https://pubsonline.informs.org/series/educ

TUTORIALS IN OPERATIONS RESEARCH
 | ISBN 979-8-9882856-1-8
https://doi.org/10.1287/educ.2023.0256

Incorporating Artificial Intelligence into Healthcare Workflows: Models and Insights

Tinglong Dai,[a,b,c,d,e,*] **Michael D. Abràmoff**[f,g]
[a] Carey Business School, Johns Hopkins University, Baltimore, Maryland 21202; [b] Hopkins Business of Health Initiative, Johns Hopkins University, Baltimore, Maryland 21202; [c] School of Nursing, Johns Hopkins University, Baltimore, Maryland 21205; [d] Institute for Data-Intensive Engineering and Science, Johns Hopkins University, Baltimore, Maryland 21210; [e] Johns Hopkins Center for Digital Health and Artificial Intelligence, Baltimore, Maryland 21202; [f] Department of Ophthalmology and Visual Sciences, University of Iowa, Iowa City, Iowa 52242; [g] Digital Diagnostics Inc., Coralville, Iowa 52241
*Corresponding author
Contact: dai@jhu.edu, https://orcid.org/0000-0001-9248-5153 (TD); michael-abramoff@uiowa.edu, https://orcid.org/0000-0002-3490-0037 (MDA)

Abstract Artificial intelligence (AI) is poised to revolutionize healthcare delivery in the United States and around the world. As AI becomes an integral part of the healthcare workflow, it will change the way we model and analyze healthcare delivery and upend the paradigm that has dictated how operations research and management science researchers interact with healthcare practitioners. In this tutorial, we demonstrate how the integration of AI into the healthcare workflow will radically transform healthcare delivery and, at the same time, require a new set of models to guide rapidly changing healthcare practices, measure productivity gains in the industry, and reduce disparities in access to care. These models should be based on a comprehensive understanding of the variables that influence various stakeholders, including patients, providers, payers, bioethicists, regulatory bodies, and investors. Although healthcare AI promises to learn and adapt based on user interactions and data, the development, validation, and approval processes require the creation of new models that generate useful insights. Finally, we discuss barriers and opportunities related to regulatory and reimbursement issues for AI in healthcare.

Funding: T. Dai acknowledges financial support he received through the Johns Hopkins Discovery Award (2022–2024) and the Hopkins Business of Health Initiative (HBHI) Pilot [Grant 2022-2023]. M. D. Abramoff was funded by Research to Prevent Blindness, based in New York; the National Eye Institute under [Grant R01EY033233]; and the Robert C. Watzke, MD, Professorship.

Keywords artificial intelligence • healthcare workflow • productivity • access to care

1. Introduction

The saga of artificial intelligence (AI), since its inception in 1956, has been a thrilling odyssey of transformation, with the term itself taking on new meanings as it is associated with an ever-evolving array of techniques and capabilities. In its early days, AI research was dominated by "symbolic AI" (also known as "rule-based AI"), which focused on identifying, formalizing, and codifying the logic and rules underlying human intelligence (Simon [74]). Over time, amid fierce competition among various disciplines and the evolution of computational infrastructure and algorithms, "learning-based AI" (also known as "connectionist AI"), which

focuses on recognizing patterns from data that capture human thought and action, has become the mainstream school of thought. In particular, deep learning, using artificial neural networks, has emerged as the workhorse of AI, although the reason for its "unreasonable effectiveness" remains more or less a mystery (Sejnowski [72]).

One thing remains constant: Medicine has always been one of the most prominent areas of AI application (Kavasidis et al. [49]), inspiring AI and medical researchers to develop new techniques in the quest for higher quality, lower cost, more accessible, more equitable, and more personalized healthcare. It is no exaggeration to say that the field of medicine has played a significant, perhaps even pivotal, role in shaping the trajectory of AI as we know it today.

In this tutorial, we provide not only an overview of AI applications in medicine, but also a preview of what we believe will be one of the most important changes in the healthcare industry—the increasingly widespread integration of AI into healthcare workflows, which will create new opportunities in healthcare delivery and require new models, new thinking, and new insights. We will use the case of AI-enabled medical devices as a focal point to guide our discussions, with a discussion of other types of medical AI toward the end of this tutorial.

1.1. A Brief History of AI and Medicine

AI and medicine have a surprisingly intimate relationship. Neural networks, the key building blocks of modern AI, were first developed by McCulloch and Pitts [60], inspired by the workings of the human brain, 12 years before the term "artificial intelligence" was coined by a group of computer scientists who met in an eight-week-long conference at Dartmouth College in 1956.

The Dartmouth meeting marked the beginning of the first AI boom, which lasted nearly two decades. During this period, the dominant paradigm of AI was symbolic AI (or "rule-based AI"). That is, the focus was on logic: identifying a set of conditional rules (coded using "`IF-THEN-ELSE`" clauses) to perform intellectual functions typically performed by humans. For example, Warner et al. [87] use the Bayes theorem to develop a set of rules for diagnosing congenital heart disease. The end result of their exercise is a checklist of symptoms, so that as the physician identifies different sets of symptoms, the model can update its prediction of whether the patient has a particular disease from a list of diseases. Beginning in 1974, an "AI winter" hit the field, mainly due to the lack of progress in developing practically usable AI systems with tangible results. By the early 1980 s, many AI projects had been canceled or postponed (Hendler [44]).

The 1980 s saw a second AI boom, characterized by various "expert systems" that attempted to represent human knowledge in a multidisciplinary way: linguists, mathematicians, philosophers, psychologists, and economists, as well as domain experts, worked alongside computer scientists and engineers to develop specialized expert systems. The focus was narrower and less ambitious than what was envisioned during the initial AI boom. Most of these expert systems remain symbolic (i.e., rule-based), despite advances in learning algorithms. For instance, the MYCIN system uses about 60 distinct rules to pinpoint bacteria responsible for severe infections. It does this by posing a series of yes/no questions to physicians, which in turn generates a list of potential bacteria that could be causing the patient's symptoms (Swartout [79]). We refer the reader to Lundsgaarde [59] for a review of notable medical expert systems from the second AI boom. However, the second AI boom came to an abrupt end in 1987 when expert systems failed to deliver tangible results in the real world (Gill [37]). This setback was compounded by a change in U.S. national funding priorities (Giacaglia [35]), resulting in the onset of the second AI winter, which lasted until the mid-1990 s. Such systems were also hampered by their reliance on what essentially amounted to a "game of noisy telephone": The patient conveys their symptoms through speech, which the physician then interprets and types into a teletype for the AI to make decisions based on. The communication can be noisy at each step, leading to potential errors.

The end of the second AI winter is followed by a period of normalization of AI that can be called the "AI spring." This was a period in which important algorithmic and hardware

breakthroughs were made, but it was not until around 2011 that it became apparent that AI was entering another, still ongoing, boom, at least partially due not so much to more sophisticated algorithms, but the availability of low-cost sensors across medicine, allowing easy access to high fidelity data.

The following developments are thought to play major roles in the current AI boom, which some refer to as the first "AI summer":

(1) Faster, widespread use of faster CPUs and GPUs, allowing more sophisticated, multilayer neural networks—deep learning revolution—allowing higher accuracy and generalization (LeCun et al. [54], Sejnowski [72]).

(2) Reinforcement learning (RL), a type of machine learning in which an agent learns to make decisions by interacting with its environment and receiving feedback in the form of rewards or punishments, is gaining popularity due to the success of projects such as `AlphaGo Zero`, a version of DeepMind's AlphaGo software that taught itself to play Go from scratch without human input, and others (Fu et al. [33]).

(3) Generative AI, a branch of AI, focused on devising models that can generate creative content such as text, images, videos, or music based on learned patterns, is expanding AI's horizons, with significant implementations like `DALL·E 2`, an advanced iteration of the `DALL·E` model known for creating original images from textual descriptions, `ChatGPT`, a powerful conversational AI model capable of generating human-like text, and `GPT-4`, the latest version of the large-scale transformer-based language model trained by OpenAI, known for its improved comprehension and generation capabilities.

(4) AI is increasingly embedded in workflows, not just confined to startups and AI labs (Agrawal et al. [5]). In manufacturing, AI-powered robots and algorithms are being used by established manufacturers and start-ups alike to increase productivity on production lines and predict maintenance needs to prevent downtime (Wuest et al. [90]). In healthcare, this trend has been enabled by (1) availability of high-quality digital data in medicine through widespread use of low-cost CMOS sensors and the like, replacing mostly noisy, primarily textual data in many areas (Abràmoff et al. [1]); and (2) ethical frameworks for AI, including metrics for those ethics, so that AI systems can be optimized to them, allowing all healthcare stakeholders (patients, physicians, payers, other providers, regulators, ethicists, AI creators, and investors) to support subsequent ethical regulation and reimbursement (Char et al. [20], Abràmoff et al. [2, 3], Youssef et al. [94]).

Notably, in the field of healthcare, which has led to the legal use of autonomous AI for patient care (i.e., a computer making a medical decision), the ethical framework and subsequent regulation and reimbursement have played central roles. Compare this to, for example, automotive AI, where the same deep learning, reinforcement learning, and generative AI technologies are used, but there is no ubiquitous autonomous car (Insurance Institute for Highway Safety [48]).

All of these developments have had a significant impact on healthcare. In particular, computer vision has arguably been the most influential AI development, resulting in the vast majority of the more than 500 AI-enabled devices approved for clinical use by the U.S. Food and Drug Administration (FDA) as of July 2022 (FDA [32]). The second development, RL and generative AI, has yet to see large-scale real-world applications in medicine (Coronato et al. [22]). The third development, generative AI, has sparked quite a lot of enthusiasm among the medical community (Ayers et al. [9]), leading to a number of proposals, such as a series of "foundation models" of "generalist medical artificial intelligence" by Moor et al. [62]. The topic of this tutorial is the fourth development—how to integrate AI into everyday healthcare workflows and make a difference in healthcare delivery.

