Proceedings of the 57th Hawaii International Conference on System Sciences | 2024

The Dynamic Integration of AI - Driven Telemedicine with In-Office Specialty Care

Opher Baron
University of Toronto
Opher.Baron@rotman.utoronto.ca

Fanying Chen
Boston University
fanying@bu.edu

Abraham Seidmann
Boston University
avis@bu.edu

Abstract

Nation-wide physician shortages and the rise of virtual access to healthcare have driven many organizations to seek innovative ways to integrate digital and automated care programs to serve their members better. We work closely with a large HMO and investigate the option of incorporating an independent third-party service that provides select patients rapid access to virtual Dermatological AI Systems. In this paper, we analyze the impact of this innovative hybrid service system on the waiting times for highly-demanded specialty medical care: here, the automated appointment scheduling system dynamically activates the virtual care only when the in-person service lines get overloaded beyond a predetermined threshold. Using extensive appointment field data, we show that our proposed hybrid service delivery policy significantly improves service quality - even when the virtual care channel serves only a tiny proportion of all patients.

Keywords: telemedicine, IT systems solution, medical AI

1. Introduction

Physician shortages and long waiting times are getting worse in most healthcare systems, especially for secondary (specialist) care. Statistics show that the average waiting times for physicians have increased by 8% from 2017 to 2022 in the U.S.; for instance, a new patient must wait 72 days to see a dermatologist in Minneapolis. The situation will not get better due to demographic changes: the U.S. expects a shortage of 21,000 to 77,100 non-primary care specialists in 2034 (IHS-Markit-Ltd., 2021).

Organizations are seeking ways to mitigate this physician shortage by introducing new digital technologies in operation and diagnostic processes. Early research (Menon et al., 2000) found that IT can positively influence service productivity. More recent research, such as Fan et al., 2023, reveals that offline appointments become more popular when physicians offer online consultations. Zhou et al., 2021 find that telehealth attracts more patients to urban hospitals, bringing the challenge for urban providers to ensure the quality of care for their face-to-face patients. When using telemedicine, patients are more likely to leave without being seen if their online physician is delayed (Qin et al., 2022). This suggests that HMOs must be cautious concerning their online channel waiting and delays. Ayabakan et al., 2023 empirically show that telehealth lowers resource utilization by reducing future visits, especially for highly virtualizable diseases such as mental issues, and suggest insurance providers should expand their plans and cover more telehealth providers. Vodrahalli et al., 2023 show by a quality improvement study that a machine learning algorithm can effectively increase the physicians' operation efficiency by reducing 68% in the number of poor-quality images because it identifies poor-quality images and the reason for poor quality.

Bestsennyy et al., 2021 point out that about $250 billion of current U.S. health spending can be virtualized. However, one challenge is integrating telehealth activities into the day-to-day clinical workflow. System-level integration is far more critical than the most-cited "point solutions", such as digital pathology or digital radiology. While these are innovative AI applications, they do not seem to have had much impact in the clinical space. In Agrawal et al., 2022, the authors claim that "point solutions" are

URI: https://hdl.handle.net/10125/107143
978-0-9981331-7-1

rather easy to implement but that their overall economic impact tends to be relatively small. In contrast, "system solutions" that affect the organizations' workflows have the highest potential impact when it comes to the successful implementation of AI in healthcare delivery. However, there are few existing guides to deploying innovative digital health solutions with existing in-office care to improve the quality of specialist services.

