyers confirmed with King's mother that she would be willing to put the full \$5,235.00 bill on her credit card. He charged the card.

Then, one by one, Dr. Meyers proceeded to wrench all 32 of Skyler King's teeth out of his head.

According to the lawsuit that he would eventually file against Dr. Meyers, King later learned that it was only the one abscessed tooth and three of its neighbors—four teeth in total—that needed to be pulled. His 28 other teeth were perfectly healthy, or near enough. (At the time, there was a second, much more treatable abscess on his lower right first molar.) King says he only came to learn these details after consulting with another dentist as part of his lawsuit.

In a deposition, Meyers claimed that it was King himself who had asked to have all of his teeth removed.

“What 23-year-old would want all of their teeth pulled when they know they have teeth in good condition?,” said King's attorney, John S. Wallach, during a 2014 interview with [Courthouse News Service](#). “Who would do that?”

It is worth noting here that Dr. Mark D. Meyers is not just a general dentist who runs a standard private practice where he administers routine cleanings and the occasional root canal. While Meyers has practiced general dentistry since 1979, he is best known as a specialist in prosthodontics. This means that he provides dentures for patients in need of full or partial replacement of missing teeth. Meyers is the proprietor of [Eat Right Dentures](#), where King sought treatment for his abscess, and where Meyers still practices dentistry to this day.

Alleged misdiagnosis, costly overtreatment, and the extreme discrepancies between Meyers' assessment and that of another dentist examining the same case—none of these is unique to

the story of Skyler King. Each is emblematic of a larger overarching problem in the field of dentistry, a phenomenon that [one 2016 research paper](#) not-so-pithily refers to as the problem of “professional variability in decision making in modern dentistry.”

Put another way, dentists have a consistency problem. Given the same patient, even the most well-trained, well-meaning clinicians will vary wildly from one to the next in their approaches to diagnosis and treatment planning. Their perceptions are indelibly influenced by the training they’ve received, the culture in which they practice, the sharpness of their senses. The methods they choose are shaped by their skill with particular procedures, by economic pressures, by the sleep they didn’t get last night.

It is, in many ways, a very human problem. And according to a growing crop of dental AI startups, it may be a problem that calls for an inhuman solution.

In the past decade, at least half a dozen dental AI companies have begun selling software they claim can help mitigate the problem of misdiagnosis in dentistry. Some of their systems attempt to accomplish this by monitoring insurance claims, and flagging suspicious patterns in patient records. Others automate the diagnosis process itself—scanning patient X-rays to identify simple issues like a cavity, for example, or a tooth abscess.

These companies say their products give dentists a second pair of eyes, so to speak, eyes that will help them better manage their practices, and avoid making flawed or inconsistent readings. For larger dental support organizations (DSOs) that manage many dentists across multiple locations, dental AI companies offers software that can analyze electronic health records to uncover insights about the health of their patients as a whole, and the performance of their clinicians.

“We can rank the dentists: How good are they prepping? ...Are they doing the correct treatment based on what is found in the [X-rays],” said Dr. Kyle Stanley, co-founder and chief clinical officer of the dental AI company Pearl, in [an appearance](#) with on the video series, “Group Dentistry Now.” “The DSO managers can then make the decision. Who needs training? Who needs a raise? Who needs to be fired?”

Speaking more recently over Zoom, Stanley suggested that the problem of inconsistency in diagnoses has led to dentists earning some “bad PR.” Indeed, the story of Skyler King provides a striking example. Stanley and his fellow co-founders at Pearl believe artificial intelligence can solve, or at least diminish, that problem.

“You can't expect an AI to be perfect,” he said. “What you *can* expect the AI to do is be more consistent—not get tired, not get hungry, always give you the same response.”

In [September 2020](#), Pearl joined forces with over a dozen leaders from major dental care providers, insurance companies, and academia to form the [Dental AI Council \(DAIC\)](#). Pearl, notably, remains the only dental AI company associated with the DAIC. According to its website, the non-profit aims to identify areas where AI will be most valuable to the field of dentistry, and

more broadly to “answer fundamental questions related to AI's role in the oral healthcare ecosystem of tomorrow.”

A few months after its founding, in late December, the DAIC published its first research report, a study showing what the organization calls “glaring inconsistency” in the diagnosis and treatment of dental pathology. The DAIC study took the form of a survey, in which 136 licensed dentists across 14 countries were asked to review a set of X-rays made from an unidentified patient, along with lists of 11 diagnoses and 29 possible treatments. Respondents were asked to identify any issues they noticed, as well as the treatments they would typically prescribe. The results, according to the DAIC, were bleak.

The DAIC found no instances whatsoever of unanimity in either diagnosis or treatment planning across the entirety of the survey. The highest level of agreement came from the identification of non-metallic fillings (81% of survey respondents identified them correctly), impacted molars (65%), and recurrent tooth decay (63%). In the overwhelming majority of cases, however, fewer than 50% of correspondents agreed upon the same diagnosis.

Variability in diagnosis leads to variability of treatment and, by extension, variability of treatment cost. Proposed treatment plans submitted by the respondents were quite varied, ranging in at just a few hundred dollars on the low end, and at as much as \$36,000 at the highest.

The DAIC report says there are two reasons for the extreme discrepancies observed both in the context of this study and anecdotally in the world of dental diagnosis and treatment planning writ large. One has to do with the proliferation of treatment options for a given diagnosis. Over the past 30 years or so, dentistry has seen rapid advancement in terms of the techniques, technologies and materials used to treat common problems, meaning that dentists have multiple options—perhaps too many options, the report suggests—for a given ailment.

The other has to do with a fundamental problem of uncertainty. The afflictions dentists identify in X-rays are often ambiguous. In most cases, they exist on some sort of spectrum in terms of category and level of severity. Microcavities, for example, are the very earliest indicator of tooth decay, and in some cases will grow into a more serious issue. However, they don't always require intervention. One patient may be perfectly fine without intervention; another may find that their teeth quickly deteriorate.

Companies like Pearl argue that artificial intelligence systems can help bring much needed consistency and standardization of care to the field of dentistry. AI-powered computer vision systems are already more than capable of scanning patient X-rays and identifying common pathologies. Since AI systems rely on repeatable, algorithmic evaluation to arrive at their diagnoses, they are much less prone to the variability of decision making found among human dentists.

However, there are some who take issue with the DAIC study's conclusions—and its methods, for that matter.

“If you submitted that to me, let's say in a journal, and I was a reviewer, I would have rejected the article—or at least told you to throw certain things out,” said Dr. Bernard Friedland, associate professor of oral medicine at the Harvard School of Dental Medicine.

For example, the DAIC study sent surveys to dentists in 14 countries, including wealthy nations like the U.S. and U.K., and developing nations like India and Peru. Friedland said that it doesn't make sense to group wealthy countries with developing nations where dentists have access to fewer resources. He observed that dentists must design their treatment plans based on the resources at their disposal, even if that means making sub-optimal clinical decisions. "They're influenced by what's available," Friedland said.

One could also quarrel with the DAIC's interpretation of the study's results. The study suggests that all instances of misdiagnosis are unintentional—caused by an overabundance of treatment options or by truly ambiguous clinical "edge cases" that any dentist would struggle to identify. However, in the real world, there is a great deal of evidence to suggest that misdiagnosis in dentistry is often intentional, with dentists increasing their profit margins by prescribing unnecessary treatments for nonexistent conditions. Profit motive might not mean much to dentists participating in an unpaid survey, but as Friedland points out, the practice of overtreatment can be so ingrained that it becomes habitual, almost reflexive.