1.2. Diabetic Retinopathy: A Case of Using AI in Healthcare Delivery

Diabetic retinopathy, a sight-threatening disease, is a serious public health issue in the United States (and much of the rest of the world). Thirty-seven million people in the United States,

or about 11% of the population, have diabetes; Black adults have a 60% higher risk of being diagnosed with diabetes than white adults (HHS [46]). Diabetic retinopathy is a common complication of diabetes (both type 1 and type 2): Between one in three and one in two patients with diabetes will develop retinopathy in their lifetime (Graue-Hernandez et al. [39]).

Diabetic retinopathy is the leading cause of blindness among working-age adults in the United States (CDC [17]). Fortunately, early detection and treatment can reduce the risk of blindness by 95% (National Eye Institute [64]). Therefore, early detection is the key to preventing vision loss from retinopathy. Yet only 15% of diabetic patients in the United States receive the recommended annual screening (Benoit et al. [12], Channa et al. [19]). The screening rate is particularly low among minority patients and those without insurance (Eppley et al. [30]).

A major reason for the low screening rate in the United States is the inconvenience associated with scheduling an appointment with an ophthalmologist or optometrist, which we will refer to as an "eye care professional" for the remainder of this tutorial. The procedure can be costly for patients, especially those who are uninsured or underinsured, and the dilated eye exam can take two hours. Because an estimated 82% of diabetic patients who are not screened for retinopathy visit a primary care physician annually, using AI to provide screening in primary care settings is an important opportunity to improve retinopathy screening rates among diabetic patients (Gibson [36]).

This is exactly what LumineticsCore (previously known as IDx-DR, as illustrated in Figure 1), the first FDA-cleared and Centers for Medicare & Medicaid Services (CMS) reimbursed autonomous AI-enabled diagnostic device, aims to do. Developed by Digital Diagnostics, it uses a hybrid biomarker-based deep learning computer vision algorithm that integrates multiple statistically partially dependent deep learning detectors. The device is designed to operate in a primary care setting by personnel with a high school diploma, eliminating the need for expert human interpretation or data entry. LumineticsCore identifies diabetic eye disease with sensitivity and specificity values exceeding 90%, and the process can be completed within a short time frame (Abràmoff et al. [1]). The term "autonomous" applies as the device does not require human interpretation of images and the medical decision making is algorithmic, aligning with the AI policy of the American Medical Association that places liability with the AI developer (American Medical Association [7]). Moreover, the device can be operated by any member of the clinical staff (with a high school diploma or higher) after a four-hour training program, indicating an upskilling opportunity in a field experiencing a shortage of medical professionals, especially in rural and inner-city regions (Butkus et al. [16]).

Figure 1. (Color online) LumineticsCore (previously IDx-Dr), the first autonomous AI-enabled diagnostic device approved by the FDA for clinical use.

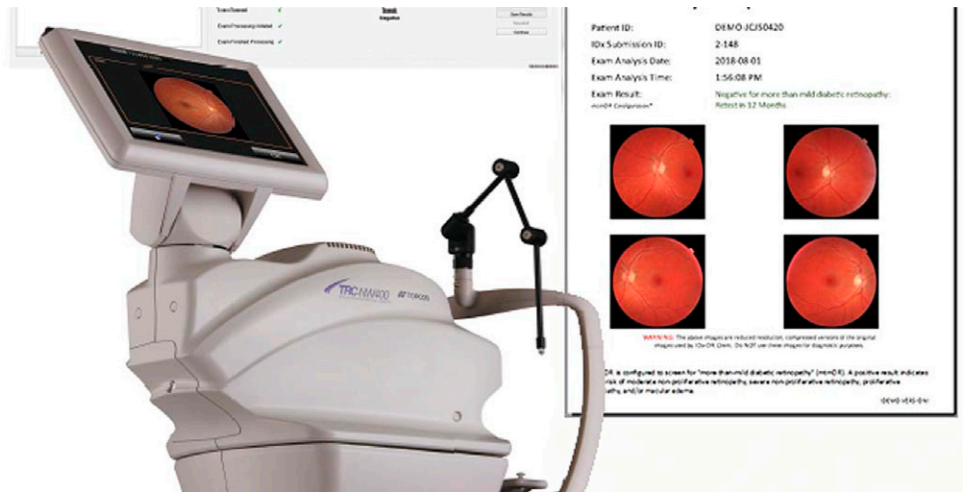

Several studies have demonstrated the beneficial impact on patient outcomes of LumineticsCore. For example, Wolf et al. [88] found that using the device in a primary care setting increased patient adherence from 49% to 95%. In a more recent investigation, Wolf and colleagues conducted a randomized controlled trial to determine if point-of-care AI screening improved screening rates compared with standard in-person eye exams by eye care professionals. They found that the group offered the AI screening option had a screening rate of 100%, whereas the group referred to an eye care professional for screening had a screening rate of 22%. These results offer hope for improving retinopathy screening rates by offering the AI screening option in a primary care setting. Another recent study, by Leong et al. [57], shows that the use of autonomous AI for retinopathy screening improves health equity for high-risk patients by facilitating evaluation of such patients before the occurrence of visual damage.

The case of LumineticsCore has several notable features with broader implications for AI-enabled healthcare.

First, LumineticsCore is an FDA-cleared medical device, the type of "tangible AI" that is certified for clinical use as medical device. As of July 2022, the FDA has approved 521 AI-enhanced medical devices across 15 medical specialties, a significant portion of which were developed using deep learning computer vision algorithms and are intended for medical diagnosis and screening (FDA [32]). The number of FDA-approved AI devices is expected to increase dramatically in the coming years, so we can expect that the experience of integrating LumineticsCore into medical practice will provide valuable lessons for using AI to bridge gaps in access to care.

Second, a shortage of healthcare workers is often the primary reason for a lack of access to essential care, especially preventive care for patients with chronic conditions. The gap between supply and demand is expected to widen in the coming decades in both developed and developing countries. Well-validated AI solutions offer an important approach to expanding access to healthcare.

Third, the cost structure of AI devices differs from that of traditional healthcare. The variable costs of using AI for screening tend to be low, and patients are often expected to pay little or no out-of-pocket costs (Abràmoff et al. [2]). This type of cost structure is particularly attractive because it has economies of scale and lends itself to large-scale deployments for public health interventions, such as disease screening. Singapore, for example, has initiated several national AI-based screening programs for diabetes-related eye disease (Ta et al. [80]). An evaluation of a semiautomated system in Singapore estimated annual savings of $15 million by 2050 (Xie et al. [91]).

1.3. AI-Augmented Healthcare: What Is Old and What Is New

In their seminal book *Prediction Machines: The Simple Economics of Artificial Intelligence*, Agrawal et al. [5] pioneer the view that AI systems can be viewed as "prediction machines" that have greatly reduced the cost of prediction, which is essential in human decision making—consisting of the prediction and judgment phases. This view essentially treats AI as an informational tool to greatly reduce the time and cost of making predictions before making decisions. In other words, when we think about the role of AI in our workflows, the focus should be on how AI generates predictions in an affordable way. In the case of healthcare, the lower cost of generating predictions—for example, about patients' disease states—means that care can be provided earlier in a patient's journey, and even prevent healthy individuals from becoming patients in the first place. In other words, medical AI can be seen as a productivity-enhancing device, as productivity loss is key issue in access and cost concerns (Helmchen et al. [43]).

The field of operations management has extensively studied information flow, particularly in the cases of supply chain management (Gavirneni et al. [34], Lee et al. [56]) and service

operations (Guo and Zipkin [40], Ibrahim et al. [47]). In fact, managing the flow of information is a central issue in the field of operations management, which focuses its mission on matching supply and demand. In this sense, the introduction of AI systems into day-to-day healthcare workflows is not new from a modeling perspective, as it falls within the domain of operations management researchers. However, it expands the scope of operations management in several ways.

First, in contrast to the classic idea of "delayed differentiation" in supply chain management, which has been extensively studied in the operations management literature (Lee and Tang [55], Swaminathan and Tayur [78]), using AI to screen patients at their point of care can be viewed as a way to achieve "*expedited* differentiation." The benefit is not only in separating different types of patients based on their risk levels. It also has implications for patient adherence behavior (Wolf et al. [88]).

Second, the implementation of AI systems has implications for the potential replacement or augmentation of human workers. Consider, for example, the case of LumineticsCore, an AI system that facilitates the screening of diabetic patients in primary care settings, enabling the early diagnosis of retinopathy, one of its explicit goals. Looking at the diagnostic process for diabetic patients through the lens of a supply chain, we identify an upstream entity (the primary care physician) and a downstream entity (the ophthalmologist). The deployment of LumineticsCore can variably increase or decrease the demand for services provided by the downstream entity, depending on the baseline screening rate—currently below 20% (Benoit et al. [12], Channa et al. [19])—and the level of adherence to medical advice prior to the deployment of the AI system. Thus, those downstream in the supply chain may be inclined to support the widespread use of AI in primary care if adherence is particularly high and the baseline rate remains low. This example highlights the power of a supply chain perspective in assessing challenges that arise within the healthcare ecosystem (Betcheva et al. [14], Dai and Tayur [25]).

Third, healthcare is a high-stakes application domain compared with many traditional operations management domains (e.g., retail operations, supply chain management, and revenue management), so it naturally imposes more organizational and operational constraints. The need for security and accountability has also inspired new research questions and can sometimes dictate the choice of AI tools and development approaches (Rudin [70]).

2. Using AI to Increase Clinical Productivity and Expand Access to Healthcare

Although the idea of using AI to improve clinical productivity and access to care is intuitive, the mechanisms underlying such improvements are not well understood. To our knowledge, there has been no formal modeling effort to help estimate such improvements. Such a model is important because it can help generate empirically testable hypotheses and inform policy and management decisions related to the use of AI in everyday healthcare.