Our research presents an innovative "system solution" which aims at reducing excessive waiting times for specialists as well as reducing geographic inequity in service level quality. In doing so, our methodology is integrating in-person specialists' care with AI-driven telemedicine is inspired by an extensive field project with a large healthcare maintenance organization (HMO) covering about 2.5 million members. At the request of that HMO, we looked at their most acute service lines, which were for their dermatologists. The average waiting time for dermatologists in that organization is about 24 business days, sometimes as long as 120 business days. After analyzing more than 10,000 patient cases, we see that the average patient no-show rate is about 20%, positively associated with longer waiting times. Given the shortage of dermatologists, the HMO leadership is looking at incorporating a third-party "subcontractor" which will offer Dermatological Artificial Intelligence (DAI) systems that can provide (for some instances) initial diagnostics based on images sent to them by patients using mobile phones (Nelson et al., 2020, Grinnell et al., 2022, Winkler et al., 2019, Brewer et al., 2013). The HMO central appointment scheduling system uses an algorithm that watches the expected wait time at each clinic. Based on the predetermined value of the queue length, for that clinic, a chat box will pop up with an offer of the TeleSkin services with an expected shorter wait time and no need to get physically to the clinic. If the chatbot finds that the patient case is medically eligible for TeleSkin, the patients are asked to take a picture of the skin lesion, a part of the skin with an abnormal growth or appearance compared to the skin around it. Once the lesion's images are uploaded to the TeleSkin system, an AI software verifies the quality of the uploaded image from the patient's mobile phone (O'Connor et al., 2017, Vodrahalli et al., 2023). And if the image is acceptable, it gets to another AI system that provides an initial diagnosis. That report with the photo is presented to the TeleSkin (subcontracted) dermatologist, which dictates the resulting outcome to the patient, to the patient's EMR, and to the Primary Care Physician of that patient.

The critical difference distinguishing our proposed IT-based mechanism from existing telemedicine offerings is that we bring extra labor into the system through an additional digital channel. Currently, the specialists working at the HMO clinics are facing a highly random demand, and at times they are running at high utilization – leading to unacceptable long waiting times for office appointments. Hence the need to cleverly add some clinical capacity. On the other hand, the critical problem with traditional telemedicine is that it often increases the workload on specialists because telemedicine brings in more patients since it cuts the patients' cost of access to care (Rajan et al., 2019, Baron et al., 2023). When the same specialist who simultaneously serves in-person patients handles these additional telemedicine requests, it only extends the already excessive waiting time for any new appointment.

Our proposed new workflow is designed to increase the overall capacity of the clinical system. The innovative idea that the HMO is experimenting with is using on-demand independent specialists as subcontractors and advanced AI systems, as explained above. The on-demand resources will be delivered by the "subcontractor" and will be handling these extra telemedicine requests through a separate service channel which does not take away from the current on-site clinical capacity. In addition, because the HMO will offer this service selectively to just some patients, it is not expected to generate extra demand.

While this new technology provides various benefits to the current care delivery system, the HMO leadership's main concern is that these virtual visits (with the subcontractor) will cost about twice as much as the current office visits of patients with the HMO's on-staff dermatologists. On the other hand, this subcontracting service will be designed by us to deliver (on demand) additional specialists' services to help mitigate unexpected peak loads. We investigate the most economical way of incorporating this "subcontracting" service option along with the current healthcare services of the HMO.

Our initial theoretical, numerical, and empirical results show that the proposed dynamic flow mechanism, which selectively allocates patients to virtual or office visits, can potentially generate dramatic improvements in the specialists' care delivery timeliness at a relatively low additional operating cost for the HMO.

2. Our model

We propose a Dynamic Flow Diversion (DFD) mechanism to control the dermatology service system. In the proposed system, the main line (ML) of service is

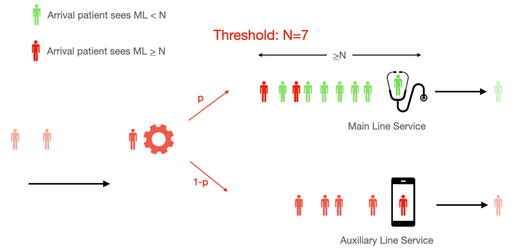

Figure 1. Dynamic flow diversion mechanism.

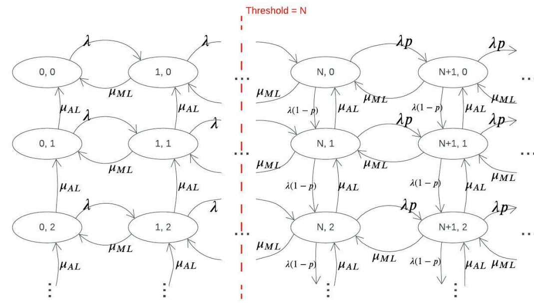

Figure 2. 2D Markov chain.

delivered by the in-office staff on monthly salaries. The on-demand virtual services at the auxiliary line (AL) are delivered by the subcontractor, which uses a mix of AI and telemedicine-driven dermatologists who are paid on a case-by-case basis. Traditionally patients log onto the HMO appointment system and join the ML while waiting to self-schedule themselves for the desired appointment time. At times of peak loads, the DFD algorithm will selectively direct some of the patients to be served by the AL, thus improving the overall service schedule quality for all patients.