"As far as I'm concerned, there is very much overtreatment in the U.S.," said Friedland, noting that some dentists genuinely believe such treatments are in the best interest of their patients, while others are more driven by profit. "It's a little hard to tease out to what extent the diagnosis is influenced by what the dentist is thinking," he said.

To illustrate his point, Friedland brought up two common treatments for cavities. Smaller cavities are often treated with inexpensive fillings that literally "fill in" the cavity. Bigger cavities call for crowns—a form of dental prosthetic that covers and protects the damaged tooth. Crowns, of course, are much more expensive than a low-end filling. The problem is that there's no firm, universal rule for deciding which cavities are "big" and which ones are "small." In many cases, the decision between a filling and a crown comes down to a dentist's purely subjective clinical judgment, and with all else being equal, Friedland says that many dentists will naturally tend towards the more expensive option.

"It's kind of like unconscious bias," he said. "They're not actually bad people. But I think in the back of their mind, [they're thinking] 'I could make a little bit more money doing a crown,' even if they're not consciously thinking of that."

Officially speaking, Pearl was born in May of 2019, when the company stepped onto the public stage for the first time to announce [its \\$11 million Series A funding round](#). However, Pearl actually began as a subunit of another, much older AI company—the somewhat ironically named [GumGum](#). While most other dental AI companies were founded by PhD researchers looking to put their conceptual work to the test, Pearl wasn't created as some academic experiment. Ultimately, it was born to seize what its founders view as a considerable market opportunity.

The idea for Pearl's dentistry-oriented solutions originated with its co-founder and CEO, the Carnegie Mellon-educated engineer Ophir Tanz. Tanz, a natural-born entrepreneur, had grown up learning a great deal about dentistry almost by osmosis thanks to his father, who had a decades-long career as a practicing clinician.

“The thing with dentistry is that it’s always somehow the forgotten medicine,” Tanz said in a [TechCrunch interview](#) shortly after the company’s debut. “But it’s such a massive market opportunity.”

Pearl’s product offerings fall into three broad categories organized by customer type. “We do a lot of work with dental laboratories,” Tanz said. “We do a lot of work with providers and DSOs. We do a lot of work with insurance carriers. And it’s really oriented around dentistry... It’s identifying untreated opportunities. It’s adjudicating insurance claims. It’s identifying fraud, waste, and abuse, all the stuff that is commercializable today, because we need to build a functional business.”

For independent dental practices and DSOs (Dental Support Organizations, which are essentially large private equity-backed dental chains), Pearl offers a “Second Opinion” diagnostic platform, which can identify common issues by scanning patient X-rays. (Pearl is currently seeking FDA approval to license the software in the U.S.) They also offer a product called Practice Intelligence, which can help dentists and DSO managers make more data-driven administrative decisions by unearthing insights from patient health records.

For dental insurance providers, Pearl provides the Claims Review platform, which can quickly inspect and approve the thousands of straightforward insurance claims that make up the bulk of a human claim adjudicator’s work, and which in many cases are simply too numerous for any human team to process. A complementary service called Pearl Protect does the opposite—sniffing out suspicious claims and flagging them for human scrutiny. Pearl also offers a pair of products for laboratories that create custom dental restorations.

At their core, each of the solutions that Pearl offers is powered by a branch of artificial intelligence called computer vision. “Computer vision is essentially teaching a machine to see things—to perceive visual information,” said Cambron Carter, Pearl’s chief technology officer, and the company’s third (and final) co-founder. Every service Pearl provides, from reviewing insurance claims to generating practice-wide administrative insights, is based on the automated analysis of patient X-rays. That analysis is enabled by computer vision.

In an interview, Pearl CTO Cambron Carter explained that computer vision, like its human counterpart, can be broken down into two core components. First, there’s the simple act of “sensing,” or taking in visual information. Then, there’s the perceptual component, which allows a digital or biological system to derive meaning from that information. “It’s more about computer visual perception,” Carter said.

To achieve visual perception, Pearl’s computer vision systems mostly rely on a kind of algorithm—or really a collection of algorithms—called a convolutional neural network. Neural networks are software programs that are used in automated data analysis and machine learning.

As the name implies, neural networks are inspired by the human brain. Picture a human dentist who spots a cavity while examining a patient’s X-ray. What’s really happening there is that the dentist has taken in visual information about the X-ray image through a powerful sensing instrument known as “the eyes,” and then passed that information to the network of neurons that

make up the brain's visual cortex, where it processes that information until it generates an output. Here, the output takes the form of the dentist's diagnosis: The patient has a cavity.

The convolutional neural networks (CNNs) used in image processing are fundamentally the same. They take information, or "input values," pass that information through "layers" of networks of artificial neurons—which mimic the function of the brain's visual cortex—and generate some kind of output value. For image processing, that usually means spitting out a label: "cat," "dog," "cavity," "tooth abscess."

"The task that we do [falls] into three subsets," said CTO Cambron Carter. "We classify images. We do what's called detection...and then also segmentation."

Image classification is the act of looking at the full digital image and declaring "this is a dental X-ray," or "this is a picture of two hamsters in a cage." Detection, or more specifically "object detection," is an image classification technique which involves placing a rectangular box around the key objects depicted in the image—a molar, a lesion, maybe a microcavity—and identifying what those objects are. Image segmentation allows the CNN to get a more granular understanding of exactly where objects in the image are located, since the boxes used in object detection aren't very precise.

But how does a CNN know whether the object it's detecting is a cat or a tooth abscess? How does it know which lines, shapes, colors, and other features belong to which kinds of objects? It all comes down to training. To teach the CNN how to recognize different objects in a dental X-ray, Dr. Stanley and a team of fellow clinicians have spent years annotating thousands upon thousands of X-ray images, many of them sourced through partnerships with dentists and other oral healthcare providers. The clinical team carefully labels all the objects they want the CNN to be able to find on its own.

From there, Carter and his technical team feed those labeled images to the CNN, and the CNN begins to teach itself how to label objects on its own. It examines the images over and over again until it is able to consistently identify key objects without relying on annotations provided by humans. Once the CNN is able to do that, it's ready to start classifying new, unlabeled images—or, well, almost.

The clinical team at Pearl also provides the system with detailed clinical guidelines, or "heuristics," shortcuts that allow developers to steer the system in the right direction without needing to figure out everything for itself. Dr. Brian Howe put the concept in more concrete terms. "If [the CNN] knows that the cavity has to be on the tooth, [and] if that box identifies it on the gums, well, it's not a cavity," he said. "You have to have the rules, obviously, for it to say where these things occur, what they are, etcetera," he said. There are also heuristics for approving insurance claims, which must be individually tailored to fit the requirements of different insurers.

If CNNs are the secret sauce that enable AI platforms like Second Opinion to perform medical diagnosis, then it's heuristics—the hand-coded rules placed atop those platforms—that enable them to meet the demands of a vast and complex healthcare system. A neural network can teach itself many things when it comes to dental diagnosis, but it cannot deduce the intricacies of privacy laws, insurance regulations, or even human biology merely by analyzing dental X-rays.

In this way, platforms like Second Opinion become a kind of avatar, perhaps even an advocate, for the rules and collected knowledge of the healthcare establishment. Where dentists are often quite insulated from that establishment, these systems bring it into the heart of their day-to-day operations.