We now provide a sketch model to capture various factors that drive clinical productivity in an autonomous AI-enhanced healthcare environment. Our model is motivated by B-PRODUCTIVE, a recent preregistered randomized controlled trial conducted in Bangladesh (Orbis [65]) using LumineticsCore to screen patients for diabetic retinopathy, and has broader implications for environments with limited healthcare resources.

Consider a system with a potential patient demand rate, denoted by Λ, that is much greater than the system capacity, denoted by μ. To simplify the analysis, we assume that the patient arrival process is a Poisson process, that service time is exponentially distributed, and that there is a single specialist in the system. We denote by s the true disease state of each patient, which can be either positive ($s = 1$) or negative ($s = 0$). The patient population has a prevalence of ρ ($0 < \rho < 1$). Following the rational queueing theory framework (Anand et al. [8], Dai et al. [27], Hassin and Haviv [42]), we consider a queue regulator (who may be, for

example, a hospital administrator) who chooses a proportion (δ) of arriving patients to serve in order to ensure a reasonable level of expected waiting time, denoted by ω. The arrival rate of patients served by the specialist is $\lambda \equiv \delta\Lambda$. The queueing system is an $M/M/1$ queue with an arrival rate of λ and a service rate of μ, that is,

$$\frac{1}{\mu - \lambda} \leq \omega,$$

which gives the maximum system throughput in this system without AI:

$$\lambda_0 = \mu - \frac{1}{\omega}. \tag{1}$$

2.1. Productivity Impact of the AI System: A Baseline Model

Now we introduce an autonomous AI system to triage patients. The AI system generates a binary signal ξ, which can be either p (which stands for "positive") or n (which stands for "negative"), which is informative about the patient's disease state but imperfect. It has a sensitivity of α and a specificity of β, that is,

$$\Pr(\xi = p \mid s = 1) = \alpha \text{ and } \Pr(\xi = n \mid s = 0) = \beta.$$

Not all patients are eligible for AI screening; Reasons why certain patients may not be eligible for AI screening include having symptoms such as known vision loss when should be seen by a specialist (Channa et al. [19]). Those who are not eligible for AI screening are directed to a specialist when there is sufficient capacity. We denote by γ the proportion of patients who are eligible for AI screening.

Suppose the goal of the queuing controller is to maintain a specific expected waiting time, represented by the variable ω. It is possible to allow a significant percentage of patients, denoted as δ', into the system under this constraint. We define λ' as the rate at which all the permitted patients arrive, which can be calculated as $\lambda' = \delta'\Lambda$. The patients who end up seeing the specialists are either those who do not qualify for AI assistance, or those who, despite qualifying for AI assistance, receive a negative outcome from the screening process. The arrival rate of patients seen by the specialist is

$$\{(1-\gamma) + \gamma[\rho\alpha + (1-\rho)(1-\beta)]\} \cdot \lambda'.$$

The queueing system (for the specialist's service) is an $M/M/1$ queue with an arrival rate of λ' and a service rate of μ, that is,

$$\frac{1}{\mu - \{(1-\gamma) + \gamma[\rho\alpha + (1-\rho)(1-\beta)]\} \cdot \lambda'} \leq \omega,$$

which gives the maximum system throughput in this AI-augmented system:

$$\lambda_1 = \frac{\mu - \frac{1}{\omega}}{1 - \gamma\rho(1-\alpha) - \gamma(1-\rho)\beta}. \tag{2}$$

Comparing (1) with (2) gives the productivity increase due to the use of AI:

$$\Delta\lambda_1 = \lambda_1 - \lambda_0 = \left(\mu - \frac{1}{\omega}\right) \cdot \frac{\gamma\rho(1-\alpha) + \gamma(1-\rho)\beta}{1 - \gamma\rho(1-\alpha) - \gamma(1-\rho)\beta}. \tag{3}$$

As a numerical example, consider the case where $\rho = 0.3$, $\gamma = 0.95$, $\alpha = 0.9$, and $\beta = 0.8$. In this case, (3) implies that $\Delta\lambda_1$ is equal to $1.275 \cdot \lambda_0$, indicating a 27.5% increase in system throughput, a primary measure of clinical productivity.

From (3), $\Delta\lambda_1$, the change in arrival rate, is directly proportional to γ (the proportion of patients eligible for AI) and β (the specificity of the AI system), but inversely proportional to α (the sensitivity of the AI system). Furthermore, as long as we can safely assume that $\alpha +$

$\beta > 1$ (suggesting that the AI system performs better than a random classifier), we observe that $\delta\lambda_1$ is also inversely correlated with ρ (disease prevalence). These observations suggest that an autonomous AI-enabled device could provide a greater increase in productivity under the following conditions: high disease prevalence, a larger proportion of patients eligible for AI, and a highly specific but not overly sensitive AI system. However, it is important not to misinterpret these findings to mean that a less sensitive AI-enabled device is universally preferable. In fact, a device with high sensitivity is advantageous because it ensures a low rate of false-negative diagnoses; the safety of the device is closely tied to its sensitivity.

2.2. Patient Heterogeneity

The basic model in Section 2.1 assumes that the service time for the human specialist to interact with each patient follows the same exponential distribution. In reality, patient conditions may have different levels of complexity. As our preliminary analysis from the B-PRODUCTIVE study shows, using AI to triage patients means that the human specialist will serve more complex patients, all else being equal.

To incorporate this aspect, consider a scenario where each patient falls into one of two categories: high complexity or low complexity. We denote by κ the probability that a randomly selected individual from the general population is classified as high complexity. Moreover, we will use ρ_h to denote the prevalence of the disease among high-complexity individuals and ρ_l for its prevalence among low-complexity individuals. Here, $\rho_h > \rho_l$, meaning the disease prevalence is higher among high-complexity individuals; κ, ρ_h, and ρ_l satisfy

$$\kappa \cdot \rho_h + (1-\kappa) \cdot \rho_l = \rho. \tag{4}$$

In addition, the specialist's service rates for patients of high complexity and low complexity are represented as μ_h and μ_l, respectively, so that $\mu_h < \mu_l$ and

$$\kappa \cdot \mu_h + (1-\kappa) \cdot \mu_l = \mu,$$

where μ is the average service rate for a patient randomly selected from the population.

To simplify the analysis, we assume that both types of patients are eligible for AI with the same probability (γ). We also assume the same level of sensitivity and specificity for both types of patients.

The hospital administrator does not have the ability to determine the complexity level of each patient upon arrival. The administrator chooses a ratio δ'' to ensure that the expected waiting time is no longer than ω. A total arrival rate of $\lambda'' = \delta'' \cdot \Lambda$ is first sent to be screened by AI and is next sent to the specialist only if the AI screening result is positive or if they are not eligible for AI.

In this case, the arrival rate of patients seen by the specialist is

$$\{(1-\gamma) + \gamma\kappa[\rho_h\alpha + (1-\rho_h)(1-\beta)] + \gamma(1-\kappa)[\rho_l\alpha + (1-\rho_l)(1-\beta)]\} \cdot \lambda'',$$

which, using (4), can be rewritten as

$$\{(1-\gamma) + \gamma[\rho\alpha + (1-\rho)(1-\beta)]\} \cdot \lambda''.$$

The average service rate, on the other hand, is no longer μ. It can now be expressed as

$$\begin{aligned} \mu'' = (1-\gamma)\mu &+ \gamma \cdot \frac{\kappa[\rho_h\alpha + (1-\rho_h)(1-\beta)]}{\kappa[\rho_h\alpha + (1-\rho_h)(1-\beta)] + (1-\kappa)[\rho_l\alpha + (1-\rho_l)(1-\beta)]} \cdot \mu_h \\ &+ \gamma \cdot \frac{(1-\kappa)[\rho_l\alpha + (1-\rho_l)(1-\beta)]}{\kappa[\rho_h\alpha + (1-\rho_h)(1-\beta)] + (1-\kappa)[\rho_l\alpha + (1-\rho_l)(1-\beta)]} \cdot \mu_l, \end{aligned} \tag{5}$$

which is less than μ because $\rho_h > \rho_l$ and $\mu_h < \mu_l$.

Using the same logic as before, to achieve an average wait time of ω, the hospital administrator will choose a λ'' such that

$$\frac{1}{\mu'' - \{(1-\gamma) + \gamma[\rho\alpha + (1-\rho)(1-\beta)]\} \cdot \lambda''} \leq \omega,$$

which implies λ'' can be at most

$$\lambda_2 = \frac{\mu'' - \frac{1}{\omega}}{(1-\gamma) + \gamma[\rho\alpha + (1-\rho)(1-\beta)]}. \tag{6}$$

Although (6) is similar to (2), it is different because in the numerator, the average service rate μ'' depends on the patient mix. Because of this difference, the productivity improvement

$$\Delta\lambda_2 = \lambda_2 - \lambda_0 = \frac{\mu'' - \frac{1}{\omega}}{(1-\gamma) + \gamma[\rho\alpha + (1-\rho)(1-\beta)]} - \left(\mu - \frac{1}{\omega}\right)$$

is less than $\lambda_1 - \lambda_0$. Contrary to the intuitive implications drawn from our basic model (see Section 2.1), comparative statics can reveal surprising results. For instance, in some cases, $\Delta\lambda_2$ might decrease as γ (the percentage of patients eligible for AI screening) increases. Unlike in the basic model, an increase in γ does not necessarily lead to improved productivity. In fact, as more patients become eligible for AI screening, the average service time can actually increase. This unexpected result is largely due to a greater proportion of high-complexity patients being seen by the specialist. Such results offer insights into the limits of AI's potential to enhance productivity and inform cost effectiveness analysis.