We assume the requests for new appointments follow Poisson arrivals at rate λ, and exponential service rates μ_{ML}, μ_{AL} for the ML and the AL, respectively. The ML and the AL are modeled as a single server. The dynamic flow diversion works as follows: once the number of patients in the ML exceeds the pre-set threshold level N, new patients are offered the optional use of the AL. Clearly, not all cases fit the AL; we denote by $(1-p)$ the probability that a new arrival is appropriate (i.e., medically eligible and that the offer for virtual service is accepted by the patient) for the AL. Figure 1 is an example of this DFD service system when N is set equal to seven.

The number of patients in the ML and AL and the birth-death process can be represented by a two-dimensional Markov Chain (2DMC). In Figure 2, each state (i, j) indicates that there are i patients at the ML and j patients at the AL; hence each column i indicates that there are i patients at the ML, and each row j means that there are j patients at the AL. Before ML reaches the pre-set threshold N (the left sub-system), all arrivals join ML and there is no arrow going down each row; after ML exceeds N (the right sub-system), the AL option is activated and new arrivals split for ML and AL following $Ber(p)$. Notice that the number of patients at ML can go beyond N in our setting. The transition rates between the states are as noted in the figure.

This MC is infinite in two-dimensions and asymmetric. We derive some analytical results to categorize the system behavior and analyze the new DFD system's performance, service quality, and system cost. We define $\pi_{i,j}$ as the stationary probability for each state (i, j), $\pi_i^{ML} = \sum_j \pi_{i,j}$ and $\pi_j^{AL} = \sum_i \pi_{i,j}$ are the stationary probabilities that there are i patients at the ML, and that there are j patients at the AL, respectively. It is also useful to know how often the ML overflows N and the system is at the right sub-system that AL needs to be activated. We define the overflow probability $P(N_{ML}^{+})$, and thus $P(N_{ML}^{-}) = 1 - P(N_{ML}^{+})$ represents the fraction of patients that are not offered the AL service. Because not all patients who receive the offer accept it, the *diversion probability* $P(AL)$ is the fraction of patients who are eventually served by the AL, and $P(ML)$ who eventually served by ML. Finally, we use L as the system size, W^Q as the time in queue, and W as the sojourn time for the patients at each service. With the setup, we are now able to describe the system behavior and its economic efficiency.

2.1. ML performance analysis

To understand the ML, we first find its steady state probabilities π_i^{ML}, the overflow probability $P(N_{ML}^{+})$, and the diversion probability $P(AL)$. For now, we have the three following assumptions: 1). every arrival will join either ML or AL, 2). there's no switch between lines once they join, and 3). they wait until they are served. (Later, we will break some of our assumptions to accommodate no shows and balking customers.) With the above three assumptions, we can say that $P(ML) = 1 - P(AL)$.

Theorem 2.1. *I. The ML Steady State Probability: The steady state distribution of the ML system is given by*

$$\pi_0^{ML} = \frac{(1-r_{ML})(1-pr_{ML})}{1-pr_{ML}-r_{ML}^{N+1}+pr_{ML}^{N+1}} \tag{1}$$

$$\pi_i^{ML} = \begin{cases} r_{ML}^{i}\pi_0^{ML}, & \textit{if } i < N \\ r_{ML}^{N}(r_{ML}p)^{i-N}\pi_0^{ML}, & \textit{if } i \geq N \end{cases} \tag{2}$$

.

II. The Overflow Probability: In the DFD system, the fraction of patients that are offered the virtual subcontractor service is given by

$$P(N_{ML}^{+}|N) = \frac{r_{ML}^{N}}{1 - pr_{ML}}\pi_0^{ML} \quad (3)$$

.

Hence, the fraction of patients who never receives the AL offer:

$$P(N_{ML}^{-}|N) = \frac{1 - r_{ML}^{N}}{1 - r_{ML}}\pi_0^{ML} \quad (4)$$

.

III. The Patient Diversion Probability:

$$P(AL) = \frac{r^N - r^{N+1} - pr^N + pr^{N+1}}{1 - pr - r^{N+1} + pr^{N+1}}$$

$$P(ML) = \frac{1 - pr - r^N + pr^N}{1 - pr - r^{N+1} + pr^{N+1}}$$

.