For decades, the overwhelming majority of dentists—including Dr. Mark D. Meyers of St. Louis, Missouri—have worked in private practice, where they generally have free rein to use antiquated techniques they may have picked up in school decades ago, or to charge patients and insurers for unnecessary treatments. One might assume that dentists commonly employ the same scientific rigor and evidence-based practices as any medical professional, or that they are subject to meaningful oversight from regulatory bodies and professional organizations, but that isn't always the case.

In 1997, *Reader's Digest* published [a bombshell article](#) in which reporter William Ecenbarger conducted an informal experiment, undergoing 50 dental exams by 50 different dentists all around the U.S. "Dentistry is a stunningly inexact science," Ecenbarger wrote in his piece. "Even expecting that different dentists would have different, yet valid, opinions did not prepare me for the astounding variation in diagnoses I received."

Ecenbarger's story was less focused on the very real ambiguities that dentists must navigate when it comes to diagnosis and treatment, and much more concerned with warning Americans about the perils of intentional overtreatment and fraud. His account overflows with examples of unscrupulous dentists urging him to spend thousands of dollars on procedures that his own dentist and a team of impartial experts agreed were unnecessary.

Dr. Michael Wahl, a practicing dentist in Wilmington, Delaware, recalls that Ecenbarger's article sent shockwaves through the dental community when it first arrived in print, and raised legitimate concerns. "I think he brought up some good points," Wahl said. "I don't know that it's quite as bad as he portrayed in his article, but certainly, there is a lot of difference of opinion [among] dentists when it comes to treatment planning."

Ecenbarger's article may have caused a stir, but it doesn't seem to have changed the way that dentists do business. More than 20 years later, in 2019, *The Atlantic* published a similar story, "[The Truth about Dentistry](#)," this time focusing on a single dental practice that was found to have overtreated and overcharged patients for years in a fraudulent scheme. While fraud was a major theme in the piece, writer Ferris Jabr says this fraud is partially enabled by a general lack of scientific rigor within the profession.

"Common dental procedures are not always as safe, effective, or durable as we are meant to believe," Jabr wrote. "As a profession, dentistry has not yet applied the same level of self-scrutiny as medicine, or embraced as sweeping an emphasis on scientific evidence."

Jabr found that while the concept of "evidence-based medicine" has been around since the 1960s, the idea of "evidence-based dentistry" has only been circulating since the 1990s. He suggests that this is why inconsistencies in diagnosis in treatment planning seem to be more common in dentistry than in general medicine.

Of course, where Jabr sees a problem, Dental AI companies see an opportunity. “There should really only be one diagnosis, no matter what the problem is,” said Pearl’s Dr. Kyle Stanley.

There is very little publicly-available data that shows just how big the problem of misdiagnosis and overtreatment in dentistry really is. [One 2019 study](#) found that roughly 25% of U.S. healthcare spending is lost to wasteful spending, including on unnecessary and erroneous treatments, but the study did not provide estimates for dentistry specifically. Many insurers estimate that fraud and waste account for roughly 10% of spending, but that number is based on a more than 20-year-old [California Medicaid study](#), and also fails to provide numbers specific to dentistry.

State dental boards often play a role in tracking complaints lodged against dentists under their jurisdiction, but only some publish that data, and even fewer provide delineate it by category. Where such figures are available, a surface-level review suggests that experiences like Skyler King’s are rare. For example, in 2019, only 15 of the 472 complaints sent to the Ohio State Dental Board fell under “Excessive/Unnecessary Treatment.”

Still, there is some anecdotal evidence which suggests that misdiagnosis and overtreatment in dentistry may be more prevalent than they appear. In 2020, *USA Today* reporter David Heath published a pair of investigative stories in which he found that large private equity-backed dental chains were pressuring dentists [“to drill healthy teeth for profit.”](#) The examples he found were far from unique. Indeed, many large national dental-care chains have been the subject of lawsuits in recent years for overtreatment and other deceptive practices.

The company DentalWorks, for example, [was sued in 2013](#) for allegedly pressuring the dentists in its franchises to make inaccurate diagnoses and push unnecessary treatments. In 2016, insurers for another dental chain called Small Smiles paid out [\\$43 million](#) to settle claims of false billing related to unnecessary procedures conducted on children—the chain itself had gone bankrupt years earlier after the FBI began investigating it for Medicaid fraud. In 2018, the dental chain Kool Smiles paid \$24 million [to settle claims of false billing](#) for unnecessary procedures conducted on children.

Skyler King’s mother may have paid for King’s unnecessary tooth extractions out of pocket, but in many cases, it’s insurance companies that bear the brunt of these. As a result, many insurance companies have enthusiastically embraced AI technology for sniffing out fraudulent claims.

Like any healthcare provider, dentists who wish to receive payment from a patient’s insurance carrier must first submit a claim, a formal request that provides proof that the dentist performed some procedure on the patient’s behalf. In dentistry, such claims usually include billing details, records of the services provided, and X-rays or other diagnostic imaging documenting those services.

Before the rise of AI, claims were reviewed by teams of human adjudicators, clinical professionals who have the expertise needed to determine whether the services provided were necessary and adequately performed. That process may seem relatively straightforward, but for

most insurance companies, it's an enormous challenge. Insurance carriers generally receive far more claims each day than human adjudicators can review, making them a ripe target for fraud.

"It's a pretty easily gamed system," said Pearl CTO Cambron Carter. "First of all they're looking at subpar information—potentially distorted, noisy X-rays—and they're only looking at 10% of all the X-rays...Based on what we hear from carriers, they just have a constant mountain of work, and they're trying to adjudicate as quickly as possible."

Carter said that malicious dentists use numerous techniques to take advantage of that system, editing X-rays with photoshop, submitting the same claim to multiple insurance providers, and more. "It's very wild, to be honest with you...to actually see it happening in practice."

Fortunately for insurance carriers, the vast majority of claims are fairly straightforward, and don't require close human inspection. This means that AI and automation are a great match for insurance adjudication. The AI system sorts through the stacks of honest claims that insurers don't have time to review by hand, and allows human adjudicators to focus their efforts on the much smaller proportion of suspicious, possibly fraudulent claims. As a result, fraudulent claims are much less likely to slip through unnoticed.

"With some of the claims that we had looked at previously, and then looked at again using the AI tool, we find things that otherwise just kind of went through the system," said Dan Williams, director of innovation at CareQuest Institute for Oral Health. In 2020, Williams' oversaw a pilot program assessing the use of AI for insurance claims adjudication. The results, he says, were impressive.

"What was surprising was that there was actual doctoring of X-rays—stuff you might not see with a normal naked eye," Williams said. "The AI tool can pick up variations of density of ink, or density of the impression that it's looking at, and say 'well, something here doesn't quite match with the rest of the pixelation of that picture from the X-ray itself.'"

However, despite the revelations of the past year, Williams—like most dental-industry professionals interviewed for this piece—maintained that the overwhelming majority of dentists are well-intentioned. He also said that fraudulent claims only make up a small portion of the claims that come in, though he said his team has a long way to go before they know how big the problem really is.

"It's too early, I think, to say 'yes, it's 10% or 15%.' I don't think we know yet," Williams said. "The indicators are that it's higher than we thought."

Pearl is far from the only company operating in the burgeoning dental AI space. Others offering similar services include [Overjet](#), a Boston-based startup founded by researchers from MIT and Harvard, [VideaHealth](#), another Boston-based startup launched around the same time, and Diagnocat, a startup based in Moscow. These companies are among the few who offer solutions for both healthcare providers and insurance carriers. Some in the industry, like [Denti.AI](#), only provide products for dentists. Others, like the San Francisco-based [Dr. Opinion](#), only provide services for insurers.