In this section, we outlined a basic model that illustrates the potential of using AI to enhance clinical productivity while recognizing the practical intricacies that extend beyond the simplified confines of our model. The strong tradition of stochastic modeling and simulation provides an opportunity to gain an insightful understanding of the impact of AI on other important facets of healthcare, such as patient safety, accessibility, and equity. More importantly, researchers in the fields of operations research and management science are well positioned to address the complex service design challenges that arise from the integration of AI into clinical practice. For a discussion of service design research opportunities in this context, see Dai and Tayur [26].

3. Stakeholder Buy-in for Medical AI

Despite the potential of AI to improve healthcare productivity, access, and equity (Channa et al. [19], Wolf et al. [88]), it must gain sufficient stakeholder acceptance to be widely used in practice (Abràmoff et al. [1], Dai and Tayur [26]). A variety of factors, including liability concerns (Price et al. [67]), data quality and privacy (Singh et al. [75]), explainability (Lebovitz et al. [52]), standards of development and validation (Lebovitz et al. [53]), algorithmic aversion (Longoni et al. [58]), perceived threats to autonomy (Henry et al. [45]), incentive misalignment (Agrawal et al. [6]), and financial factors, especially reimbursement, may influence whether and how physicians and patients adopt AI. Each of these factors provides a rich context for theoretical development and related empirical, laboratory, and field work. In this section, we draw on a recent working paper by Dai and Singh [24] to discuss the impact of physician buy-in on the use of AI in healthcare, focusing on liability implications when physicians use AI in their medical practices.

Liability concerns have always been a major driver of medical decision making, resulting in an annual cost to the medical liability system (including malpractice insurance premiums, malpractice litigation costs, and defense medicine) that has been estimated at $55 billion per year (Mello et al. [61]). The ever-increasing use of AI in medicine may increase or decrease physicians' malpractice liability when adverse patient outcomes occur (Sullivan and

Schweikart [77]). Accordingly, such liability implications may cause the physician to over- or under-use AI.

To make things more concrete, consider the scenario of treatment planning that was initially discussed by Price et al. [67]. In this particular scenario, physicians have the choice of prescribing either a standard or a nonstandard treatment plan for a patient. A standard treatment plan is defined as one that follows the standard of care. For example, in the treatment of ovarian cancer, the dosage of the chemotherapy drug bevacizumab may vary between a standard plan of 15 mg/kg every three weeks and a nonstandard plan of 75 mg/kg every three weeks. Under the existing liability system, a physician is liable if deviation occurs from the standard of care and harms the patient. Conversely, the physician is protected from liability if the standard of care was followed. This liability system, as illustrated in Table 1, incentivizes physicians to err on the side of a low or standard dosage because this choice reduces their potential liability. As a result, physicians may be inclined to prescribe a standard regimen even when a nonstandard regimen may be more effective.

Now we introduce an AI system into the scenario. The AI system is capable of generating a binary signal that advises the physician to provide either a standard or a nonstandard treatment plan. The AI system is considered "assistive AI" in the sense that it guides the physician in making the final decision but cannot make the decision on its own.[1] Therefore, the physician is the one who makes the final decision about the treatment plan and is solely responsible for potential malpractice liability (Abràmoff et al. [3], Price et al. [67], American Medical Association [7]). Depending on the AI system's recommendation (which can be either standard or nonstandard), the physician's final decision (which, again, can be standard or nonstandard), and the patient outcome (which can be either good or bad), we have a total of eight scenarios (see Table 2).

It is safe to eliminate cases 1, 3, 5, and 7 because the physician is generally not expected to be liable if the patient's outcome is good. We are left with cases 2, 4, 6, and 8. Of these, it is safe to rule out case 2, in which the AI system merely reinforces the standard treatment plan, so the physician is unlikely to be liable for acting in accordance with *both* the AI system and the standard of care. Therefore, we have only three cases (4, 6, and 8) to weigh. We have two reasonable liability schemes to consider.

- First, as per Russell and Norvig ([71], p. 1051), when the AI system demonstrates substantial accuracy, it could potentially *override* the current standard of care. In such a scheme, a physician who consults AI for treatment planning would be held accountable for poor patient outcomes if they deviate from the AI system's recommendation, as illustrated in cases 4 and 8 in Table 2. Conversely, if the physician adheres to the AI system's recommendation (as depicted in case 6 in Table 2), they would not bear the liability for subpar patient outcomes.
- Second, if the Russell-Norvig liability scheme seems excessively radical, a more tempered perspective could be proposed, viewing the AI system as a tool to *enforce* the current standard of care, rather than overriding it (Price et al. [67], Tanenbaum et al. [81]). Under this liability framework, a physician would not be liable for adhering to the standard of care (case 8 in Table 2), but would still be accountable for poor patient outcomes when straying from the standard of care (cases 4 and 6 in Table 2). One can refine this argument by adding

Table 1. Prevailing liability scheme (without the option of using AI).

Physician decision	Patient outcome	Physician liable?
Standard	Good	No
Standard	Bad	No
Nonstandard	Good	No
Nonstandard	Bad	Yes

Table 2. Eight different scenarios when the physician uses AI to assist in treatment planning.

Case number	AI recommendation	Physician decision	Patient outcome
1	Standard	Standard	Good
2	Standard	Standard	Bad
3	Standard	Nonstandard	Good
4	Standard	Nonstandard	Bad
5	Nonstandard	Nonstandard	Good
6	Nonstandard	Nonstandard	Bad
7	Nonstandard	Standard	Good
8	Nonstandard	Standard	Bad

that a physician's liability might be lessened when the AI recommends a nonstandard treatment plan (case 6 in Table 2), whereas it might be heightened when the AI recommends the standard treatment plan (case 4 in Table 2).

Tables 3 and 4 present the two liability schemes discussed previously. Based on each of these liability schemes, a decision analysis model can be developed to analyze the physician's decision about whether and how to use AI. In the U.S. healthcare environment where the physician is compensated on a fee-for-service basis for both in-office drug delivery and the use of AI, Dai and Singh [24] find that under both liability schemes, physicians may use AI for low uncertainty cases and avoid using AI for high uncertainty cases. This finding, consistent with what Price et al. [67] argue, means that physicians may avoid using AI in situations where it can help them generate better information because of the possibility that AI's recommendation may contradict their ultimate decision. In contrast, physicians may use AI in cases with little or no uncertainty, essentially using AI to confirm their own intuition. This result enriches the literature on algorithmic aversion and provides theoretical support for measures to reduce such aversion, such as the introduction of "AI insurance" (Bertsimas and Orfanoudaki [13], Stern et al. [76]).

A key driver of AI adoption is its safety and efficacy. It is often argued that physicians are more likely to use AI when it makes sufficiently accurate predictions. Counterintuitively, Dai and Singh [24] show that as the precision of AI improves, the physician's tendency to avoid AI in high-uncertainty cases may be amplified. This is because as AI becomes more accurate, all else being equal, the physician is more likely to have deviated from the AI recommendation when patient harm occurs. This, in turn, creates a disincentive for the physician to use AI because both liability systems penalize the physician for deviating from AI recommendations, albeit in different ways.

The section has focused on liability concerns connected to the use of medical AI. Other than liability issues, other primary drivers of AI adoption include ease of use, seamless integration into workflows and, most importantly, both direct and indirect reimbursement methods. The direct method involves reimbursement, while the indirect approach relies on quality measures such as Healthcare Effectiveness Data and Information Set (HEDIS) and Risk Adjustment Factor (RAF)/Hierarchical Condition Category (HCC) (see Abràmoff et al. [3]

Table 3. Second liability scheme where the physician consults AI before deciding on the treatment plan.

AI recommendation	Physician decision	Patient outcome	Physician liable?
Standard	Nonstandard	Bad	Yes
Nonstandard	Nonstandard	Bad	No
Nonstandard	Standard	Bad	Yes

Table 4. First liability scheme where the physician consults AI before deciding on the treatment plan.

AI recommendation	Physician decision	Patient outcome	Physician liable?
Standard	Nonstandard	Bad	Yes and more
Nonstandard	Nonstandard	Bad	Yes and less
Nonstandard	Standard	Bad	No

for details). A case in point is the recent demise of Pear Therapeutics. Despite having a clinically validated AI that was proven to benefit patients and had received FDA approval (Waltz [86]), the company lacked an appropriate reimbursement strategy, leading to its eventual collapse (Kellaher [50]).

4. Regulatory and Payment Models for Medical AI

The emergence of AI in the healthcare industry, particularly in the form of medical devices, has spurred significant change. The transformative impact of this technology is evidenced by the 521 AI-enabled medical devices approved for clinical use by the FDA as of July 2022 (FDA [32]), with more expected in the future. Notably, approximately 75% of these devices are focused on a single medical specialty—radiology, followed by cardiovascular disease and hematology. Certain specialties, including dermatology, have yet to see an FDA-approved AI-enabled medical device as of 2023 (Young et al. [93]). A detailed breakdown of FDA-approved AI-enabled medical devices by specialty is shown in Figure 2.

The primary goal of AI in healthcare is to improve patient and population outcomes, reduce costs, and increase health equity and access, necessitating its widespread adoption under a sustainable business model. From an operational perspective, a constellation of stakeholders—including patients and their organizations, providers and their organizations, payers, value-based care organizations, regulators, ethicists, and AI developers—are on the critical path to widespread and sustainable adoption of healthcare AI. It is important to recognize the nuanced interdependencies among their interests, which we will attempt to

Figure 2. (Color online) Breakdown of FDA-approved AI-enabled medical devices by specialty (as of July 2022).