We write $P(N_{ML}^{-})$ and $P(N_{ML}^{+})$ short for the overflow probabilities; N_{ML}^{-} and N_{ML}^{+} represent the situations that ML has less than N patients and ML has greater than or equal to N patients, respectively.

The above important probabilities help us state our main findings that characterize the ML performance. Recall that L_{ML} is the average system size at the ML and W_{ML} is the average sojourn time that a patient spends in the ML (that is, the time he waits to see the doctor plus his service time). Here we present the closed-form expressions for $E[L_{ML}]$, $E[W_{ML}^{Q}]$ and $E[W_{ML}]$.

Theorem 2.2 (ML Performance Measures)**.** *I. The expected number of patients in the ML as a function of N is given by*

$$E[L_{ML}|N] = \pi_0^{ML}\Big[\frac{(N-1)r_{ML}^{N+1} - Nr_{ML}^{N} + r_{ML}}{(r_{ML}-1)^2} - \frac{(N-1)pr_{ML}^{N+1} - Nr_{ML}^{N}}{(pr_{ML}-1)^2}\Big] \quad (5)$$

.

II. The expected system time for patients in the ML of the system as a function of N is given by

$$E[W_{ML}|N] = \pi_0^{ML}\Big[\frac{(N-1)r_{ML}^{N+1} - Nr_{ML}^{N} + r_{ML}}{\lambda(r_{ML}-1)^2} - \frac{(N-1)pr_{ML}^{N+1} - Nr_{ML}^{N}}{\lambda p(pr_{ML}-1)^2}\Big] \quad (6)$$

.

III. Accordingly, the expected waiting time W_{ML}^{Q} is

$$\begin{aligned} E[W_{ML}^{Q}] &= E[W_{ML}] - \frac{1}{\mu_{ML}} \\ &= \pi_0^{ML}\Big[\frac{(N-1)r_{ML}^{N+1} - Nr_{ML}^{N} + r_{ML}}{\lambda(r_{ML}-1)^2} \\ &\quad - \frac{(N-1)pr_{ML}^{N+1} - Nr_{ML}^{N}}{\lambda p(pr_{ML}-1)^2}\Big] - \frac{1}{\mu_{ML}} \end{aligned} \quad (7)$$

.

2.2. AL performance analysis

For the AL system not to be overwhelmed, we need the AL to work no slower than its in-flow rate, which does not depend on the AL. In fact, how fast the AL has to serve is given by

$$\mu_{AL} > \frac{1-p}{1 - pr_{ML}} r_{ML}^{N} \lambda \pi_0^{ML}$$

.

Let π_j^{AL} be the steady-state probability that the AL's system size is j. Since $\pi_{(i,j)}$ is the steady state probability that the ML's system size is i and the AL system size is j, $\pi_j^{AL} = \sum_{i=0}^{\infty} \pi_{(i,j)}$.

Theorem 2.3 (AL Idle Probability)**.** *Suppose partial the traffic intensity for the AL $r_{AL} = \frac{\lambda(1-p)}{\mu_{AL}} < 1$. Then the AL server's idle probability is given by*

$$\pi_0^{AL} = 1 - r_{AL} + r_{AL}P(N_{ML}^{-}) \quad (8)$$

.

2.3. AL approximation

The arrival process to the AL is not a renewal process, making it hard to get closed-form solutions for the AL performance measures. Since the AL is expected to operate at relatively low utilization, we approximate the AL arrival into a Poisson process, whose equivalent arrival rate λ' is the same as that of the actual AL arrival rate spread out on the whole time horizon. This AL approximation allows us to model AL as an M/M/1 system instead of the G/M/1 system and get closed-form expressions for its essential performance measures. We later verified by a simulation that the key performance measures under our AL approximation are close to those of the actual system behavior.

In the approximation, the equivalent normalized arrival rate λ' is λ times the overall diversion probability $P(AL)$: $\lambda' = \lambda\frac{(1-p)(1-r_{ML})r_{ML}^{N}}{1-pr_{ML}-r_{ML}^{N+1}+pr_{ML}^{N+1}}$. The service rate is still μ_{AL} and the normalized traffic intensity $d' = \frac{\lambda'}{\mu_{AL}}$.