Notably, nearly every dental AI company that is in business today was founded in the year 2017 or later. This makes intuitive sense given the evolution of CNNs and computer vision technology.

In 2011, image recognition technology was still fairly immature, with the most sophisticated systems boasting an error rate of around 25% when attempting to identify even simple objects, like a dog or a strawberry. This essentially means that their best guesses for identifying everyday objects depicted in publicly-available image data sets were incorrect 25% of the time. For context, humans—who will also make mistakes on image classification tasks when they're tired, or when an object could fit into multiple categories—have an error rate of about 5% for similar tasks.

In 2012, a CNN called AlexNet won an annual international visual recognition competition, identifying images of everyday objects with an error rate of just 16%. In the years that followed, the techniques used to create AlexNet—along with the proliferation of large, high-quality, open-source image data sets—helped other researchers make rapid advancements in the field. By 2017, more than three-quarters of all entrants in the annual visual recognition competition had achieved error rates of less than 5%. Today, the average image recognition system boasts an error rate of roughly 3%.

Modern AI systems don't just perform better than humans when it comes to labeling large sets of random images, they're also better than humans at medical diagnosis—at least when making diagnoses based on diagnostic imaging like X-rays. [A 2019 meta-analysis](#) found that, in general, AI consistently outperforms human clinicians in detecting diseases based on imaging across all illness types. In dentistry, studies have found that AI outperforms human clinicians at identifying [bone loss](#), [oral lesions](#), and [dental implants](#)—among other things.

These differences are often quite moderate. [One study](#) found that a diagnostic system for identifying cavities on patient X-rays had an average accuracy rate of about 80%, as opposed to 71% for dentists. That may not be an enormous improvement, but it's not nothing. And while the technology behind computer vision platforms is [improving all the time](#), the same cannot be said for human eyesight, or our ability to process visual information.

Based on this data, AI companies argue that automated diagnostic platforms can bring tremendous value to oral healthcare providers, and to insurers as well. While there isn't much available data showing how well AI systems perform against their human counterparts in the context of insurance claims adjudication, insurers themselves are very much convinced that it will offer significant advantages. The industry has been [investing heavily in AI technology](#) for years, both in and out of healthcare.

Still, experts warn that AI has many limitations, and broadly agree that humans still have an important role to play in both clinical diagnoses and insurance claim review.

“When I'm looking at an X-ray, [it's] one piece of the puzzle to help me make a decision. [It's] a 2D film of a 3D structure,” said University of Iowa family dentistry professor Dr. Brian Howe. “I also need to rely on my clinical exam, to determine what I think.”

As Howe explains, today's AI diagnostic platforms make diagnosis based only on the analysis of patient X-rays, but human clinicians must consider many other sources of information.

“We're using information based off of what we collected from multiple sources,” Howe said. “We're looking at X-rays. We're looking at the exam...[There's] diagnostic testing that we may do on different teeth, sometimes advanced imaging if we need it...What does the patient want? One of the questions we ask is what's the patient's chief concern?...We're balancing those things.”

Dentists also raised concerns about the effect that algorithmic bias could have as AI begins to play a bigger role in adjudicating insurance claims. Computer vision systems, like all machine learning systems, learn to make decisions by analyzing data. In most cases, that data is created by humans, and when it comes to healthcare, it also *describes* humans—meaning that it contains all of our human prejudices, limitations, and mistakes. Like impressionable children, AI systems inevitably reproduce, and even amplify, those failings.

[One infamous 2019 study](#) found that millions of black people had been impacted by racial bias embedded in a widely-used decision-making algorithm that was designed to identify at-risk hospital patients most in need of advanced care. Trained on decades of hospital records, the algorithm began giving black patients lower risk scores than white patients with the same conditions.

This wasn't because the black patients were less likely to grow seriously ill, but simply because the algorithm was making decisions based on historical records that showed black patients consistently receiving less care than their white counterparts. It seems the developers behind the algorithm failed to provide it with any heuristics for the systemic racism, generational poverty, and history of exploitation and mistreatment in the medical system that the black community has traditionally faced in the U.S.

Algorithmic bias in automated claims review may manifest in a similar way, or it may arrive in a form that's much more subtle, and which companies will have less incentive to eliminate. “When we look at AI, we have to take a look at how AI is being implemented,” said Dr. Chris Smiley, a practicing dentist and editor-in-chief of the *Journal of the Michigan Dental Association*. “What's it being implemented for? What are the objectives?”

Dr. Smiley, whose surname may very well have destined him for the profession, observed that insurers may program their automated systems with cost-cutting measures that aren't always in the patient's best interest. “If AI is simply geared towards serving specific metrics that a benefit plan has, and those metrics are for cost containment, those metrics are maybe not geared towards having the service that has the best outcome,” he said.

Smiley suggested that such a system might, for example, refer patients to in-network dentists who offer the lowest rates, regardless of quality of service. “You're looking at bias,” he said, “anything that's made by man has bias.”

Fortunately, the conversation around algorithmic bias has grown significantly in recent years, and most AI companies are on high alert for instances of discriminatory bias in their software. This is especially important for companies in the healthcare sector, whose patient-facing products require regulatory approval.

“I deal with this daily for our FDA pursuits,” said Cambron Carter. “They want [you] to show that you are not biased with respect to gender, you’re not biased with respect to age, geography, ethnicity.”

If algorithmic bias can be dealt with, and if AI can be integrated into medicine without diminishing other modes of clinical examination and decision-making, many clinicians say that the technology has enormous potential. “This could be one of those tools that I use, and it could certainly make my life a lot easier,” said University of Iowa’s Dr. Brian Howe.

Today, systems like Second Opinion are mostly focused on diagnosis, but Howe envisions AI potentially offering a helpful shortcut in treatment planning as well. “[Maybe the AI] spits out a treatment plan to say, ‘here are some options that you should consider for these teeth.’ That would be awesome,” he said.

Current computer vision diagnostic systems are also mostly limited to identifying what’s already there, but Pearl CTO Cambron Carter looks forward to a day when they can do more. Carter suggested that, by analyzing a compendium of patient X-rays from childhood to present day, a system like Second Opinion might be able to model what the patient’s X-ray will look like in the future—enabling better preventative care.

“It’s a little *Minority Report*,” said Carter, referencing the 2002 science fiction thriller about a society in which crimes are punished before they happen. “There’s a whole host of ethical [issues], because there’s nothing wrong yet. But really leaning into the forecasting power of neural networks, I think, is something that stands to benefit medicine.”

As AI continues to make its way into dentistry, and into the broader healthcare system, clinicians and technologists agree that the most important thing is keeping human judgment “in the loop”—so that automated systems don’t take over patient care entirely. “I think that AI in dentistry is a tool that can be used...but it’s one tool to be used,” said Howe, “just like I use the multiple parts of my exam, to make a decision on diagnosis and treatment for my patients.”

Technologists like Cambron Carter echo this line of thinking, albeit for different reasons. “There’s this fear out there about replacement,” said Carter, using the example of a system like Second Opinion replacing the human radiologists who interpret X-ray imaging devices. “Most of the medical devices I see being introduced into the market are meant to assist you, to be an aid. Not to say that they couldn’t outperform radiologists, but I think that sort of [assistive] philosophy with the go-to-market strategy is very helpful.”