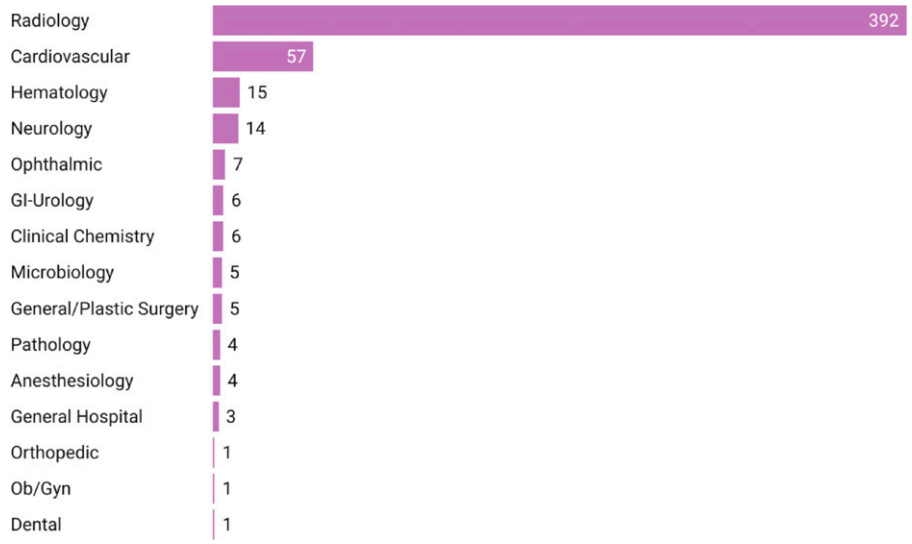

Source. FDA [32].

articulate. The path from conceptualization to widespread deployment of healthcare AI is characterized by a time lag that is the result of various clinical, scientific, and regulatory requirements. Importantly, each stakeholder can influence this time lag.

In this section, we highlight the roles of two key stakeholders whose policies significantly influence the development, validation, and timeline for widespread deployment of AI-enabled medical devices: FDA regulatory policies and payer reimbursement policies.

4.1. Regulatory Models for AI in Medicine

The FDA's regulatory framework significantly influences the adoption and development of AI in the medical sector, particularly through its 510(k) and de novo pathways (Benjamens et al. [11]). This role is underscored by the agency's commitment to ethical principles in AI regulation, a stance echoed by many stakeholders (Abràmoff et al. [3]). Nevertheless, issues such as overfitting, data shift susceptibility, and bias against underrepresented subgroups persist (Wu et al. [89]).

A prevailing debate centers around the "AI flywheel effect" (Gurkan and de Véricourt [41]), which suggests that an AI algorithm will improve with more use, as it generates and learns from more data. Critics argue that current regulatory measures do not fully exploit this effect. To address this, the FDA introduced the Predetermined Change Control Plan (PCCP) and the Precertification Program (FDA [31]) as part of their continual learning policy. The PCCP facilitates the reporting of anticipated modifications to AI devices, while the Precertification Program encourages innovative regulatory oversight approaches for AI-enabled medical devices developed by organizations maintaining high standards of quality and excellence. However, these programs do not fully resolve the complexities involved in evaluating real-world performance, which often prompts requirements not anticipated in advance. According to the principle of "metrics for ethics" (Abràmoff et al. [3]), it is imperative to quantify the impact of specific changes to a validated and authorized AI on safety, efficacy, and equity. This includes understanding the variability in an AI's sensitivity due to changes in its training set (Abràmoff et al. [3]).

Although these initiatives incentivize continual learning and foster AI development, systematic evaluation and stakeholder involvement remain key to their success (Vokinger et al. [85]). Critics voice concerns about potential patient harm due to decreased accuracy and the introduction of undesirable bias (Abràmoff et al. [3], Char et al. [20]). They propose a cautious approach favoring rigorous validation, controlled "episodic" updates, and comprehensive risk assessments. These critics fear that continuous learning toward a potentially flawed physician's determination could inadvertently compromise the accuracy and safety of AI systems (Youssef et al. [94]). This ongoing debate underscores the need for a balanced approach that harnesses the potential of continual learning without compromising patient safety.

Here, we briefly outline two directions for theoretical modeling.

4.1.1. Moral Hazard in AI Development Practice. A key factor in the FDA's approach to authorization is its concern for algorithm quality. Modern AI algorithms based on neural networks are prone to overfitting and unstable performance if best development practices are not followed (Chollet [21]). Thus, under the PCCP program described above, one concern is that the AI developer may not follow best development practices, resulting in overfitted algorithms that are advertised as improved but actually suffer from poor real-world performance. This can be addressed by having a metric for performance improvement based on changes.

On a conceptual level, concerns about the quality of the AI development process are related to the well-known concept of moral hazard in the principal-agent theory literature (Laffont and Martimort [51]). Almost all of the moral hazard literature assumes that the outcome is perfectly and immediately observable. Several recent papers at the intersection of economics

and operations management explore the effect of imperfect outcome visibility, due to, for example, limited inventory (Dai et al. [28]) and a finite queuing buffer (Baiman et al. [10]). However, these papers consider the lack of inventory or queue buffering as the central driver of outcome invisibility. In the case of AI devices, the lack of visibility is often caused by the lack of a sufficiently large and diverse user base. Thus, a low-demand AI device may fall into a low-accuracy trap, where its low accuracy makes it difficult to develop a large and diverse user base; the small and relatively homogeneous user base, in turn, makes it difficult to evaluate its real-world performance and thus improve the algorithm. In addition, even with a large enough user base, it often takes a long time to discover the real-world performance of a device; The lagging nature of results reporting provides another venue for future research. This significant time lag underscores the importance of rigorous premarket validation, using reference standards that prioritize actual outcomes over comparisons with physicians or other experts (Abràmoff et al. [3]).

4.1.2. Market Expansion and Research and Development Funding. The transition from the conceptualization of AI to its widespread, sustainable deployment requires significant funding, typically provided by angel investors and venture capital (VC). The VC investment model, typically organized into a series of funding stages, each with distinct timelines and milestones, has been instrumental in driving growth in the broader technology sector (Cumming and Johan [23], Gilson [38]).

However, the funding model faces unique challenges when applied to AI-enabled medical devices. These challenges stem from the uncertainties surrounding the required "runway" or time lag for these devices to become financially self-sustaining. AI-enabled devices differ from other high-risk sectors, such as pharmaceuticals and traditional medical devices, due to their continuous learning and adaptation (FDA [32]), data network effects (Varian [83]),[2] interaction with human decision makers (Reverberi et al. [69]), and ethical considerations (Char et al. [20], Shachar and Gerke [73]).

The role of healthcare AI investors is to bridge the gap between the conceptualization of an AI system and its widespread adoption. The length of this runway is a critical factor for investors, and all stakeholders can influence it. For example, safety and lack of bias are important considerations for many stakeholders. Strong FDA oversight in demonstrating safety or lack of bias provides assurance to other stakeholders, such as CMS, that certain safety parameters have been met (Centers for Medicare & Medicaid Services [18]).

However, if FDA regulatory scrutiny is low and minimal safety evidence is required, other stakeholders such as CMS, the American Medical Association (AMA), or patient organizations such as the American Diabetes Association may be less comfortable with AI, potentially lengthening the runway. Thus, shortening the "runway" is critical to the success of AI in healthcare, but it can only be achieved by carefully navigating the critical path. Reducing the scientific, regulatory, or clinical requirements for one stakeholder without understanding the interdependencies with all other stakeholders will not significantly impact the overall duration of the "runway."

The uncertainty in AI funding is primarily related to the unknown time lag. If the duration of this time lag were known, the net present value of an investment aimed at bridging this gap for healthcare AI could be accurately analyzed and calculated. However, uncertainty hinders such financial analysis. A shorter time lag would require less funding to bridge the gap and would also yield a higher internal rate of return on the funds invested.

Despite these challenges, calls for a lower regulatory burden for AI should be interpreted within this framework of critical path operational research. Furthermore, even if the burden of FDA regulation were reduced, continued uncertainties about the value, safety, and desirability of AI from a payer perspective could potentially delay or reduce reimbursement rates, thereby impeding the achievement of widespread, sustainable adoption in a meaningful way.

4.2. Reimbursement Models for AI in Medicine

Despite the ongoing calls for reform in healthcare payment methods, the U.S. healthcare industry largely operates on the fee-for-service model (Burns and Pauly [15]). Critics of this model argue that it potentially fuels higher healthcare costs and perpetuates a system in which healthcare providers are incentivized to conduct additional procedures, irrespective of their necessity or appropriateness, as determined by stakeholders other than patients and providers. Conversely, proponents of this model suggest that it plays a crucial role in maintaining accessibility, as it allows both the patient and provider to decide on the provision of a specific service, independent of other stakeholders. However, as an increasing number of healthcare services are funded by entities other than patients, these stakeholders are demanding a say in the care decisions traditionally made between the provider and patient, a context within which the fee-for-service system originally evolved. This external influence often manifests as "value-based care," where other stakeholders, including payers, regulators, ethicists, and others, insist that reimbursement should be optimized based on their perception of "value." This value could be defined in terms of improved patient outcomes, health equity, overall expenditures, and quality-adjusted life years, among other metrics. The challenge for policymakers and healthcare leaders lies in striking a balance between cost containment and the pursuit of quality within the U.S. healthcare system.

Two primary approaches currently exist for reimbursing physicians who use AI. In the case of LumineticsCore, the AI creator established a charge using an equity-based, per-use framework for AI reimbursement (Abràmoff et al. [2]). This AI system was the first FDA-cleared autonomous diagnostic device. By using the standard mechanisms of the CPT Editorial Panel and CMS for determining appropriate reimbursement based on charges sent to billing physicians, CMS was able to establish, for the first time, national reimbursement for autonomous AI (Abràmoff et al. [1], Centers for Medicare & Medicaid Services [18]). Starting in May 2019, the CPT code 92229, specifically created for autonomous AI, allowed for per-use reimbursement (Abràmoff et al. [1]). This model, which aligns charges and reimbursements on a per-patient basis, invoices the charges per use. As a result, the physician using the autonomous AI, typically in primary care clinics, is both charged and reimbursed per patient, eliminating the need for upfront investment. This allows even smaller clinics to use this form of autonomous AI.

