



KARIM R. LAKHANI
AMY KLOPFENSTEIN

VideaHealth: Building the AI Factory

AI is not magic. It is pattern recognition in a very robust and reliable way.

– Florian Hillen, Founder and CEO

Florian Hillen, founder and CEO of VideaHealth (Videa), a Boston-based start-up that used artificial intelligence (AI)^a to provide dental diagnoses, examined the latest round of user feedback from the company's pilot programs. Since Videa's founding, he had assembled a team of dental industry veterans, software engineers, and machine learning experts to build products that analyzed x-rays and identified dental conditions such as abscesses, cavities, and tumors. He was confident that several potential markets would benefit from Videa's products and had the vision to adopt them.

Participants on both sides of the dental market—providers and payors—had an interest in Videa's products. Yet, it was not clear which would be the primary driver of adoption. The U.S. dental industry was disrupted during the COVID-19 pandemic, and sales cycles among both potential customer bases were uncertain. Thus, Hillen shifted more of his focus to building Videa's AI factory, enabling him to develop, iterate, and experiment with robust AI algorithms for both sides of the market.

An organizational framework rather than a physical factory, the AI factory lay at the heart of Videa's operations. It consisted of the company's data pipeline, labeling operations, software infrastructure, and machine learning programs. Hillen envisioned releasing software solutions that could detect a range of dental pathologies, which he and his team could tailor to a client's specific needs. He wanted to ensure that the AI factory's foundation was in place before investing in specific algorithms, as changes to the AI infrastructure would be more difficult—and expensive—as the company matured.

The team had devoted more than two years to building and iterating on the AI factory and had developed the technology in continuous dialogue with stakeholders in the dental industry. After the performance of Videa's AI surpassed human-level performance, and with recent successful pilots in both the provider and insurance markets, Hillen wanted to more actively pursue a specific market segment. He wondered which customer base would be faster in adopting Videa's technology, and how the AI factory might evolve after its launch.

^a AI was "the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings." AI engaged in learning, problem solving, perception, and reasoning, but could not fully replicate human intelligence. Source: B.J. Copeland, "Artificial Intelligence," *Britannica*, August 11, 2020, <https://bit.ly/2ZQhnc1>, accessed August 2020.

Professor Karim R. Lakhani and Case Researcher Amy Klopfenstein (Case Research & Writing Group) prepared this case. It was reviewed and approved before publication by a company designate. Funding for the development of this case was provided by Harvard Business School and not by the company. Karim R. Lakhani is an angel investor and board member of VideaHealth. HBS cases are developed solely as the basis for class discussion. Cases are not intended to serve as endorsements, sources of primary data, or illustrations of effective or ineffective management.

Copyright © 2021 President and Fellows of Harvard College. To order copies or request permission to reproduce materials, call 1-800-545-7685, write Harvard Business School Publishing, Boston, MA 02163, or go to www.hbsp.harvard.edu. This publication may not be digitized, photocopied, or otherwise reproduced, posted, or transmitted, without the permission of Harvard Business School.

Videa History: Scanning the Dental Market

Hillen's interest in AI emerged while pursuing a Master's in Computer Science and Technology and Policy at the Massachusetts Institute of Technology (MIT). During his time at MIT and as a Harvard researcher at the Laboratory for Innovation Science, he worked on projects related to genomics and radiology, which sparked his curiosity regarding AI's ability to make medical diagnoses more accurate. Confident that imaging AI had matured enough to be commercialized, Hillen drew on his medical background to analyze healthcare's different verticals. He specifically focused on those that used images, such as dentistry, genomics, and radiology.

He ultimately turned his attention toward dentistry, as most dental diagnoses and treatments relied on x-ray imaging. However, more than 30% of certain conditions were misdiagnosed.¹ Many dental pathologies also went undetected; for instance, the average clinician failed to identify almost 50% of periapical radiolucencies, darkened areas on x-rays that resulted from conditions such as cysts, infections, and tumors.² Further, multiple dentists were likely to come to different conclusions after reviewing the same x-ray, weakening diagnostic consensus. This, in turn, eroded trust among patients who received conflicting diagnoses and treatment plans from different dentists.

The high error rates were partially due to the many demands facing dentists, who were expected to perform the roles of primary care physician, radiologist, and surgeon. As a result, many dental students received a single semester of training on radiographs, even though these formed the basis for most dental diagnoses. Other errors stemmed from flaws in the x-rays themselves, such as lighting inconsistencies and equipment defects that made the images difficult to read.

Unlike many medical professionals, most dentists worked individually, which did not give them an easy opportunity to consult colleagues regarding difficult cases. Videa Clinical Operations Lead Reha Jhunjhunwala, a trained dentist, reflecting on her experience, noted: "When you make a diagnosis in dental school, you get challenged by your teachers and peers. It makes you pause and wonder if you are correct. Once you start clinical practice in dentistry, you don't have that feedback loop. There's more communication in medicine that allows radiologists to confirm or look over diagnoses."

Hillen believed that AI could address many of these issues by introducing a transparent, data-driven approach to dental diagnostics. He explained, "While AI might still have some errors, it provides dentists with an accuracy and a consistency that they currently lack."

A properly-constructed AI system would remove individual subjectivity from diagnoses, as well as environmental factors that could skew dentists' judgment. For example, radiologists were more likely to misread x-rays in the afternoon, perhaps because doctors grew tired throughout the day.³ Improving diagnostic accuracy would also translate to a higher rate of correct treatments and payments.