Following the average system size, average sojourn time and waiting time for M/M/1 system, we have the theorem:

Theorem 2.4 (Approximated AL Performance Measure)**.** *I. The stationary distribution* π'_{AL}

$$\pi_0'^{AL} = (1 - d') \tag{9}$$

$$\pi_i'^{AL} = (1 - d')d'^i \tag{10}$$

.

II. The average system size is

$$E[L'_{AL}] = \frac{d'}{1 - d'} = \frac{\lambda'}{\mu_{AL} - \lambda'} \tag{11}$$

.

III. The average time with the AL is

$$E[W'_{AL}] = \frac{1}{\mu_{AL} - \lambda'} \tag{12}$$

.

IV. The average waiting time for the AL is

$$E[W'^{Q}_{AL}] = E[W'_{AL}] - \frac{1}{\mu_{AL}} = \frac{\lambda'}{\mu_{AL}(\mu_{AL} - \lambda')} \tag{13}$$

.

3. The economics of the optimal system

The interesting managerial issue is to identify the threshold level N, beyond which eligible patients could be offered the AL virtual service. Below we present a theorem that computes the expected system cost per patient given the parameter set (N, p).

In general, a patient's expected cost has two parts: the amount of time he spends waiting and the service fee that he pays. Notice that the ML servers (for example, the HMO specialists) are employees who receive a fixed amount of full-time salary. We normalize the ML service cost to zero.

$$\begin{aligned} E[cost] =& P(ML)(\text{ML waiting cost} \cdot \text{ML waiting time}) \\ &+ P(AL)(\text{AL waiting cost} \cdot \text{AL waiting time} \\ &\quad + \text{AL service cost}) \end{aligned}$$

Theorem 3.1. *Given N and p, the system's expected per patient cost is*

$$\begin{aligned} E[cost] =& \frac{r_{ML}^N(p-1) - pr_{ML} + 1}{(1 - r_{ML})(1 - pr_{ML})}\pi_0^{ML}. \\ & (w_{ML}\pi_0^{ML}[\frac{(N-1)r_{ML}^{N+1} - Nr_{ML}^N + r_{ML}}{\lambda(r_{ML} - 1)^2} \\ & + \frac{(1-N)pr_{ML}^{N+1} + Nr_{ML}^N}{\lambda p(r_{ML}p - 1)^2}]) \\ & + \frac{r_{ML}^N(1-p) + pr_{ML}^2 + r_{ML}}{(1 - r_{ML})(1 - pr_{ML})}\pi_0^{ML}. \\ & (w_{AL}\frac{1}{\mu_{AL} - \lambda'} + c_{AL}) \end{aligned} \tag{14}$$

.

In the above theorem, $r_{ML} := \frac{\lambda}{\mu_{ML}}$ is the ML's traffic intensity; π_0^{ML} is the ML's idle probability; w_{ML} and w_{AL} respectively are the patient's unit waiting costs in the ML and the AL, which capture the lost goodwill for the HMO due to excessive waiting; c_{AL} is the per-case cost for the AL that is charged by the subcontractor; and, finally, λ' is the approximated arrival rate for the AL. While all our results for the ML are exact, we use fluid flow approximation to study the performance of the AL. The expected cost function in theorem 3.1 shows that the threshold N is a function of p and we can optimize the expected per patient cost by taking derivative with respect to p.

3.1. Numerical exploration

We developed a simulation program based on the basic flow model and the DFD mechanism that allows us to modify the input, the mechanism parameters, and cost parameters and compare the results. We also use a numerical example below to illustrate that a small addition to the AL helps to keep the ML stable.

Example 3.1. *The HMO clinic we have studied can be modeled as a conventional M/M/1 system with no DAI help. According to our data, the patients' arrival rate* $\lambda = 39.95Pt/day$ *and the specialist serves them in that clinic at rate* $\mu = 40Pt/day$*. Then* $E[L_{conv}] = \frac{r_{ML}}{1-r_{ML}} = 799\ pts$*, where* $r_{ML} = \frac{\lambda}{\mu}$ *and* $E[W_{conv}] = \frac{1/\mu}{1-r_{ML}} = 20\ days$*. On average, as we saw in our data, and as the model accurately predicts, there are about 800 patients in this conventional service system and each patient spends on average 20 (working) days waiting to be seen by the specialist. (We exclude weekend waiting days as they distort the measurement due to the "day of the weekend arrival" effect.)*