Carter stressed that AI healthcare tools are meant to be used by healthcare providers, meaning that companies like Pearl must support those providers if they want to sell their products. “Really your target users are doctors,” Carter said. “You don’t want to scare them away by saying, ‘Hey, I’m going to take your job in five years,’ which is certainly not the case—absolutely not the case—and I think there are tons of unforeseen jobs that [AI in healthcare] will yield.”

Still, history provides many reasons to doubt just how sincere healthcare AI companies are being when they say they don’t want to automate clinicians out of the job. From factory workers to switchboard operators, the past century is littered with examples of occupations that have

been lost to automation. Now, with millions of Americans out of work due to the COVID-19 pandemic, some fear that automation trends [will only accelerate](#).

“Even if they claim they always want the clinician in the loop, capitalism is gonna make that not happen,” said Dr. Robin Zebrowski. Zebrowski is an associate professor who researches AI at Beloit College. “If I’m a clinician and... [my competitor] is staying in the loop and only getting through 10 patients a day, [but] I’m just letting this algorithm work and I’m getting the money from 40 patients in a day, the 40 patients are going to be the ones we go with.”

Another consideration is the impact that AI will have on dentists’ clinical decision making. On paper, dental AI companies like Pearl pitch their products as helpful tools that can support dentists by automating busywork and providing a backstop against innocent human error. They suggest that computer vision technology can help ensure that every dentist has access to the latest best practices and evidence-based care recommendations.

However, AI technology has the potential to do much more in dentistry than merely play the role of digital assistant. In many ways, AI systems can also serve as a kind of hall monitor, providing tacit enforcement of the clinical best practices that inform their programming. In some respects, this may be a good thing for patients.

“What we’re talking about is like a baseline,” said Dr. Wardah Inam, CEO of the dental AI company Overjet. “Whether you are going to the best dentist in Boston, or you’re going to a rural village in the middle of nowhere, you should be able to get at least the basic standard of care.”

Perhaps Skyler King would still have his teeth today if Eat Right Dentures had employed some kind of AI-powered diagnostics platform. Such a platform might have scanned King’s X-rays, indicated that the majority of his teeth were quite healthy, and dissuaded Dr. Meyers from moving forward with his extreme treatment plan. If it is the case that Dr. Meyers chose to extract Skyler King’s teeth for less-than-virtuous reasons, perhaps having a record of an automated diagnosis would have discouraged any attempted fraud, or at least made such actions more difficult to defend in court.

At the same time, AI systems have the potential to exert less positive influence. Depending on the insurance heuristics their systems are outfitted with, the presence of AI could pressure dentists into selecting cheaper, less effective treatments. It could also pressure dentists to disregard the personal preferences of their patients—especially if those preferences somehow conflict with standard clinical best practices, or with an insurance company’s particular policies.

Adviser, witness, custodian, enforcer—AI may go on to play any combination of these roles. Unencumbered by outdated knowledge, unmoved by human greed, the technology could bring much-needed consistency and oversight to the surprisingly messy world of dentistry. The question is whether it will be a benevolent force, or a kind of tyrant, bullying human clinicians into accepting its automated diagnoses. As artificial intelligence continues its entrée into the world of dentistry, and into the field of medicine more generally, only time will tell whether it arrives at the right balance.

Bibliography

Interviews:

- Barrow, Jane R. (Associate Dean for Global and Community Health, Harvard School of Dental Medicine, MA). Audio interview: February 10, 2021.
- Blatz, Markus (Assistant Dean for Digital Innovation and Professional Development, University of Pennsylvania School of Dental Medicine, PA). Audio interview: February 5, 2021.
- Carter, Cambron (Chief technical officer & co-founder, Pearl, CA). Video interview: January 27, 2021.
- Carter, Cambron (Chief technical officer & co-founder, Pearl, CA). Video interview: April 23, 2021.
- Claffey, Noel (Former professor, Trinity College Dublin, Dublin). Audio interview: January 28, 2021.
- Fine, Noah (Post-doctoral fellow, University of Toronto, Toronto). Video interview: March 1, 2021.
- Friedland, Bernard (Associate Professor of Oral Medicine, Infection, and Immunity, Harvard School of Dental Medicine, MA). Video interview: April 20, 2021.
- Gibson, Barry (Professor in medical sociology, University of Sheffield School of Clinical Dentistry, Sheffield). Video interview: February 11, 2021.
- Grover, Jane (Director of the Council on Access Prevention and Interprofessional Relations, American Dental Association, IL). Audio interview: February 18, 2021.
- Howe, Brian (Assistant professor and director of clinics, University of Iowa College of Dentistry and Dental Clinics, IA). Video interview: May 7, 2021.
- Inam, Wardah (CEO, Overjet, MA). Video interview: February 19, 2021.
- Inam, Wardah (CEO, Overjet, MA). Video interview: April 16, 2021.
- Messina, Matthew (Assistant professor and clinical director, Ohio State University College of Dentistry). Video interview: February 8, 2021.
- Richens, Jonathan (Senior research scientist, Babylon Health, London). Video interview: March 1, 2021.
- Skrocki, Marilyn (Professor of Health Sciences, Saginaw Valley State University, MI). Audio interview: February 4, 2021.

Smiley, Chris (Owner and general dentist, Smiley Family Dentistry, MI). Audio interview: April 29, 2021.

Stanley, Kyle (Chief clinical officer & co-founder, Pearl, CA). Video interview: January 20, 2021.

Stanley, Miguel (Clinical Director, White Clinic, Lisbon). Video interview: February 18, 2021.

Tanz, Ophir (CEO & co-founder, Pearl, CA). Video interview: January 14, 2021.

Thompson, Murray (Professor of Dental Epidemiology and Public Health, University of Otago, Dunedin). Video interview: February 17, 2021.

Wahl, Michael J. (Owner, Wahl Family Dentistry, DE). Video interview: January 29.

Williams, Dan (Executive director of innovation, CareQuest Institute for Oral Health, MA). Video interview: April 16, 2021.

Zebrowski, Robin (Associate Professor of Cognitive Science, Beloit College, WI). Video interview: January 29, 2021.

Journal articles, online sources, government documents, and books:

Beck, Julie. "Why Dentistry Is Separate From Medicine." *The Atlantic*, Atlantic Media Company, 9 Mar. 2017, www.theatlantic.com/health/archive/2017/03/why-dentistry-is-separated-from-medicine/518979/.

Bindushree, V., et al. "Artificial Intelligence: In Modern Dentistry." *Journal of Dental Research & Review*, vol. 7, no. 1, 2020, pp. 27–21., doi:10.4103/jdrr.jdrr_2_20.

Burnett, Brigette. "Settlement for Kids Taken Advantage of by Small Smiles Dental Clinics." <https://www.13abc.com>, WGTV Toledo Ohio, 7 Oct. 2016, www.13abc.com/content/news/Settlement-for-kids-taken-advantage-of-by-Small-Smiles-Dental-Clinics-396257021.html.

Catani, Marco. "A Little Man of Some Importance." *Brain*, vol. 140, no. 11, 2017, pp. 3055–3061., doi:10.1093/brain/awx270.

Chan, Chun Hei Michael. "Classical Neural Network: What Really Are Nodes and Layers?" *Medium*, Towards Data Science, 29 Mar. 2020, towardsdatascience.com/classical-neural-network-what-really-are-nodes-and-layers-ec51c6122e09.

Chen, Po-Hsuan Cameron, et al. "How to Develop Machine Learning Models for Healthcare." *Nature Materials*, vol. 18, no. 5, 2019, pp. 410–414., doi:10.1038/s41563-019-0345-0.