Dental AI had the potential to have a broader social impact. Dental conditions that went untreated could ultimately escalate into more serious issues. Some dental diseases could migrate to the rest of the body, while other oral health conditions were warning signs of systemic health problems such as diabetes and high blood pressure.⁴ Hillen noted, "AI can allow dentists to detect diseases earlier, before these diseases require serious interventions, and before oral health has an impact on systemic health."

After completing his Master's degree, Hillen raised \$1.1 million from angel investors to launch Videa. He subsequently raised \$5.4 million in seed funding from Zetta Venture Partners in San Francisco, California and Pillar Venture Capital in Boston, Massachusetts—both venture capital firms in the healthcare and AI space—and began laying the groundwork for Videa's AI factory. The company's earliest AI models focused on detecting cavities and periodontal diseases.

The U.S. Dental Industry

Hillen believed that the dental industry's fragmented nature presented unique opportunities and challenges for an AI start-up. In 2019, more than 200,000 dentists, also known as providers, practiced in the United States, while roughly two-thirds of American adults visited a dentist each year.⁵ The average general dental practice generated roughly \$700,000 in annual revenue, and specialist practices earned an average annual revenue of \$1 million, for net incomes of approximately \$200,000 and \$340,000, respectively.⁶ In 2017, slightly more than 50% of dentists practiced on their own.⁷

Some dentists who did not own their practices worked for professional corporations (PCs), supported by dental service organizations (DSOs). These numerous private-equity owned conglomerates purchased DSOs supporting dental practices and then streamlined their equipment and operations. Major DSOs included Heartland Dental, which consisted of more than 1,000 dental practices, operated in 37 states, and reported more than \$1.5 billion in annual revenue.⁸ In 2019, DSO-supported offices represented roughly 15% of U.S. dental practices, with some observers anticipating that they would represent up to 30% of the market by 2025.⁹

Another major industry player were insurance companies, or payors. In 2018, the dental insurance market was worth roughly \$187 billion worldwide, and analysts projected it would exceed \$217 billion by 2025.¹⁰ More than three-quarters of the U.S. population, or almost 250 million people, had access to dental coverage in 2016. The majority (66%) of insured individuals had private coverage, while the remainder leveraged public healthcare programs such as Medicare.¹¹ Some large U.S. insurance companies processed more than 120 million claims each year.¹²

Large insurance companies often hired licensed dentists, known as dental consultants, to facilitate claim review. A dental consultant typically assessed whether treatment was appropriate based on a patient's medical file and x-rays, and also determined whether the condition was covered under the patient's insurance plan. However, dental consultants were often only able to review a portion of total submitted claims, which resulted in accepting many claims that should have been rejected. Some experts estimated that fraudulent or unnecessary procedures comprised 7% of total U.S. dental spending, or roughly \$18 billion; of this, as much as \$6 billion stemmed from insurance fraud.¹³

The Medical AI Industry

Videa was one of many new entrants to the medical AI industry. Valued at roughly \$5 billion in 2020, analysts projected that the global healthcare AI market would exceed \$31 billion by 2025.¹⁴ In addition to diagnostics, AI could optimize repetitive administrative tasks, which increased the amount of time medical professionals could dedicate to patient care.¹⁵ Other AI applications allowed patients to self-monitor conditions from home.¹⁶ Some experts anticipated that the efficiencies gained through AI would be necessary to meet the challenges facing the global healthcare community, including a growing elderly population, who often needed more intensive forms of care, and a projected shortage of nurses and care workers.¹⁷

The industry consisted of both small start-ups and established technology companies. In the start-up space, major players included Zebra Medical Vision, an Israel-based company that applied AI to patient x-rays to identify conditions such as cancer and bone fractures, and Atomwise, which used AI to make pharmaceutical research more efficient.¹⁸ Meanwhile, companies such as Microsoft created technologies to map patients' immune systems and provide customized health recommendations.¹⁹ Alphabet (Google's parent company) developed products such as Debug, which leveraged automation, computer science, and sensors to prevent the spread of mosquito-borne illnesses.²⁰

Within dentistry, AI-based ventures included Denti.AI and Pearl, both of which used x-rays to provide diagnoses.²¹ Australia-based start-up SmileMate leveraged AI to perform dental screenings through patients' smartphones,²² while Boston-based Overjet automatically vetted dental claims for insurance companies.²³ Meanwhile, teledentistry company Candid used AI to remotely monitor patients' orthodontics plans and claimed to cost \$2,500 to \$4,500 less than traditional teeth-straightening solutions such as braces or Invisalign.²⁴

The AI Factory at Videa

While Videa faced multiple rivals in the dental AI market, Hillen believed that developing his company as an AI factory would provide a competitive advantage. It enabled him to iterate on products faster than his competitors and calibrate solutions to specific client needs. One source defined the AI factory as "the scalable decision engine that powers the digital operating model of the twenty-first-century firm."²⁵ In practice, the AI factory consisted of data, which made up the "fuel" for the factory, software infrastructure, the "pipes" that carried the data to other units, and machine learning algorithms, the "machines" that processed the data and produced the factory's outputs.²⁶ Examples of AI factories included Netflix, which developed an automated recommendation system based on users' personal viewing behavior, Google's AdWords business, which used predictive software to manage its advertising pricing auctions across its entire search infrastructure, and Ant Group's entire business supporting its financial services strategy for over 1.4 billion users.²⁷

While most companies employed some form of analytics and predictive technologies, they were typically conducted manually, perhaps by a team of internal data analysts. However, a properly-constructed AI factory automated these processes, which companies could then efficiently scale across different business units. Just as a car factory had the machinery to create distinct automobile models based on different blueprints, an AI factory could create multiple models using the same automation infrastructure. For example, Videa's AI factory could develop products that identified various dental pathologies based on the same underlying technologies.