An alternative reimbursement model for medical AI involves the new technology add-on payments (NTAP) within the inpatient prospective payment system (IPPS). This method has been used for the reimbursement of several FDA-approved AI systems, such as Viz.ai's Viz LVO. This system analyzes CT images of the brain and alerts hospital staff when a suspected large-vessel occlusion (LVO) is identified. However, this reimbursement approach is only valid for four years, casting doubt on its long-term sustainability given the likely ongoing costs of the service. Other AI systems reimbursed using the NTAP model include the following (Parikh and Helmchen [66]):

- Rapid AI's Rapid LVO, an AI-guided medical imaging acquisition system designed to help acquire cardiac ultrasound images
- PROCEPT BioRobotics Corporation's AQUABEAM system, an autonomous tissue removal robot for the treatment of benign prostatic hyperplasia–related symptoms of the lower urinary tract
- Caption Health's Caption Guidance, which is a medical imaging acquisition system that uses AI to assist in the acquisition of cardiac ultrasound images
- Rapid AI's Rapid Aspects, a computer-assisted diagnostic device designed to characterize abnormalities in brain tissue on brain computed tomography images
- AIDoc's Briefcase for PE, a radiological computer-assisted triage and notification software that detects pulmonary embolism

The NTAP model differs from the per-use model in that it does not reimburse providers for the use of the AI system per se. Rather, an additional payment is made only when the AI system provides "substantial clinical benefit," such as early detection of medical conditions (Murray et al. [63]). In the case of Viz LVO, for example, the NTAP adds a maximum hospital payment of $1,040 for managing a stroke episode. Although this model may seem well aligned with the philosophy of reimbursing a procedure based on the value it creates according to the payer, its real-world impact is more or less mixed: Viz.ai has marketed the device as a way to detect missed stroke cases and generate additional revenue for hospitals, estimating that for a health system with six spoke centers, each with an estimated 31 untreated stroke cases per year and a reimbursement of $28,000 per case, use of the device can generate a total of $5.2 million in "potential new revenue opportunity" (Vizai [84]).

The NTAP model, as Abràmoff et al. [1] note, has some major limitations. First, NTAP only covers inpatient services under the current IPPS, leaving outpatient and physician services outside its scope. Second, NTAP's time limit is another significant limitation. The add-on payments are only available for three years for a given technology or indication. Given the significant resources required to develop, acquire, and implement technologies, this time limit may discourage long-term investments. Third, the financial structure of NTAP is constraining. The maximum add-on payment is limited to 65% of the incremental cost of the new technology. The hospital's costs for a particular stay must exceed the bundled payment amount under the Medicare severity diagnosis-related group (MS-DRG) system for this payment to be made. Therefore, to qualify for the add-on payment, hospitals must demonstrate that they have suffered a financial loss. Finally, many technologies may find it difficult to meet NTAP's stringent eligibility requirements. A new technology must be novel, cost-effective, and demonstrate substantial clinical improvement over existing technologies. Although achieving FDA breakthrough status may satisfy the newness and clinical improvement requirements, newness is a temporary state that exists only until utilization data are available.

To our best knowledge, even with a recent stream of medical literature that provides the guiding framework for reimbursement design (Abràmoff et al. [1], Parikh and Helmchen [66]), no formal research has been conducted into the optimal design of payment schemes for AI-based medical devices (Dai and Tayur [26]). Among many venues for future research, one entails understanding how the use of AI impacts physicians' attention or effort toward individual cases. A paper somewhat related to this idea is by Adida and Dai [4], who study the impact of the physician payment scheme on a physician's diagnostic effort and testing decisions. They present a model in which a patient visits a physician to seek a diagnosis for a given medical condition. The physician estimates the prior probability that the patient is suffering from a serious condition before deciding the level of effort (high or low). The physician may also perform a confirmatory test, which is a diagnostic test to confirm a serious condition. The diagnostic test, which may be a conventional one or AI enabled, is confirmatory because the physician cannot diagnose a serious condition without performing the test. The model assumes that the baseline examination provides an unbiased prior of the patient's condition, and if the physician exerts high diagnostic effort, a private signal informative of the patient's true state will be generated. The tradeoff here is rather subtle: High diagnostic effort can either reduce the need for testing (because the effort may generate a signal that leads to a lower updated likelihood of disease) or increase the propensity to test (because the effort may generate a signal that increases the likelihood of disease and thus necessitates testing). Thus, the relationship between diagnostic effort and testing can be either complementary or substitutive. Under a fee-for-service payment scheme, the relationship depends on the revenue generated from testing. Adida and Dai [4] conclude that there does not exist a fee-for-service payment level that induces the socially optimal amount of effort and testing. Instead, they show a promising payment model entails providing the physician with an additional payment for testing only if the test produces a positive result. This payment model has the same spirit as the NTAP model, but it is for physician payment (not hospital payment). This distinction

is important because physicians are the ones who ultimately determine the allocation of healthcare resources (Dranove and Satterthwaite [29]). Ultimately, the validity of these payment models hinges on who defines the "value" of a service or even a health outcome. In this context, it is the external experts, rather than the provider or patient, who make these determinations.

5. Conclusions

AI has come a long way since its inception in the 1950 s, and in every incarnation of AI, it has had deep and growing ties to medicine. In its current boom, which began in 2011, AI is increasingly being integrated into healthcare workflows in the tangible form of medical devices. The integration of AI into healthcare is still in its early stages, but it has the potential to revolutionize the way medicine is practiced. AI-powered medical devices can help healthcare workers make more accurate and timely diagnoses, prescribe more effective and personalized treatments, and operationalize public health measures. Purposefully designed AI-augmented healthcare systems will help improve the productivity, access, and equity of healthcare delivery.

The integration of AI into healthcare workflows presents a number of challenges that need to be addressed. Key among these is the need for sustained support from all stakeholders, which is complicated by the interdependencies between them. An improvement that appears beneficial to one stakeholder may actually result in less support from another. Ongoing research has identified a number of barriers to the adoption and use of AI. These include liability (Price et al. [67]), data quality and privacy (Singh et al. [75]), explainability (Lebovitz et al. [52]), standards for development and validation (Lebovitz et al. [53]), algorithmic aversion (Longoni et al. [58]), perceived threats to autonomy (Henry et al. [45]), and misaligned incentives (Agrawal et al. [6]). However, some argue that these barriers can be mitigated in specific cases, as evidenced by the widespread stakeholder support for LumineticsCore and similar autonomous AI systems (Abràmoff et al. [1]). More broadly, these barriers need to be thoroughly examined through theoretical, empirical, laboratory, and field work.

To ensure safety and efficacy, AI-enabled medical devices must undergo the same rigorous approval process as other medical devices. This process can be time consuming and costly. Because these devices have the potential to affect the health and well-being of patients, the regulatory approval process for AI-enabled medical devices requires a high level of scrutiny. This requires extensive testing and validation to ensure that the devices meet regulatory requirements for safety and effectiveness. However, because of the unique characteristics of AI-based medical devices, the traditional regulatory approval process may not be readily applicable. A key role is expected to be played by a better quantification of the impact on safety, efficacy, and equity of a previously validated AI. The FDA has proposed an action plan that includes the development of new regulatory approaches tailored to the specific characteristics of AI-based medical devices (FDA [31]). Accordingly, sophisticated modeling efforts are in order for performance analysis and validation of the new regulatory approaches in view of the interests of multiple stakeholders.

Another significant challenge to integrating AI into healthcare workflows is establishing a sustainable payment model, particularly for AI that positively impacts patient outcomes, costs, and health equity. Currently, there are two primary methods of reimbursing physicians for the use of AI: the equity-based model and the NTAP. The equity-based model prioritizes access and health equity, ensuring that beneficial AI technologies are widely available. The NTAP model, on the other hand, reimburses providers only when the AI system delivers "significant clinical value." Despite the proliferation of these models, there appears to be a lack of formal research on the optimal design of payment models for AI-enabled medical devices. This knowledge gap indicates the need for further research to establish more effective payment structures that facilitate the effective use of AI in healthcare.

This tutorial focuses on medical AI, which has the potential to help clinicians screen, diagnose, and treat patients. AI can be used to automate tasks such as note taking, patient messaging, and integration across multiple data collection systems. The rapid development of generative AI has also led to the possibility of generative AI systems trained on a medical knowledge base performing some reasoning activities (Moor et al. [62]). It is also possible to use federated learning (Rajpurkar et al. [68]), a method for training machine learning models on distributed data across multiple devices or servers without directly sharing the data, thereby preserving privacy, and federated causal inference (Xiong et al. [92]), an emerging technique similar to federated learning but focused on causal inference, to improve AI algorithms across multiple locations and data sources.

The integration of AI into healthcare workflows ushers in a new era of possibilities, where the convergence of technology and compassion has the potential to change the very nature of healthcare (Topol [82]). It is time for the operations research and management science community to chart a course for a future in which artificial intelligence serves as a beacon of hope and healing, transforming the medical landscape as we know it.

Acknowledgments

The authors thank Hari Balasubramanian and Ebru Bish for the invitation to write this tutorial and the two external reviewers, Douglas Shier, and the editorial team for their invaluable feedback. This tutorial is the fruit of several past and ongoing research initiatives in collaboration with Elodie Adida, Roomasa Channa, Nathan Congdon, Jiatao Ding, Xin Fang, Michael Freeman, Shubhranshu Singh, Simrita Singh, Leon Xu, Haiyang Yang, Noelle Whitestone, and Risa Wolf. T. Dai benefited from comments and suggestions from seminar participants at Cambridge University, Judge Business School; Emory University, Goizueta Business School; The George Washington University, School of Business; Harvard University, Harvard Business School; Institute for Mathematical and Statistical Innovation; Johns Hopkins University, Bloomberg School of Public Health; Johns Hopkins University, School of Medicine, Department of Plastic and Reconstructive Surgery; The Ohio State University, Fisher College of Business; Singapore Management University; University of California, Los Angeles, Anderson School of Management; University of California, Riverside, School of Business; University of Kansas, School of Business; University of Nebraska-Lincoln, College of Business; Neeley Analytics Initiative; University of North Carolina at Chapel Hill, Kenan-Flagler Business School; University of Pennsylvania, The Wharton School; University of Southern California, Marshall School of Business; and University of Toronto, Rotman School of Management.