Childress, Sarah. "DentalWorks Chain Misdiagnosed for Money, Dentists Say." *PBS*, Public Broadcasting Service, 13 Mar. 2013, www.pbs.org/wgbh/frontline/article/dentalworks-chain-misdiagnosed-for-money-dentists-say/.

Christopoulos, Apostolos. "Mouth Anatomy." *Overview, Gross Anatomy: Oral Vestibule, Gross Anatomy: Oral Cavity Proper*, Medscape, emedicine.medscape.com/article/1899122-overview#a2.

Claffey, Noel M., et al. "Chapter 4: History of the Oral-Systemic Relationship." *Periodontal Disease and Overall Health: A Clinician's Guide*, Professional Audience Communications, Inc., 2010, pp. 42–54.

GumGum, gumgum.com.

Demush, Rostyslav. "A Brief History of Computer Vision (and Convolutional Neural Networks)." *HackerNoon*, 26 Feb. 2019, hackernoon.com/a-brief-history-of-computer-vision-and-convolutional-neural-networks-8fe8aacc79f3.

Dental AI Council, www.dentalaicouncil.org/.

"Dental AI Council Launches to Help Define the Future of AI in Dentistry." *PR Newswire*, 29 Sept. 2019, www.prnewswire.com/news-releases/dental-ai-council-launches-to-help-define-the-future-of-ai-in-dentistry-301140014.html.

Denti.AI, www.denti.ai/.

Deo, Priya Nimish, and Revati Deshmukh. "Oral Microbiome: Unveiling the Fundamentals." *Journal of Oral and Maxillofacial Pathology : JOMFP*, vol. 23, no. 1, 2019, pp. 122–128., doi:10.4103/jomfp.JOMFP_304_18.

DeStefano, Amelia Williamson. "Exclusive: Dental AI Council Launched by Industry Leaders from Delta Dental, Heartland, Schein, Walmart, Others." *Dentistry IQ*, 28 Sept. 2020, www.dentistryiq.com/dentistry/article/14184293/exclusive-dental-ai-council-launched-by-industry-leaders-from-delta-dental-heartland-schein-walmart-others.

Dewhirst, Floyd E., et al. "The Human Oral Microbiome." *Journal of Bacteriology*, vol. 192, no. 19, 2010, pp. 5002–5017., doi:10.1128/jb.00542-10.

Dickson, Ben. "3 Ways AI Is Transforming the Insurance Industry." *TNW*, 27 Apr. 2021, thenextweb.com/news/3-ways-ai-is-transforming-the-insurance-industry.

Donoff, R. Bruce. "It's Time to Break down the Wall between Dentistry and Medicine." *STAT*, STAT, 17 July 2017, www.statnews.com/2017/07/17/dentistry-medicine-division/.

Dr. Opinion - Dental Insurance Fraud Detection, dropinion.ai/.

"Eat Right Dentures: Same Day Dentures." *Eat Right Dentures*, Dr. Mark D. Meyers, www.eatrightdentures.com/index.html.

Ecenbarger, William. "I Went to 50 Different Dentists and Almost All of Them Gave Me a Different Diagnosis." *Reader's Digest*, Reader's Digest, Feb. 1997, www.rd.com/article/how-honest-are-dentists/.

"Epithelium." *Encyclopædia Britannica*, Encyclopædia Britannica, Inc., www.britannica.com/science/epithelium.

- Eramo, Stefano, et al. "Mozart: the Toothache of a Genius." *Journal of the History of Dentistry*, vol. 59, no. 3, 2011, pp. 129–134.
- Favaretto, Maddalena, et al. "Big Data and Digitalization in Dentistry: A Systematic Review of the Ethical Issues." *International Journal of Environmental Research and Public Health*, vol. 17, no. 7, 2020, p. 2495., doi:10.3390/ijerph17072495.
- Fine, N., et al. "Periodontal Inflammation Primes the Systemic Innate Immune Response." *Journal of Dental Research*, vol. 100, no. 3, 2020, pp. 318–325., doi:10.1177/0022034520963710.
- Flowers, Jonathan C. "Rethinking Algorithmic Bias through Phenomenology and Pragmatism." *Computer Ethics - Philosophical Enquiry (CEPE) Proceedings*, vol. 2019, 29 May 2019, doi:10.25884/mh5z-fb89.
- Frellick, Marcia, and Charles P. Vega. "How Big a Problem Is Misdiagnosis in Medicine?" *Medscape*, Medscape Education, 2 July 2020, www.medscape.org/viewarticle/933116.
- Gibson, Barry J., and C. Exley. "The Mouth." *Social Science and Dentistry*, vol. 2, no. 2, 2013, pp. 50–57.
- Gomaa, Noha, et al. "How Does the Social 'Get under the Gums'? The Role of Socio-Economic Position in the Oral-Systemic Health Link." *Canadian Journal of Public Health*, vol. 108, no. 3, 2017, doi:10.17269/cjph.108.5930.
- Grubbe, Blanche D., et al. "The Link Between Oral Health and Whole-Body Health: Is It Really, 'All in Your Head'?" *Townsend Letter*, 22 May 2021, www.townsendletter.com/article/443-oral-health-whole-body-health-connection/.
- Haider, Monica. "New Jersey Dentist Suspended after 1 Patient Dies and 14 Others Suffer Heart Infections." *CNN*, Cable News Network, 5 Feb. 2020, www.cnn.com/2020/02/04/us/nj-dentist-suspended/index.html.
- Harris, Joe. "Man Sues Dentist Who Yanked All His Teeth." *Courthouse News Service*, 23 Apr. 2014, www.courthousenews.com/man-sues-dentist-who-yanked-all-his-teeth/.
- Heath, David, and Jill Rosenbaum. "IMPACT: 'Dentists' Probe Leads to New Texas Law." *Center for Public Integrity*, 26 June 2013, publicintegrity.org/health/impact-dentists-probe-leads-to-new-texas-law/.
- Heath, David, et al. "Dentists under Pressure to Drill 'Healthy Teeth' for Profit, Former Insiders Allege." *USA Today*, Gannett Satellite Information Network, 19 Mar. 2020, www.usatoday.com/in-depth/news/investigations/2020/03/19/dental-chain-private-equity-drills-healthy-teeth-profit/4536783002/.
- Heath, David. "Dental Boards Rarely Punish Dentists for Unnecessary Treatment." *USA Today*, Gannett Satellite Information Network, 19 Mar. 2020, www.usatoday.com/in-depth/news/investigations/2020/03/19/dental-board-rarely-suspend-dentist-license-bad-dentistry/5014443002/.