Videa Head of AI Christian Ledig explained, "The AI factory is a paradigm that we chase. When I think about the AI factory, I think about acquiring the data that feeds into the factory, bringing it into a standardized format, and applying automated routines that normalize the data at hand. Then we automate processes that improve the quality of the data, create the machine learning model, evaluate the results, and then deploy." (See **Exhibit 1**.)

The AI factory also enabled engineers to quickly build, test, and iterate on AI models that provided dental diagnoses, which could save time and money and allowed the Videa team to rapidly respond to dentists' feedback. Hillen noted, "Clients are often very new to AI, and are not always sure what they need, so there will always be back and forth during product development. AI models require a lot of upfront investment. Some varieties can lock you in to a specific way of doing things, so you want to test them out before investing in them."

The Data Pipeline

Acquiring and cleaning data The data acquisition process was the first stage of the AI factory. At Videa, data consisted of x-rays and patient records, which the company received from DSO-supported offices.

Certain features within patient records, such as clinician labeling conventions and image formatting, varied across different dental practices. In contrast, industry-wide reporting procedures

such as Current Dental Terminology (CDT) codes^b were subject to change annually. Other records occasionally contained errors or incomplete information. To ensure consistency within its data, the Videa team developed software programs that standardized records and images before integrating them into the rest of the AI factory. Uniform data was essential for training AI models, as inconsistencies in the data sets could generate inaccurate results. The Videa team also anonymized the dental records to ensure patient confidentiality.

Over time, Videa acquired almost 10 million records that showed the trajectory of patients' teeth over a 10-year period through their partnerships with various participants in the dental market. The team prioritized data diversity during the acquisition process, as research suggested that biased data sets could prompt AI systems to produce outputs that harmed certain populations. For example, some law enforcement agencies used facial recognition software to apprehend suspects, but such technologies were more likely to accurately identify fair-skinned men than dark-skinned women.²⁸

Hillen observed, "There should be something similar to the Hippocratic Oath^c for AI. Just like dentists and doctors swear to 'first, do no harm,' diagnostic AI providers should vow to develop products that will first and foremost act in the best interest of the patient." To reduce the likelihood of bias within the AI models, the team ensured that dental records represented the entire population—including age, gender, geography, race, and socioeconomic status.

Labeling operations Following the data acquisition process, the team labeled the images to identify dental pathologies. The labels included a description of the dental condition and the features that made up the condition. For example, a labeled x-ray might have a circle or box around a carious tooth with the accompanying text 'cavity.' Annotated images were later used to teach the AI models to associate the images with the labels, which enabled them to subsequently detect similar patterns in unmarked x-rays.

Before she began labeling, Jhunjhunwala interviewed at least three experts to establish a standard process for annotating a given condition, which minimized the risk of individual bias. She then took the first pass at annotating the company's x-rays. She liaised with Ledig's machine learning team to determine the best way to annotate x-rays that could then be read by Videa's software. After she annotated enough x-rays to create a minimum viable AI model, she connected with external dentists to annotate a larger data set and scale the model. Videa hired a team of over 25 dentists to label x-rays on a part-time basis. Each dentist typically worked 20 hours per month and contributed between 2,000 and 2,500 images per condition. The Videa team then collected and combined the annotated images for the machine learning data sets.

Hillen noted that the data acquisition and annotation stages were the most critical elements of the AI factory. He explained, "Even though machine learning development seems fancy, it is just a small step in building AI products. The majority of the process is the data you get, and how you decide to label it. This will determine most of your product decisions."

To determine which pathologies to focus on, Jhunjhunwala consulted with DSOs regarding the most common conditions that their supported clinicians treated. She then analyzed Videa's internal

^b CDT codes applied a specific identifier to each dental procedure. Each year, the American Dental Association (ADA) reviewed and updated CDT codes. In 2020, the ADA removed six codes, revised five codes, and added 37 new codes. Source: Susan Rohde, "The Who, What, When and How of Dental Codes," Eide Bailly, September 18, 2019, <https://www.eidebailly.com/insights/articles/2019/9/dental-coding>, accessed September 2020.

^c The Hippocratic Oath was a set of ethical guidelines for physicians. Originally developed in ancient Greece, it was a common ethical framework among medical practitioners.

data to identify the prevalence of those conditions within its patient records. Based on her research, she determined the pathologies that Videa's first batch of AI models should detect. Clinicians used different types of dental imaging technologies to collect x-rays, so she also examined multiple image formats to determine which would be compatible with Videa's software.

Algorithm Development

Supervised learning The next facet of the AI factory built the machine learning models, drawing on the annotated data to automate tasks, detect patterns, or make predictions. The models' building blocks were **algorithms**, i.e., instructions regarding how the software should use the data.²⁹

After deciding which dental condition to target within a given AI model, Ledig and his team used a **training dataset** of annotated x-rays to teach the algorithms to recognize the features of the pathology, a process known as **supervised learning**.^d The AI then reviewed a second set of expert-labeled x-rays, known as a **validation dataset**, to assess the model's accuracy. As a final step, the Videa team evaluated the model on a **test dataset**, which simulated its performance in a real-world setting. If the model's predictions failed to consistently match the experts' diagnosis, the team adjusted the algorithms or sought new training and validation datasets (see **Exhibit 2**).