Endnotes

[1]The discussion in this section primarily pertains to assistive AI systems. However, it is important to note that when dealing with autonomous AI systems, the majority of the liability tends to fall on the creator of the AI (Abràmoff et al. [1]).

[2]The "data network effect" refers to the effect that algorithm quality may increase with the number of users (Varian [83]). If an AI developer succeeds in further expanding its user base and leveraging the data network effect, a virtuous cycle may follow, often referred to as the "AI flywheel effect" (Gurkan and de Véricourt [41]).

References

[1] M. D. Abràmoff, P. T. Lavin, M. Birch, N. Shah, and J. C. Folk. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. *npj Digital Medicine* 1(39): 2018.

[2] M. D. Abràmoff, C. Roehrenbeck, S. Trujillo, J. Goldstein, A. S. Graves, M. X. Repka, and E. Silva. A reimbursement framework for artificial intelligence in healthcare. *npj Digital Medicine* 5(1): 2022.

[3] M. D. Abràmoff, B. Cunningham, B. Patel, M. B. Eydelman, T. Leng, T. Sakamoto, B. Blodi, et al. Foundational considerations for artificial intelligence using ophthalmic images. *Ophthalmology* 129(2):e14–e32, 2022.

[4] E. Adida and T. Dai. Impact of physician payment scheme on diagnostic effort and testing. *Management Science.* Forthcoming, 2023.

[5] A. Agrawal, J. Gans, and A. Goldfarb. *Prediction Machines: The Simple Economics of Artificial Intelligence.* Harvard Business Review Press, Boston, 2018.

[6] A. Agrawal, J. Gans, and A. Goldfarb. *Power and Prediction.* Harvard Business Review Press, Boston, 2022.

[7] American Medical Association. Policy summary. augmented intelligence in healthcare. Accessed May 28, 2023, https://www.ama-assn.org/system/les/2019-08/ai-2018-board-policy-summary.pdf, 2019.

[8] K. S. Anand, M. F. Paç, and S. Veeraraghavan. Quality–speed conundrum: Trade-offs in customer-intensive services. *Management Science* 57(1):40–56, 2011.

[9] J. W. Ayers, A. Poliak, M. Dredze, E. C. Leas, Z. Zhu, J. B. Kelley, D. J. Faix, et al. Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA Internal Medicine.* 183(6):589–596, 2023.

[10] S. Baiman, S. Netessine, and R. Saouma. Informativeness, incentive compensation, and the choice of inventory buffer. *Accounting Review* 85(6):1839–1860, 2010.

[11] S. Benjamens, P. Dhunnoo, and B. Meskó. The state of artificial intelligence-based FDA-approved medical devices and algorithms: An online database. *npj Digital Medicine* 3(1), 2020.

[12] S. R. Benoit, B. Swenor, L. S. Geiss, E. W. Gregg, and J. B. Saaddine. Eye care utilization among insured people with diabetes in the U.S., 2010–2014. *Diabetes Care* 42(3):427–433, 2019.

[13] D. Bertsimas and A. Orfanoudaki. Pricing algorithmic insurance. Working paper, Massachusetts Institute of Technology, Cambridge, MA, 2023.

[14] L. Betcheva, F. Erhun, and H. Jiang. OM forum—Supply chain thinking in healthcare: Lessons and outlooks. *Manufacturing & Service Operations Management* 23(6):1333–1353, 2021.

[15] L. R. Burns and M. V. Pauly. Transformation of the healthcare industry: Curb your enthusiasm? *The Milbank Quarterly* 96(1):57–109, 2018.

[16] R. Butkus, K. Rapp, T. G. Cooney, and L. S. Engel. Envisioning a better U.S. healthcare system for all: Reducing barriers to care and addressing social determinants of health. *Annals of Internal Medicine* 172(2, suppl):S50, 2020.

[17] CDC. Diabetes and vision loss. Accessed April 16, 2023, https://www.cdc.gov/diabetes/managing/diabetes-vision-loss.html, 2022.

[18] Centers for Medicare & Medicaid Services. Proposal to establish values for remote retinal imaging (CPT code 92229). Policy proposal, Centers for Medicare & Medicaid Services, Washington, DC, 2021.

[19] R. Channa, R. M. Wolf, M. D. Abràmoff, and H. P. Lehmann. Effectiveness of artificial intelligence screening in preventing vision loss from diabetes: A policy model. *npj Digital Medicine* 6(53): 2023.

[20] D. Char, M. D. Abràmoff, and C. Feudtner. Identifying ethical considerations for machine learning healthcare applications. *American Journal of Bioethics* 20(11):7–17, 2020.

[21] F. Chollet. *Deep Learning with Python*, 2nd ed. Manning Publications, New York, 2021.

[22] A. Coronato, M. Naeem, G. D. Pietro, and G. Paragliola. Reinforcement learning for intelligent healthcare applications: A survey. *Artificial Intelligence in Medicine* 109:101964, 2020.

[23] D. J. Cumming and S. A. Johan. *Venture Capital and Private Equity Contracting: An International Perspective.* Academic Press, San Diego, 2013.

[24] T. Dai and S. Singh. Artificial intelligence on call: The physician's decision of whether to use AI in clinical practice. Working paper, Johns Hopkins University, Baltimore, MD, 2023.

[25] T. Dai and S. Tayur. OM forum—Healthcare operations management: A snapshot of emerging research. *Manufacturing & Service Operations Management* 22(5):869–887, 2020.

[26] T. Dai and S. Tayur. Designing AI-augmented healthcare delivery systems for physician buy-in and patient acceptance. *Production and Operations Management* 31(12):4443–4451, 2022.

[27] T. Dai, M. Akan, and S. Tayur. Imaging room and beyond: The underlying economics behind physicians' test-ordering behavior in outpatient services. *Manufacturing & Service Operations Management* 19(1):99–113, 2017.

[28] T. Dai, R. Ke, and C. T. Ryan. Incentive design for operations-marketing multitasking. *Management Science* 67(4):2211–2230, 2021.

[29] D. Dranove and M. A. Satterthwaite. The industrial organization of healthcare markets. Culyer AJ, Newhouse JP, eds. *Handbook of Health Economics*, vol. 1B. Elsevier, Amsterdam, 1093–1139, 2000.

[30] S. E. Eppley, S. L. Mansberger, S. Ramanathan, and E. A. Lowry. Characteristics associated with adherence to annual dilated eye examinations among US patients with diagnosed diabetes. *Ophthalmology* 126(11):1492–1499, 2019.

[31] FDA. Artificial intelligence/machine learning (AI/ML)-based software as a medical device (SaMD) action plan. Accessed April 16, 2023, https://www.fda.gov/media/145022/download, 2021.

[32] FDA. Artificial intelligence and machine learning (AI/ML)-enabled medical devices. Accessed April 16, 2023, https://www.fda.gov/medical-devices/software-medical-device-samd/articial-intelligence-and-machine-learning-aiml-enabled-medical-devices, 2022.

[33] M. C. Fu, H. S. Chang, J. Hu, and S. I. Marcus. Google DeepMind's AlphaGo. *OR/MS Today* 43(5):24–29, 2016.

[34] S. Gavirneni, R. Kapuscinski, and S. Tayur. Value of information in capacitated supply chains. *Management Science* 45(1):16–24, 1999.

[35] G. Giacaglia. *Making Things Think: How AI and Deep Learning Power the Products We Use.* Holloway, San Francisco, 2022.

[36] D. M. Gibson. Estimates of the percentage of US adults with diabetes who could be screened for diabetic retinopathy in primary care settings. *JAMA Ophthalmology* 137(4):440, 2019.

[37] T. G. Gill. Early expert systems: Where are they now? *MIS Quarterly* 19(1):51–81, 1995.

[38] R. J. Gilson. Engineering a venture capital market: Lessons from the American experience. *Stanford Law Review* 55:1067–1103, 2003.

[39] E. O. Graue-Hernandez, D. Rivera-De-La-Parra, S. Hernandez-Jimenez, C. A. Aguilar-Salinas, D. Kershenobich-Stalnikowitz, and A. Jimenez-Corona. Prevalence and associated risk factors of diabetic retinopathy and macular oedema in patients recently diagnosed with type 2 diabetes. *BMJ Open Ophthalmology* 5(1):e000304, 2020.

[40] P. Guo and P. Zipkin. Analysis and comparison of queues with different levels of delay information. *Management Science* 53(6):962–970, 2007.

[41] H. Gurkan and F. de Véricourt. Contracting, pricing, and data collection under the AI flywheel effect. *Management Science* 68(12):8791–8808, 2022.

[42] R. Hassin and M. Haviv. *To Queue or Not to Queue.* International Series in Operations Research & Management Science. Springer, New York, 2002.

[43] L. A. Helmchen, H. P. Lehmann, and M. D. Abràmoff. Automated detection of retinal disease. *American Journal of Managed Care* 20(11), 2014.

[44] J. Hendler. Avoiding another AI winter. *IEEE Intelligent Systems* 23(02):2–4, 2008.

[45] K. E. Henry, R. Kornfield, A. Sridharan, R. C. Linton, C. Groh, T. Wang, A. Wu, et al. Human–machine teaming is key to AI adoption: Clinicians' experiences with a deployed machine learning system. *npj Digital Medicine* 5(1), 2022.