- Heaven, Will Douglas. "An Algorithm That Can Spot Cause and Effect Could Supercharge Medical AI." *MIT Technology Review*, MIT Technology Review, 2 Apr. 2020, www.technologyreview.com/2020/02/05/349131/an-algorithm-that-can-spot-cause-and-effect-could-supercharge-medical-ai/.
- Herrera, Julianne. "Arizona Dental Office Sued for Wrongful Death after Girl, 4, and Boy, 2, Die after Visits." *WATE 6 On Your Side*, WATE 6 On Your Side, 7 Jan. 2018, www.wate.com/news/national-world/arizona-dental-office-sued-for-wrongful-death-after-girl-4-and-boy-2-die-after-visits/905388357/.
- Hung, Man, et al. "Artificial Intelligence in Dentistry: Harnessing Big Data to Predict Oral Cancer Survival." *World Journal of Clinical Oncology*, vol. 11, no. 11, 2020, pp. 918–934., doi:10.5306/wjco.v11.i11.918.
- "Inconsistency in Radiographic Dental Diagnostics & Treatment Planning." *DAIC*, Dental AI Council, Dec. 2020, [static1.squarespace.com/static/5f23743e90bfa971db0fe5b1/t/5fe0ccb786144b2be1731386/1608568403250/DAIC - TreatmentStudy - 12.16.2020.pdf](http://static1.squarespace.com/static/5f23743e90bfa971db0fe5b1/t/5fe0ccb786144b2be1731386/1608568403250/DAIC-TreatmentStudy-12.16.2020.pdf).
- Jabr, Ferris. "The Truth About Dentistry." *The Atlantic*, Atlantic Media Company, May 2019, www.theatlantic.com/magazine/archive/2019/05/the-trouble-with-dentistry/586039/.
- Khanagar, Sanjeev B., et al. "Scope and Performance of Artificial Intelligence Technology in Orthodontic Diagnosis, Treatment Planning, and Clinical Decision-Making - A Systematic Review." *Journal of Dental Sciences*, vol. 16, no. 1, 2021, pp. 482–492., doi:10.1016/j.jds.2020.05.022.
- Khanagar, Sanjeev B., et al. "Developments, Application, and Performance of Artificial Intelligence in Dentistry – A Systematic Review." *Journal of Dental Sciences*, vol. 16, no. 1, 2021, pp. 508–522., doi:10.1016/j.jds.2020.06.019.
- Kim, Jaeyoung, et al. "DeNTNet: Deep Neural Transfer Network for the Detection of Periodontal Bone Loss Using Panoramic Dental Radiographs." *Scientific Reports*, vol. 9, no. 1, 26 Nov. 2019, doi:10.1038/s41598-019-53758-2.
- Klingberg, G. "Dental Anxiety and Behaviour Management Problems in Paediatric Dentistry — a Review of Background Factors and Diagnostics." *European Archives of Paediatric Dentistry*, vol. 9, no. S1, 2008, pp. 11–15., doi:10.1007/bf03262650.
- Krizhevsky, Alex, et al. "ImageNet Classification with Deep Convolutional Neural Networks." *NIPS'12: Proceedings of the 25th International Conference on Neural Information Processing Systems*, vol. 1, Dec. 2012, pp. 1097–1105., doi:10.1145/3065386.
- Ledford, Heidi. "Millions of Black People Affected by Racial Bias in Health-Care Algorithms." *Nature News*, Nature Publishing Group, 24 Oct. 2019, www.nature.com/articles/d41586-019-03228-6.
- Lee, Jae-Hong, et al. "A Performance Comparison between Automated Deep Learning and Dental Professionals in Classification of Dental Implant Systems from Dental Imaging: A

- Multi-Center Study.” *Diagnostics*, vol. 10, no. 11, 2020, p. 910., doi:10.3390/diagnostics10110910.
- Lewis, Donald W., et al. “Dentists Variability in Restorative Decisions, Microscopic and Radiographic Caries Depth.” *Community Dentistry and Oral Epidemiology*, vol. 24, no. 2, 1996, pp. 106–111., doi:10.1111/j.1600-0528.1996.tb00824.x.
- Li, Guo-zheng, and Bao-yan Liu. “Big Data Is Essential for Further Development of Integrative Medicine.” *Chinese Journal of Integrative Medicine*, Chinese Association of Traditional and Western Medicine, 3 May 2015, link.springer.com/article/10.1007/s11655-015-2169-3.
- Listl, S, et al. “Professional Variability in Decision Making in Modern Dentistry: A Pilot Study.” *Operative Dentistry*, vol. 41, no. S7, 2016, doi:10.2341/14-369-o.
- 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 Digital Health*, vol. 1, no. 6, Oct. 2019, pp. e271–e297., doi:10.1016/s2589-7500(19)30123-2.
- Lovejoy, Chris. “From Correlation to Causation in Machine Learning: Why and How.” *Medium*, Towards Data Science, 31 Oct. 2020, towardsdatascience.com/from-correlation-to-causation-in-machine-learning-why-and-how-4485bca8d145.
- Mahony, Edmund H. “Dental Chain Will Pay \$24M To Settle False Billing Claims Over Unnecessary Work On Children.” *Courant.com*, Hartford Courant, 6 Dec. 2018, www.courant.com/news/connecticut/hc-dental-settlement-20180110-story.html.
- Mailonline, Claire Toureille For. “Woman Was Left TOOTHLESS after Her Toothache Turned out to Be Cancer.” *Daily Mail Online*, Associated Newspapers, 12 Oct. 2020, www.dailymail.co.uk/femail/article-8830629/Woman-left-TOOTHLESS-toothache-turned-cancer.html.
- McGregor, Milecia. “What Is a Convolutional Neural Network? A Beginner's Tutorial for Machine Learning and Deep Learning.” *FreeCodeCamp.org*, 28 Apr. 2021, www.freecodecamp.org/news/convolutional-neural-network-tutorial-for-beginners/.
- United States, Congress, Cong. House, Committee on Commerce. *Medicaid Fraud and Abuse: Assessing State and Federal Responses: Hearing before the Subcommittee on Oversight and Investigations of the Committee on Commerce, House of Representatives, One Hundred Sixth Congress, First Session, November 9, 1999*, U.S. G.P.O., 1999. 106th Congress, 1st session, document.
- Messina, Matthew J. “Why Older Adults Have Teeth: A History of Dental Care and Its Improvements.” *Generations: Journal of the American Society on Aging*, vol. 40, no. 3, 2016, pp. 19–24. *JSTOR*, www.jstor.org/stable/26556212.
- Mudrak, Jörg. “Artificial Intelligence and Deep Learning in Dental Radiology.” *Oral Health*, Oral Health Group, 10 Dec. 2019, www.oralhealthgroup.com/features/artificial-intelligence-and-deep-learning-in-dental-radiology-a-way-forward-in-point-of-care-radiology/.