Images in training and validation datasets were annotated through a single and multi-clinician approach to ensure quality control, especially for edge cases. The machine learning team also developed testing protocols that accounted for subjectivity within dental diagnoses. Ledig noted:

It is challenging to come up with the truth when examining x-rays. There is not always a black-and-white answer. There is a large grey zone where different experts would disagree, or where the conclusion cannot be made from the image itself, and you would need to look at patient history or additional information. In these cases, we consult multiple experts and then aggregate their responses to create a reference point. We use these reference points to automatically test our algorithms as we develop them, and monitor whether our algorithms improve over the development cycle.

The accuracy of AI models tended to improve over time as they were exposed to more data (see **Exhibit 3**). Thus, machine learning initially used training sets of several hundred images to train and calibrate a model, then expanded to larger sets as the AI became more sophisticated. Videa's AI ultimately aimed to match or exceed the diagnostic accuracy of dentists.

Modular design Wherever possible, Videa's machine learning engineers prioritized simplicity within algorithm design, as complex models were often more challenging to adapt to new use cases. Ledig also encouraged his team to break down complex algorithms into smaller components that they could quickly test and adjust.

"There are big advantages to modularization," explained Ledig. "You can mix and match different building blocks as the company scales. If you want to apply a similar machine learning problem to a different medical use case, ideally, you would only replace your raw data and annotations, and the rest of the model would be largely the same."

An emphasis on AI systems at the micro level also enabled the Videa team to identify potential improvements within the algorithm-development process. By constructing, testing, and adjusting

^d Other AI approaches included unsupervised learning, in which an AI model drew insights from large data sets, and reinforcement learning, which used software to optimize a model's performance within a pre-determined set of parameters.

individual pieces of an AI model, the team determined which tasks they could accomplish more efficiently. Ideally, these building blocks of the AI infrastructure would ultimately be automated, which would allow the AI factory to expand at a faster rate.

The Experimentation Platform

Real-time experimentation As product accuracy improved, Videa tested its models with clinical practitioners. The early-stage pilot recruited a single clinician to use the AI to evaluate x-rays after a patient visit. Later pilots expanded to multiple clinicians who used the products as a real-time diagnostic tool with patients.

“We want to get the product into the hands of final users sooner rather than later so that we can get feedback, make changes, and then give them a new version,” explained Jhunjhunwala. “It keeps users engaged, because they like to see the changes they requested appear in the product.” For a pilot, she acted as a liaison between clinical practitioners and the machine learning team, translating the technical aspects of dentists’ feedback so that engineers could incorporate it in a new version of an AI model.

To test models on human subjects, Videa had to secure permission from an Institutional Review Board (IRB), an independent governing body that ensured that researchers adhered to ethical research standards. After initially approving an experiment, the IRB had to approve any subsequent changes within the research model. To approve major changes within the pilots, the team submitted amendments to Videa’s original research proposal and regularly communicated with the IRB to expedite the review process. This minimized the regulatory delays that the company encountered and allowed them to continue to rapidly iterate on new products.

Model sensitivity During the model-evaluation process, the Videa team had to determine the ideal tradeoff between sensitivity and precision in their models. As sensitivity increased, AI models became more likely to detect pathologies in x-rays. While higher degrees of sensitivity reduced the likelihood of missed diagnoses, also known as false negatives, they increased instances of misdiagnoses, or false positives. Thus, a highly sensitive AI model would be more likely to incorrectly detect cavities among patients with healthy teeth, while a less sensitive version of the same model would be more inclined to overlook cavities.

To determine the appropriate balance, the machine learning team proposed different sensitivity levels at which Videa’s models could operate, then solicited feedback from clinicians to determine the best setting for a given product. “If false positives can be easily ruled out at a low overhead cost, there are use cases where high sensitivity and low precision are acceptable,” explained Ledig. For example, a dentist could quickly check the accuracy of an AI-based cavity diagnosis during a patient visit, while a false positive for an oral cancer diagnosis would subject a patient to an additional round of tests.

Software Infrastructure

The final element of Videa’s AI factory consisted of its software infrastructure, which determined the process by which Videa’s patient files and x-rays were cleaned, stored, analyzed, and transmitted throughout the data pipeline. Videa Head of Engineering Shahid Jabbar explained, “We are responsible for getting the data from our partners and putting it in a form that the AI team can query.” Like the AI models, the algorithms within the software infrastructure were constructed in a modular format, which allowed teams to work at their own pace and adjust individual algorithms as needed.

The Videa software team wrote the code that cleaned and normalized new patient files. The programs had to have the flexibility to accept changes within the data pipeline, such as new CDT codes

or file formats, and account for potential pain points in the process, such as uploads for different patients with the same name. The engineers stored the clean, unlabeled data in a company-wide, searchable repository called a **data lake**, which aggregated all data within the company and ensured that all departments in the AI factory were working with the same information. It also enabled the Videa team to retrieve large data sets quickly; for example, Jhunhunwala could hypothetically query the data lake for all records containing impacted wisdom teeth, which she could then retrieve for labeling operations. The software systems were also responsible for maintaining data security across every stage of the AI factory.