Now we integrate it with the DAI system using our

DFD mechanism. We add an AL, which serves at the same speed as the ML ($\mu_{AL} = 40Pt/day$), and $c_{AL} = \$75Pt/case$, is the per-case cost that the DAI subcontractor charges the HMO. Suppose only 30% of the patients who get the offer will be eligible and will choose the AL (i.e., the offer refusal probability $p = 0.7$). In this example, we model the patients' waiting cost as $w_{ML} = \$5/day$ for the ML and $w_{AL} = \$40/day$ for the AL. Notice that $w_{AL} >> w_{ML}$ because patients diverted to the AL expect an expedited service. Figure 3 illustrates the HMO's cost as a function of N.

Our theoretical results indicate that the optimal threshold level for cost per patient is at $N = 45$, where $E[L_{ML}] = 24.85$ and $E[W_{ML}] = 0.62\ day = 5\ hr$. The AL operates at a low utilization of 2%. The overall probability that any patient gets the AL offer and accepts it is $P[Join\ AL] = 0.02$. The resulting new average system size $E[L] = 24.36$ and $E[W] = 0.61\ day = 4.8\ hr$. The average waits thus decreases by 97%.

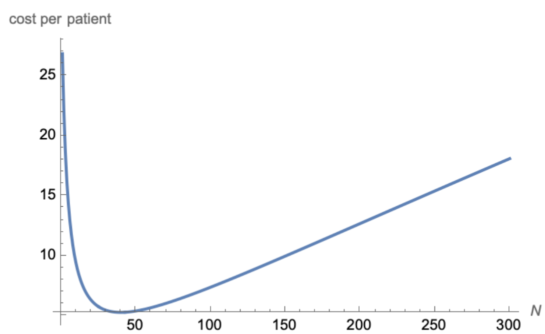

Figure 3. Total Cost as N Changes

Figure 3 shows how the HMO's total cost per period varies with the threshold level, N, using representative data from the HMO. As expected, initially, as N increases, more and more patients are served by the on-staff dermatologist, which is the least-cost option. However, the increased imputed waiting cost for the ML dominates these savings beyond a certain point. Figure 3 shows that the optimal value for the diversion threshold in this example is: $N = 45$ patients. It is important to note that the queue length for the ML can be more than 45 since not all arriving cases are eligible for the virtual dermatology care. Also, as N becomes sufficiently large, almost all patients stay with the ML, and the HMO cost per patient rapidly escalates due to the excessive waiting times.

4. Conclusions and managerial insights

Physician shortages and increasingly long waiting times for specialist (secondary) care have become a significant policy and managerial concern. Political leaders and large healthcare delivery organizations seek creative ways to address this growing public policy issue. Many are looking at the latest developments in virtual care, telemedicine, and medical AI as a way to address these concerns cost-effectively. Yet, one emerging issue is that many medical AI innovations have become straightforward point solutions; they replace one clinical function (manual or automated) with a far more advanced AI solution. On the other hand, as cited earlier, it is the system solutions that also require new or modified service process workflows, are those which generate the most significant returns for AI and telemedicine investments.

We present a novel model of a medical service system that uses Dynamic Flows Diversion to reduce waiting times for outpatient, dermatology specialty care. Our proposed DFD online scheduling model controls the appointments' service system by dynamically activating the AI-driven virtual care service line only when the (in-person) mainline gets overloaded. Our analytical results, along with our field data from a large HMO, show that this proposed mechanism achieves significant service improvements even when AI-driven virtual care serves just a small proportion of patients. Our field data indicates that by optimally using the DFD model, we can dramatically reduce overall waiting times by as much as 97%, even when the (far more expensive) AI-driven virtual dermatology care line handles just 2% of the total volume of patients. The DFD concept, therefore, offers a cost-effective solution to the problem of excessive waits for dermatologists, and perhaps for other specialties as well.

We also find that the optimal value of N (the threshold lever for activating the DFD) is not very sensitive to p, the probability of case diversion, when the queue length exceeds N. However, it is susceptible to the ratio of the cost structure of the in-person, AI services by the subcontractors, and the imputed waiting cost.