- Nazir, Milady. "Improving Computer Vision for AI." *ScienceDaily*, ScienceDaily, 27 May 2021, www.sciencedaily.com/releases/2021/05/210527091439.htm.
- Nettleton, Sarah. "Power and Pain: The Location of Pain and Fear in Dentistry and the Creation of a Dental Subject." *Social Science & Medicine*, vol. 29, no. 10, 1989, pp. 1183–1190., doi:10.1016/0277-9536(89)90361-4.
- Nettleton, Sarah. "Protecting a Vulnerable Margin: towards an Analysis of How the Mouth Came to Be Separated from the Body." *Sociology of Health & Illness*, vol. 10, no. 2, 2008, pp. 156–169., doi:10.1111/1467-9566.ep11339934.
- O'Neill, Tim. "Did People in the Middle Ages Have Bad Teeth?" *Slate Magazine*, Slate, 10 Apr. 2015, slate.com/human-interest/2015/04/dental-hygiene-did-people-in-the-middle-ages-have-bad-teeth.html.
- "Oral-Systemic Health." *American Dental Association*, Department of Scientific Information, Evidence Synthesis & Translation Research, ADA Science & Research Institute, LLC., www.ada.org/en/member-center/oral-health-topics/oral-systemic-health.
- Ostermann, Thomas. "Digital Integrative Medicine – Just Knocking on Wood or a Bit More?" *Complementary Medicine Research*, vol. 26, no. 1, 2019, pp. 1–2., doi:10.1159/000497050.
- Otto, Mary. "For Want of a Dentist." *The Washington Post*, WP Company, 28 Feb. 2007, www.washingtonpost.com/wp-dyn/content/article/2007/02/27/AR2007022702116.html.
- Overjet*, www.overjet.ai/.
- Pallasch, Thomas J., and Michael J. Wahl. "Focal Infection: New Age or Ancient History?" *Endodontic Topics*, vol. 4, no. 1, 2003, pp. 32–45., doi:10.1034/j.1601-1546.2003.00002.x.
- "Patient Wins £35k Compensation after Losing 13 Teeth When Dentists Misdiagnosed Her Gum Disease." *The Telegraph*, Telegraph Media Group, 4 July 2016, www.telegraph.co.uk/news/2016/07/04/patient-awarded-35k-compensation-after-losing-13-teeth-when-dent/.
- Pauwels, Ruben, et al. "Artificial Intelligence for Detection of Periapical Lesions on Intraoral Radiographs: Comparison between Convolutional Neural Networks and Human Observers." *Oral Surgery, Oral Medicine, Oral Pathology and Oral Radiology*, vol. 131, no. 5, 2021, pp. 610–616., doi:10.1016/j.oooo.2021.01.018.
- Peckham, Michelle, et al. "The Leeds Histology Guide." *The Histology Guide*, University of Leeds, 2003, www.histology.leeds.ac.uk/.
- Penfield, Wilder, and Edwin Boldrey. "Somatic Motor And Sensory Representation In The Cerebral Cortex Of Man As Studied By Electrical Stimulation." *Brain*, vol. 60, no. 4, 1937, pp. 389–443., doi:10.1093/brain/60.4.389.
- Perkes, Courtney. "Newport Beach Celebrity Dentist from TV's 'The Swan' Surrenders License." *Orange County Register*, Orange County Register, 23 Feb. 2017,

www.ocregister.com/2017/02/23/newport-beach-celebrity-dentist-from-tvs-the-swan-surrenders-license/.

Pihlstrom, Bruce L., et al. "Promoting Oral Health Care Because of Its Possible Effect on Systemic Disease Is Premature and May Be Misleading." *The Journal of the American Dental Association*, vol. 149, no. 6, 2018, pp. 401–403., doi:10.1016/j.adaj.2018.03.030.

Reisman, Miriam. "EHRs: The Challenge of Making Electronic Data Usable and Interoperable." *P & T : a Peer-Reviewed Journal for Formulary Management*, vol. 42, no. 9, Sept. 2017, pp. 572–575.

Richens, Jonathan G., et al. "Publisher Correction: Improving the Accuracy of Medical Diagnosis with Causal Machine Learning." *Nature Communications*, vol. 11, no. 1, 11 Aug. 2020, doi:10.1038/s41467-020-18310-1.

Russell, Emma. "Whanganui Dad Dying of Brain Cancer after Being Misdiagnosed with a Tooth Abscess." *NZ Herald*, NZ Herald, 23 Dec. 2020, www.nzherald.co.nz/nz/whanganui-dad-dying-of-brain-cancer-after-being-misdiagnosed-with-a-tooth-abscess/SG646QAR56F2OVDXZIFZMMQODE/.

Schwendicke, F., et al. "Cost-Effectiveness of Artificial Intelligence for Proximal Caries Detection." *Journal of Dental Research*, vol. 100, no. 4, 2020, pp. 369–376., doi:10.1177/0022034520972335.

Semuels, Alana. "Machines and AI Are Taking Over Jobs Lost to Coronavirus." *Time*, Time, 6 Aug. 2020, time.com/5876604/machines-jobs-coronavirus/.

Shaw, Alan. "Mouth Cancer Was Missed by My Dentist." *The Sunday Post*, The Sunday Post, 23 Nov. 2016, www.sundaypost.com/fp/mouth-cancer-was-missed-by-my-dentist/.

Shieber, Jonathan. "Pearl, the Healthcare Spinout from LA-Based AI Startup, GumGum, Raises \$11 Million." *TechCrunch*, TechCrunch, 8 May 2019, techcrunch.com/2019/05/08/pearl-the-healthcare-spinout-from-la-based-ai-startup-gumgum-raises-11-million/.

Sim, Ida. "Two Ways of Knowing: Big Data and Evidence-Based Medicine." *Annals of Internal Medicine*, vol. 164, no. 8, 2016, p. 562., doi:10.7326/m15-2970.

Skrocki, Marilyn. "Standardization Needs for Effective Interoperability." *Transactions of the International Conference on Health Information Technology Advancement*, vol. 2, no. 1, Oct. 2013, pp. 76–83., scholarworks.wmich.edu/ichita_transactions/32.

Smarsh, Sarah. "There Is No Shame Worse than Poor Teeth in a Rich World – Sarah Smarsh: Aeon Essays." *Aeon*, Aeon, 28 May 2021, aeon.co/essays/there-is-no-shame-worse-than-poor-teeth-in-a-rich-world.

Stueck, Wendy, and Xiao Xu. "Public Hearing Scheduled for B.C. Dentist Alleged to Have Overbilled and Misdiagnosed Patients." *The Globe and Mail*, 5 Oct. 2018, www.theglobeandmail.com/canada/british-columbia/article-public-hearing-scheduled-for-bc-dentist-alleged-to-have-overbilled/.

- Stureborg, Rich. "Conv Nets for Dummies: a Bottom-up Approach." *Medium*, Towards Data Science, 17 May 2020, towardsdatascience.com/conv-nets-for-dummies-a-bottom-up-approach-c1b754fb14d6.
- Tandon, Divya, et al. "Present and Future of Artificial Intelligence in Dentistry." *Journal of Oral Biology and Craniofacial Research*, vol. 10, no. 4, 2020, pp. 391–396., doi:10.1016/j.jobcr.2020.07.015.
- Tch, Andrew. "The Mostly Complete Chart of Neural Networks, Explained." *Medium*, Towards Data Science, 4 Aug. 2017, towardsdatascience.com/the-mostly-complete-chart-of-neural-networks-explained-3fb6f2367464.
- Thorogood, Nicki. "Mouthrules and the Construction of Sexual Identities." *Sexualities*, vol. 3, no. 2, 2000, pp. 165–182., doi:10.1177/136346000003002004.
- Videa Healthcare*, www.videa.ai/.
- Walter, Katie L. *Middle English Mouths*. 2018.
- Wei, Jerry. "AlexNet: The Architecture That Challenged CNNs." *Medium*, Towards Data Science, 25 Sept. 2020, towardsdatascience.com/alexnet-the-architecture-that-challenged-cnns-e406d5297951.
- Weller, Frances. "Wilmington Dentist Stripped of License by NC Dental Board." <https://www.wect.com>, WECT News 6, 20 Mar. 2020, www.wect.com/2020/03/20/local-dentist-stripped-license-by-nc-dental-board/.
- Wiggers, Kyle. "Pearl Raises \$11 Million to Analyze Dental Scans with AI." *VentureBeat*, VentureBeat, 9 May 2019, venturebeat.com/2019/05/08/pearl-raises-11-million-to-analyze-dental-scans-with-ai/.
- William H. Shrank, MD. "Waste in the US Health Care System." *JAMA*, JAMA Network, 15 Oct. 2019, jamanetwork.com/journals/jama/fullarticle/2752664.
- YouTube*, Group Dentistry Now, 17 Feb. 2020, www.youtube.com/watch?v=GY1WyFtNmXc&t=646s.