In addition to the company's internal programs, the software team developed its external user interfaces, such as how clients uploaded x-rays to Videa and received diagnoses. After a client uploaded a file, it was sent to Videa's servers, which processed the file, generated a diagnosis, and sent the findings back to the client (see **Exhibit 4**). The Videa team aimed to send results in real-time; however, processing time depended on clients' system bandwidth. Processing batches on the magnitude of 100,000 or more images could also delay the turnaround time for AI results. To manage client expectations, the software team developed interfaces to communicate the AI system's progress while it was working. The team also developed distinct user interfaces for Videa's different customer segments. They tailored each interface to individual client needs and continually made adjustments.

As the product development progressed, the team grappled with the appropriate way in which Videa should communicate with users. Jabbar explained, "One of our greatest challenges has been determining the persona of the AI. Is the AI an assistant, a collaborator, or the expert? We don't want the AI system to give the feeling of being a direct competitor to the dentist, so the challenge is to communicate the AI findings in a way that is respectful to dentists' expertise." To empower users, Jabbar and his team focused on positioning Videa's AI as an assistant offering a second opinion.

While Jabbar led the development of Videa's software systems, he was mindful that they would need to scale alongside the rest of the AI factory. Thus, he developed programs that were resilient and could adapt to the changing needs of the company. For some systems, such as the process for normalizing data, the software team wrote code with "blank" steps that, while not necessary during the early days of the AI factory, would be essential once Videa grew.

Taking Videa to Market

Hillen considered his company's progress. He had spent years laying the groundwork for its AI factory and iterating with early customers on the product value. Videa's products were able to detect 25% more diseases on average than dentists, and were able to detect many conditions that might appear on a radiograph. Satisfied that Videa's AI factory was in a position to successfully scale, he considered which of the company's potential customer bases would be most receptive to Videa's dental AI technology. He was confident that Videa's AI factory would enable his company to rapidly develop the right product for any given client. Still, the lack of industry-wide understanding of AI and the pandemic-related disruptions to dentistry posed potential barriers to uptake. He knew that Videa's products would benefit both dental payors, which consisted primarily of insurance companies, and providers, that is, dentists and DSOs. Weighing Videa's capabilities against the industry's realities, Hillen considered the opportunities of each of the company's potential customer bases.

The Payor Market

Hillen believed that Videa might appeal to insurance companies, who could use the AI to triage submitted claims. The AI systems could automatically approve some claims while flagging more

complicated cases for further review. Under such a scenario, dental consultants would have a greater bandwidth to detect and reject false claims. Insurance companies would also save money, as consultants would spend less time reviewing claims that would ultimately be approved. Hillen noted, “Right now, 60% to 80% of claims that are reviewed are approved, and then insurers have to pay for both the claim review and the treatment.” He estimated that Videa could generate 30% to 60% in claim review cost savings for insurers. Conversations with insurance companies led Hillen to believe that dental consultants were paid roughly \$6 per claim and could review an average of 30 claims per hour; he estimated that the dental insurance industry spent \$300-\$500 million on dental consultants per year.

There were strong signals by the dental insurance market that it was ready to adopt Videa’s AI technologies. Steve Pollock, President and CEO of DentaQuest, a U.S. dental insurance company and early Videa client, noted, “AI has the potential to transform the dental industry. There are opportunities across all segments of the industry, but from a dental insurer perspective, AI holds the potential to drive greater consistency and efficiency while improving the dental experience for members and providers. From stronger detection of fraud, waste, and abuse or real-time review and approval for covered benefits, we are excited about the opportunities that AI will unlock.”

However, Michael Hahn, former Chief Dental Officer of Cigna, one of the largest U.S. dental insurers, and Chief Clinical Advisor at Videa, believed that AI was unlikely to eliminate the need for consultants. In some states, he noted, “denial of a service can only be done by a licensed dental consultant. States do not want people denying claims who are not experts or who have an incentive to deny as many claims as possible. AI could present the same types of regulatory issues.”

The Provider Market

Dentists Hillen considered selling the product directly to dentists, who could use the AI to identify issues they may have overlooked when they initially evaluated x-rays. Videa also could increase trust between dentists and patients: if dentists confirmed their diagnoses through AI technology, patients might be more likely to accept their treatment plans. Hillen estimated that Videa could increase dentists’ revenue by as much as 15% each year through increased disease detection and higher patient acceptance of proposed elective treatments. Videa could also be a marketing tool, as patients might gravitate towards providers who guaranteed transparency.

Videa technology could also expedite communications between dentists and insurers. Dentists often submitted proposed treatment plans to insurers before a procedure. Dental consultants then reviewed the plans and provided cost estimates based on patients’ insurance coverage. The process typically took between one to three weeks, which could result in lost revenue for dentists.

Hahn explained, “Patients sometimes change their minds between the time they receive the diagnosis and when the pre-treatment estimates come back. The condition might not be bothering them anymore, and they may decide to delay the treatment.” Videa’s AI could provide dentists with real-time claims reviews and cost estimates, reducing the risk that patients would opt to defer treatment.

Unlike some emerging medical AI products, positioned as a substitute for doctors, Videa was a tool to make dentists more accurate.³⁰ The team believed that this could increase its appeal among dentists. Jhunjhunwala explained, “No one’s job is threatened here. Part of a dentist’s role is diagnosis, but the lucrative part of their job is treatment. An AI system can’t pull a tooth or do a filling.” While dentists tended to embrace new technologies, many were unaware of the use of AI within dentistry. Thus, there was a risk that demand for Videa would be low among individual practitioners.