[46] HHS. Diabetes and African Americans. Accessed June 13, 2023, https://minorityhealth.hhs.gov/omh/browse.aspx?lvl=4 &lvlid=18, 2023.

[47] R. Ibrahim, M. Armony, and A. Bassamboo. Does the past predict the future? The case of delay announcements in service systems. *Management Science* 63(6):1762–1780, 2017.

[48] Insurance Institute for Highway Safety. IIHS creates safeguard ratings for partial automation. Accessed May 28, 2023, https://www.iihs.org/news/detail/iihs-creates-safeguard-ratings-for-partial-automation, 2022.

[49] I. Kavasidis, F. P. Salanitri, S. Palazzo, and C. Spampinato. History of AI in clinical medicine. Byrne MF, Parsa N, Greenhill AT, Chahal D, Ahmad O, Bagci U, eds. *AI in Clinical Medicine.* John Wiley & Sons, Hoboken, NJ, 39–48, 2023.

[50] C. Kellaher. Pear Therapeutics files for bankruptcy, to pursue sale. *Wall Street Journal* (April 2023). Accessed June 13, 2023, https://www.wsj.com/articles/pear-therapeutics-les-for-bankruptcy-to-pursue-sale-e16dd591, 2023.

[51] J. J. Laffont and D. Martimort. *The Theory of Incentives: The Principal-Agent Model.* Princeton University Press, Princeton, NJ, 2002.

[52] S. Lebovitz, H. Lifshitz-Assaf, and N. Levina. To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science* 33(1):126–148, 2022.

[53] S. Lebovitz, H. Lifshitz-Assaf, and N. Levina. The no. 1 question to ask when evaluating AI tools. *MIT Sloan Management Review* 64(3):27–30, 2023.

[54] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature* 521(7553):436–444, 2015.

[55] H. L. Lee and C. S. Tang. Modelling the costs and benefits of delayed product differentiation. *Management Science* 43(1):40–53, 1997.

[56] H. L. Lee, K. C. So, and C. S. Tang. The value of information sharing in a two-level supply chain. *Management Science* 46(5):626–643, 2000.

[57] A. Leong, J. Wang, R. Wolf, R. Channa, M. D. Abràmoff, H. Lehmann, and T. Y. A. Liu. Autonomous artificial intelligence increases health equity for patients who are more at risk for poor visual outcomes due to diabetic eye disease. *Proceedings of the 15th Annual Medical Student Research Symposium of the Johns Hopkins University School of Medicine.* Accessed June 13, 2023. https://msrs.som.jhmi.edu/media/DL/2023/Ariel_Leong_CR_MSRS.pdf, 2023.

[58] C. Longoni, A. Bonezzi, and C. K. Morewedge. Resistance to medical artificial intelligence. *Journal of Consumer Research* 46(4):629–650, 2019.

[59] H. P. Lundsgaarde. Evaluating medical expert systems. *Social Science & Medicine* 24(10):805–819, 1987.

[60] W. S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics* 5:115–133, 1943.

[61] M. M. Mello, A. Chandra, A. A. Gawande, and D. M. Studdert. National costs of the medical liability system. *Health Affairs* 29(9):1569–1577, 2010.

[62] M. Moor, O. Banerjee, Z. S. H. Abad, H. M. Krumholz, J. Leskovec, E. J. Topol, and P. Rajpurkar. Foundation models for generalist medical artificial intelligence. *Nature* 616(7956):259–265, 2023.

[63] N. M. Murray, P. Phan, G. Hager, A. Menard, D. Chin, A. Liu, and F. K. Hui. Insurance payment for artificial intelligence technology: Methods used by a stroke artificial intelligence system and strategies to qualify for the new technology add-on payment. *The Neuroradiology Journal* 35(3):284–289, 2022.

[64] National Eye Institute. People with diabetes can prevent vision loss. Accessed April 16, 2023, https://www.nei.nih.gov/sites/default/les/2019-06/diabetes-prevent-vision-loss.pdf, 2019.

[65] Orbis. Bangladesh productivity in eyecare trial. Accessed April 16, 2023, https://clinicaltrials.gov/ct2/show/NCT05182580, 2023.

[66] R. B. Parikh and L. A. Helmchen. Paying for artificial intelligence in medicine. *npj Digital Medicine* 5(1):63 2022.

[67] W. N. Price, S. Gerke, and I. G. Cohen. Potential liability for physicians using artificial intelligence. *JAMA* 322(18):1765 2019.

[68] P. Rajpurkar, E. Chen, O. Banerjee, and E. J. Topol. AI in health and medicine. *Nature Medicine* 28(1):31–38, 2022.

[69] Reverberi C, Rigon T, Solari A, Hassan C, Cherubini P, Antonelli G, Awadie H, et al. Experimental evidence of effective human–AI collaboration in medical decision-making. *Scientific Reports* 12(14952): 2022.

[70] C. Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1(5):206–215, 2019.

[71] S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*, 3rd ed. Prentice Hall Series in Artificial Intelligence. Prentice Hall, Upper Saddle River, NJ, 2015.

[72] T. J. Sejnowski. The unreasonable effectiveness of deep learning in artificial intelligence. *Proceedings of the National Academy of Sciences* 117(48):30033–30038, 2020.

[73] C. Shachar and S. Gerke. Prevention of bias and discrimination in clinical practice algorithms. *JAMA* 329(4):283–284, 2023.

[74] H. A. Simon. *Models of My Life.* MIT Press, London, 1996.

[75] R. P. Singh, G. L. Hom, M. D. Abràmoff, J. P. Campbell, and M. F. Chiang. Current challenges and barriers to real-world artificial intelligence adoption for the healthcare system, provider, and the patient. *Translational Vision Science and Technology* 9(2):45, 2020.

[76] A. D. Stern, A. Goldfarb, T. Minssen, and W. N. Price. AI insurance: How liability insurance can drive the responsible adoption of artificial intelligence in healthcare. *NEJM Catalyst* 3(4), 2022.

[77] H. R. Sullivan and S. J. Schweikart. Are current tort liability doctrines adequate for addressing injury caused by AI? *AMA Journal of Ethics* 21(2):160–166, 2019.

[78] J. M. Swaminathan and S. R. Tayur. Managing broader product lines through delayed differentiation using vanilla boxes. *Management Science* 44(12-part-2):S161–S172, 1998.

[79] W. R. Swartout. Rule-based expert systems: The MYCIN experiments of the Stanford heuristic programming project. *Artificial Intelligence* 26(3):364–366, 1985.

[80] A. W. A. Ta, H. L. Goh, C. Ang, L. Y. Koh, K. Poon, and S. M. Miller. Two Singapore public healthcare AI applications for national screening programs and other examples. *Healthcare Science* 1(2):41–57, 2022.

[81] W. A. Tanenbaum, K. Song, and L. A. Malek. *Theories of AI Liability: It's Still About the Human Element.* Reuters, 2022.

[82] E. Topol. *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again.* Basic Books, London, 2019.

[83] H. Varian. The economics of artificial intelligence: An agenda. *Artificial Intelligence, Economics, and Industrial Organization.* University of Chicago Press, Chicago, 399–422, 2019.

[84] Vizai. Viz.ai: Financial opportunity. AccessedApril 16, 2023, https://www.viz.ai/roi, 2023.

[85] K. N. Vokinger, S. Feuerriegel, and A. S. Kesselheim. Continual learning in medical devices: FDA's action plan and beyond. *The Lancet Digital Health* 3(6):e337–e338, 2021.

[86] E. Waltz. Pear approval signals FDA readiness for digital treatments. *Nature Biotechnology* 36(6):481–482, 2018.

[87] H. R. Warner, A. F. Toronto, and L. G. Veasy. Experience with Baye's theorem for computer diagnosis of congenital heart disease. *Annals of the New York Academy of Sciences* 115(2):558–567, 1964.

[88] R. M. Wolf, R. Channa, M. D. Abràmoff, and H. P. Lehmann. Cost-effectiveness of autonomous point-of-care diabetic retinopathy screening for pediatric patients with diabetes. *JAMA Ophthalmology* 138(10):1063, 2020.

[89] E. Wu, K. Wu, R. Daneshjou, D. Ouyang, D. E. Ho, and J. Zou. How medical AI devices are evaluated: Limitations and recommendations from an analysis of FDA approvals. *Nature Medicine* 27(4):582–584, 2021.

[90] T. Wuest, A. Kusiak, T. Dai, and S. R. Tayur. Impact of COVID-19 on manufacturing and supply networks—The case for AI-inspired digital transformation. Working paper, 2020.

[91] Y. Xie, Q. D. Nguyen, H. Hamzah, G. Lim, V. Bellemo, D. V. Gunasekeran, M. Y. T. Yip, et al. Artificial intelligence for teleophthalmology-based diabetic retinopathy screening in a national programme: An economic analysis modelling study. *The Lancet Digital Health* 2(5):e240–e249, 2020.

[92] R. Xiong, A. Koenecke, M. Powell, Z. Shen, J. T. Vogelstein, and S. Athey. Federated causal inference in heterogeneous observational data. Working paper, 2023.

[93] A. T. Young, J. Y. Chen, A. Bhattarcharya, and M. L. Wei. AI in dermatology. Byrne MF, Parsa N, Greenhill AT, Chahal D, Ahmad O, Bagci U, eds. *AI in Clinical Medicine: A Practical Guide for Healthcare Professionals.* John Wiley & Sons, Hoboken, NJ, 165–175, 2023.

[94] A. Youssef, M. D. Abràmoff, and D. S. Char. Is the algorithm good in a bad world, or has it learned to be bad? The ethical challenges of "locked" vs. "continuously learning" and "autonomous" vs. "assistive" AI tools in healthcare. *American Journal of Bioethics* 23(5):43–45, 2023.