Our initial findings already provide valuable insights to clinical organizations seeking to enhance their service level quality with the addition of AI-driven on-demand virtual services in the most cost-efficient way.

References

Agrawal, A., Gans, J., & Goldfarb, A. (2022). From prediction to transformation. *HARVARD BUSINESS REVIEW*, *100*(11-12), 100–+.

Ayabakan, S., Bardhan, I. R., & Zheng, Z. ((2023). Impact of Telehealth and Process Virtualization on Healthcare Utilization [Publisher:

INFORMS]. *Information Systems Research*. https://doi.org/10.1287/isre.2023.1220

Baron, O., Chen, X., & Li, Y. (2023). Omnichannel services: The false premise and operational remedies. *Management Science*, *69*(2), 865–884.

Bestsennyy, O., Gilbert, G., Harris, A., Rost, J., et al. (2021). Telehealth: A quarter-trillion-dollar post-covid-19 reality. *McKinsey & Company*, *9*.

Brewer, A. C., Endly, D. C., Henley, J., Amir, M., Sampson, B. P., Moreau, J. F., & Dellavalle, R. P. (2013). Mobile applications in dermatology. *JAMA dermatology*, *149*(11), 1300–1304.

Fan, W., Zhou, Q., Qiu, L., & Kumar, S. (2023). Should doctors open online consultation services? an empirical investigation of their impact on offline appointments. *Information Systems Research*, *34*(2), 629–651.

Grinnell, M., Concha, J., Feng, R., Keyes, E., Okawa, J., Diaz, D., Vazquez, T., & Werth, V. (2022). 278 photo validation study using cutaneous dermatomyositis disease area and severity index in dermatomyositis patients. *The Journal of Investigative Dermatology*, *142*(8), S48.

IHS-Markit-Ltd. (2021). *The Complexities of Physician Supply and Demand: Projections From 2019 to 2034*. https://www.aamc.org/media/54681/download

Menon, N. M., Lee, B., & Eldenburg, L. (2000). Productivity of Information Systems in the Healthcare Industry [Publisher: INFORMS]. *Information Systems Research*, *11*(1), 83–92. https://doi.org/10.1287/isre.11.1.83.11784

Nelson, C. A., Pérez-Chada, L. M., Creadore, A., Li, S. J., Lo, K., Manjaly, P., Pournamdari, A. B., Tkachenko, E., Barbieri, J. S., Ko, J. M., et al. (2020). Patient perspectives on the use of artificial intelligence for skin cancer screening: A qualitative study. *JAMA dermatology*, *156*(5), 501–512.

O'Connor, D. M., Jew, O. S., Perman, M. J., Castelo-Soccio, L. A., Winston, F. K., & McMahon, P. J. (2017). Diagnostic accuracy of pediatric teledermatology using parent-submitted photographs: A randomized clinical trial. *JAMA dermatology*, *153*(12), 1243–1248.

Qin, J., Chan, C. W., Dong, J., Homma, S., & Ye, S. (2022). *Waiting online versus in-person in outpatient clinics: An empirical study on visit incompletion* (tech. rep.). working paper, Columbia University. URL: http://www.columbia. edu/˜ cc3179 . . .

Rajan, B., Tezcan, T., & Seidmann, A. (2019). Service systems with heterogeneous customers: Investigating the effect of telemedicine on chronic care. *Management Science*, *65*(3), 1236–1267.

Vodrahalli, K., Ko, J., Chiou, A. S., Novoa, R., Abid, A., Phung, M., Yekrang, K., Petrone, P., Zou, J., & Daneshjou, R. (2023). Development and clinical evaluation of an artificial intelligence support tool for improving telemedicine photo quality. *JAMA dermatology*, *159*(5), 496–503.

Winkler, J. K., Fink, C., Toberer, F., Enk, A., Deinlein, T., Hofmann-Wellenhof, R., Thomas, L., Lallas, A., Blum, A., Stolz, W., et al. (2019). Association between surgical skin markings in dermoscopic images and diagnostic performance of a deep learning convolutional neural network for melanoma recognition. *JAMA dermatology*, *155*(10), 1135–1141.

Zhou, M., Li, X., & Burtch, G. (2021). Healthcare across boundaries: Urban-rural differences in the financial and healthcare consequences of telehealth adoption. *Boston University Questrom School of Business Research Paper*, (3807577).