Dental Service Organizations Hillen believed that Videa's products could benefit DSOs and their supported doctors who had a mission to ensure consistently high levels of patient care across different offices and practitioners. Videa could identify dentists who were not providing quality levels of care, which would allow DSOs to offer them training and mentoring programs to improve their accuracy and quality. Dr. Seth Gibree, Senior Director of Clinical Advocacy at Heartland Dental, one of Videa's early clients, elaborated:

Our supported dentists and their patients are already seeing the value, from our pilots with Videa, of bringing AI into the "chair." Our newer supported dentists share that they feel more confident, their patients' trust in the diagnoses is enhanced, and it helps them move forward with the treatment they need and desire. We see it becoming an advantage as we help recruit dentists and attract new patients for the supported offices. Ultimately our goal is to support dentists and their teams as they deliver high quality dental care and experiences to their communities—this technology has great potential to assist in achieving that goal.

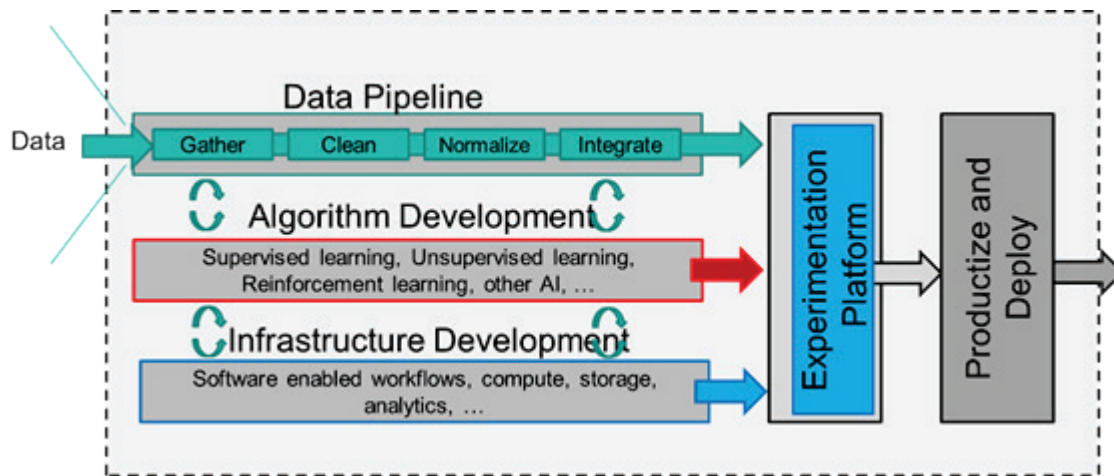
Hillen also thought AI would appeal to the dentists who tended to gravitate toward DSO-supported offices. He explained, "DSOs attract a lot of young dentists. They tend to lack experience, and they do not want to come across to patients as evil doctors suggesting a lot of treatments. So they tend to be very conservative when discussing diagnoses with patients." Hillen believed that Videa could bolster the confidence of less experienced dentists by providing a second opinion on diagnoses. This, in turn, could increase the rate of diagnoses and treatments and improve patient care levels and consistency across DSO-supported offices.

DSOs had the organizational infrastructure to scale the Videa technology across their numerous offices, perhaps making them and their supported doctors a more appealing market than individual clinicians. Hillen also suspected that consolidation within the dental industry would increase in the immediate future. He believed that the COVID-19 pandemic might prompt some veteran dentists to retire early and sell their practices to PCs supported by DSOs. Selling to a few large organizations rather than thousands of small offices seemed more feasible for a start-up like Videa, particularly as the DSO model gained market share within dentistry. Nevertheless, the DSO model represented a smaller market opportunity than individual dentists.

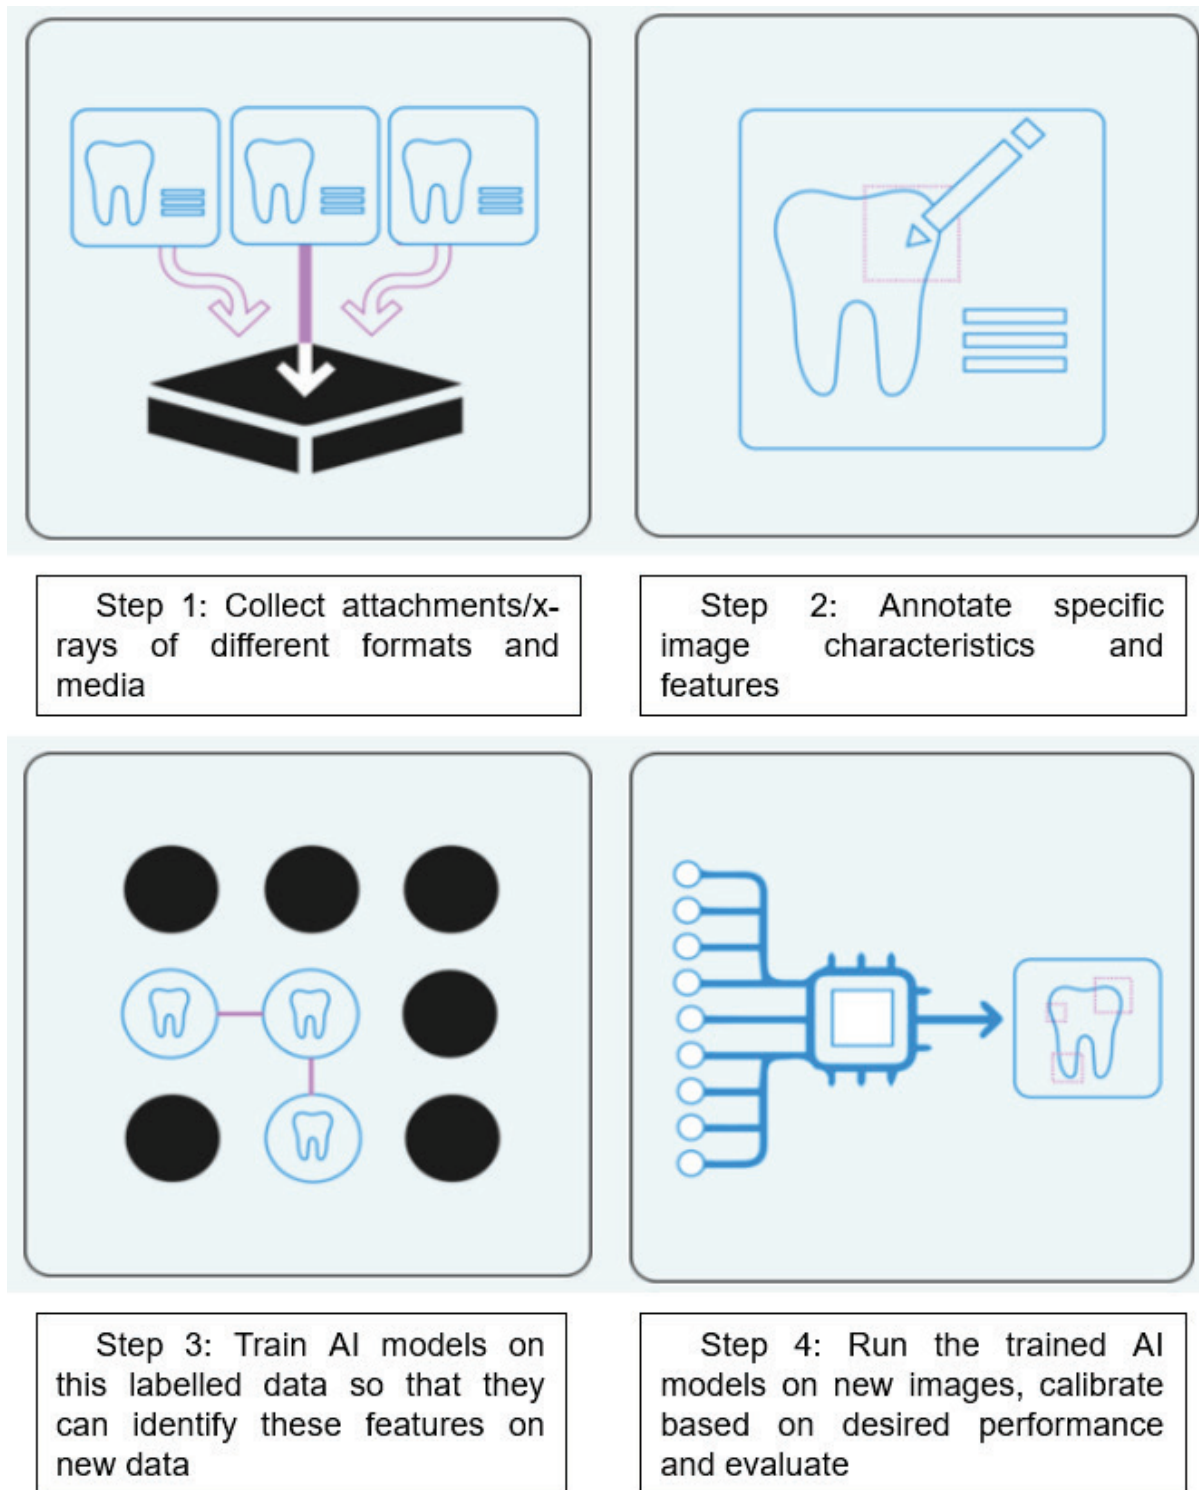
Moving Videa Forward

Hillen took another look through the results of Videa's recent pilots. As he reviewed the latest feedback, he became convinced that his company was ready to go to market. Videa's products would add value to any of its potential customer bases, but he was uncertain which would be the most likely adopter. Should he focus on dentists, DSO-supported dentists, or insurance companies?


He also reflected on the AI factory he and his team had built. How would Videa's AI infrastructure adapt as the company scaled? Would the AI factory allow it to stand out in the crowded dental AI field? Hillen was sure of one thing: Videa's AI systems would continue to evolve as they expanded to new use cases, ingested more data, and incorporated new user feedback. "AI and data analytics will improve the care of every dental patient in the not-too-distant future," he said. "But the transformation ahead involves many players, and we have to chart the right path to drive the change carefully."

Exhibit 1 AI Factory Components


Source: Marco Iansiti and Karim R. Lakhani, *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World* (Boston: Harvard Business Review Press, 2020), p. 53.

Exhibit 2 Steps of the Videa Machine Learning Process

Source: Company website.

Exhibit 3 The AI Factory's Virtuous Cycle

Source: Marco Iansiti and Karim R. Lakhani, *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World* (Boston: Harvard Business Review Press, 2020), p. 54.

Exhibit 4 Example of Videa AI Output

Source: Company documents.

Endnotes

- ¹ Michael G. Endres, et al., "Development of a Deep Learning Algorithm for Periapical Disease Detection in Dental Radiographs," *Diagnostics* vol. 10 no. 6 (June 2020): 430.
- ² Endres, et al., "Development of a Deep Learning Algorithm for Periapical Disease Detection in Dental Radiographs."
- ³ A.S. Alshabibi et al., "Effects of time of day on radiological interpretation," *Clinical Radiology*, vol. 75, no. 2 (February 2020): 148-155.
- ⁴ Prashant Kumar, KMK Mastan, Ramesh Chowdhary, and K. Shanmugam, "Oral manifestations in hypertensive patients: A clinical study," *Journal of Oral Maxillofacial Pathology* 16(2) (Summer 2012):215-221.
- ⁵ American Dental Association, "Workforce," <https://bit.ly/3uxqxsA>; America's Health Rankings, 2019 Annual Report, <https://bit.ly/3qRDqLh>, both accessed August 2020.
- ⁶ "Income and Gross Billings," American Dental Association, <https://bit.ly/3aOuiBo>, accessed August 2020.
- ⁷ "HPI: Fewer dentists are practicing solo," American Dental Association, October 10, 2018, <https://bit.ly/3qNSKbI>, accessed August 2020.
- ⁸ Moody's, "HEARTLAND DENTAL, LLC - Moody's assigns a B3 rating to Heartland Dental's new \$200 million incremental First Lien Term Loan," Yahoo News, August 11, 2020, <https://yhoo.it/3pSWvvl>, accessed September 2020.
- ⁹ Agam Shah, "AI Can Lead to Lower Dental Bills," *The Wall Street Journal*, July 16, 2019, <https://on.wsj.com/3qSbjLL>, accessed August 2020.
- ¹⁰ "According to the report, the global dental insurance market was approximately USD 135.5 billion in 2018 and is expected to generate around USD 217.6 billion by 2027, at a CAGR of around 5.4% between 2019 and 2027," GlobalNewswire, May 20, 2019, <https://bit.ly/3armGmt>, accessed August 2020.
- ¹¹ "Who has dental benefits today?," National Association of Dental Plans, <https://bit.ly/3dIELQN>, accessed December 2020.
- ¹² "The Delta Dental Organization," Delta Dental, 2018, <https://bit.ly/3uoBA6R>, accessed October 2020.
- ¹³ "Pearl AI Computer Vision Breakthrough Allows Dental Insurers To Review 100% of Claims-based Oral Imagery With Superhuman Precision," PR Newswire, January 27, 2020, <https://prn.to/2Mlx5sn>, accessed September 2020.
- ¹⁴ "Artificial Intelligence in Healthcare Market Worth \$45.2 Billion by 2026-Exclusiv Report by MarketsandMarkets," press release, PRNewswire, June 9, 2020, <https://prn.to/2Mkl52>, accessed August 2020.
- ¹⁵ Angela Spatharou, Solveigh Hieronimus, and Jonathan Jenkins, "Transforming healthcare with AI: The impact on the workforce and organizations," McKinsey & Company, March 10, 2020, <https://mck.co/3sqhabC>, accessed August 2020.
- ¹⁶ Spatharou, Hieronimus, and Jenkins, "Transforming healthcare with AI: The impact on the workforce and organizations."
- ¹⁷ Spatharou, Hieronimus, and Jenkins, "Transforming healthcare with AI: The impact on the workforce and organizations."
- ¹⁸ Zebra Medical Vision, "AI1 Solutions," <https://www.zebra-med.com/solutions>; Atomwise, "Our Technology," <https://www.atomwise.com/our-technology/>, both accessed August 2020.
- ¹⁹ Microsoft, "Microsoft Immunomics," January 4, 2018, <https://bit.ly/3r4Xf1X>, accessed August 2020.
- ²⁰ "Stopping mosquitos that can spread disease with mosquitos that can't," Verily, October 6, 2016, <https://verily.com/stories/stopping-mosquitoes-that-can-spread-disease-with-mosquitoes-that-can-t/>, accessed August 2020.
- ²¹ Denti.AI, <https://www.denti.ai/>, Pearl, "Products," <https://hellopearl.com/products>, both accessed August 2020.
- ²² Angela Tufvesson, "Artificial intelligence in dentistry," *Bite Magazine*, December 16, 2019, <https://bit.ly/3pN21zD>, accessed August 2020.
- ²³ Alex Wilhelm, "Overjet raises \$7.85M for its dental-focused AI tech," Techcrunch, June 2, 2020, <https://tcrn.ch/3sIHvaX>, accessed August 2020.
- ²⁴ Candid, "How It Works," <https://www.candidco.com/how-it-works>, accessed August 2020.

²⁵ Marco Iansiti and Karim R. Lakhani, *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World* (Boston: Harvard Business Review Press, 2020), p. 53.

²⁶ Iansiti and Lakhani, *Competing in the Age of AI*, p. 58.

²⁷ Iansiti and Lakhani, *Competing in the Age of AI*, p. 53.

²⁸ Joy Buolamwini, "Artificial Intelligence Has a Problem With Race and Gender Bias. Here's How to Solve It," *Time*, February 7, 2019, <https://bit.ly/37GHWog>, accessed August 2020.

²⁹ Iansiti and Lakhani, *Competing in the Age of AI*, p. 62.

³⁰ Jonathan Shieber, "Using AI to improve dentistry, VideaHealth gets a \$5.4 million polish," *Techcrunch*, September 25, 2019, <https://tcrn.ch/3pRBjFY>, accessed August 2020.