



**TAKEMI PROGRAM  
IN INTERNATIONAL HEALTH**

HARVARD T.H. CHAN SCHOOL OF PUBLIC HEALTH

1983-2023

# The Digital Transformation in Health: How AI can Improve the Performance of Health Systems

África Periañez, Ana Fernandez Del Río, Ivan Nazarov, Enric Jané, Moi Hassan,  
Aditya Rastogi, Dezian Tang

40th Anniversary Takemi Symposium in International Health

**DIGITAL HEALTH: OPPORTUNITIES AND CHALLENGES FOR GLOBAL HEALTH**

**KEYNOTE ADDRESS**

October 21, 2023  
Boston, MA, USA



**HARVARD T.H. CHAN**  
SCHOOL OF PUBLIC HEALTH

**TAKEMI PROGRAM**  
IN INTERNATIONAL HEALTH

# The Digital Transformation in Health: How AI can Improve the Performance of Health Systems

ÁFRICA PERIÁÑEZ, Causal Foundry, Inc,

ANA FERNANDEZ DEL RÍO, Causal Foundry, Inc,

IVAN NAZAROV, Causal Foundry, Inc,

ENRIC JANÉ, Causal Foundry, Inc,

MOIZ HASSAN, Causal Foundry, Inc,

ADITYA RASTOGI, Causal Foundry, Inc,

DEXIAN TANG, Causal Foundry, Inc,

Mobile health tools have the potential to revolutionize healthcare delivery and patient engagement globally. We discuss how the integration of AI into digital health tools—focused on supply chain management, patient engagement, and capacity building, among others use cases—can enable improved performance of the health system and public health. We present an AI platform that allows the delivery of data-driven adaptive interventions whose impact can be optimized through experimentation and real-time monitoring. The system can integrate multiple data sources and digital health applications. The flexibility of this platform to connect to a variety of mobile health applications and send personalized recommendations based on predictions can significantly improve the impact of digital tools in improving health system outcomes.

Additional Key Words and Phrases: digital transformation, global health, adaptive interventions, reinforcement learning, artificial intelligence

---

Authors' addresses: África Perriáñez, [africa@causalfoundry.ai](mailto:africa@causalfoundry.ai), Causal Foundry, Inc, , ; Ana Fernandez del Río, [ana@causalfoundry.ai](mailto:ana@causalfoundry.ai), Causal Foundry, Inc, , ; Ivan Nazarov, [ivan@causalfoundry.ai](mailto:ivan@causalfoundry.ai), Causal Foundry, Inc, , ; Enric Jané, [enric@causalfoundry.ai](mailto:enric@causalfoundry.ai), Causal Foundry, Inc, , ; Moiz Hassan, [moiz@causalfoundry.ai](mailto:moiz@causalfoundry.ai), Causal Foundry, Inc, , ; Aditya Rastogi, [aditya@causalfoundry.ai](mailto:aditya@causalfoundry.ai), Causal Foundry, Inc, , ; Dexian Tang, [dexian@causalfoundry.ai](mailto:dexian@causalfoundry.ai), Causal Foundry, Inc, ,

## 1 INTRODUCTION

Developing countries face multiple challenges in health systems and public health, mostly caused by factors such as constrained resources, inadequate infrastructure, and socio-economic disparities. One major issue is poor access to basic healthcare services, often resulting from a scarcity of healthcare facilities and trained healthcare professionals.<sup>1</sup> Ineffective health information systems hamper the effective collection and utilization of data for decision-making and resource allocation.<sup>2</sup> Health disparities and inequalities are aggravated by socio-economic factors, resulting in the poorest populations having the least access to essential healthcare services and an increased burden of disease.<sup>3</sup> Infectious diseases like HIV/AIDS, tuberculosis, and malaria disproportionately affect developing nations, straining their health systems and diverting scarce resources from other important public health initiatives.<sup>4</sup> Furthermore, the rise of non-communicable diseases presents a dual burden, complicating efforts to enhance public health amidst already constrained circumstances.<sup>5</sup> Despite advancements in healthcare, developed countries also face significant challenges. Pervasive issues related to patient engagement, adherence to treatments, and adoption of public health interventions prevail.<sup>6</sup>

Accelerated by the COVID-19 pandemic, the global healthcare sector is experiencing rapid growth in the use of digital health tools as promising solutions to some persistent challenges encountered by patients, caregivers, and suppliers.<sup>7</sup> As part of the impetus towards universal health coverage, the World Health Organization has recognized health worker decision support and targeted provider and patient communication via mobile applications as part of a set of prioritized digital interventions for health system strengthening.<sup>8</sup> These platforms can support communication, management of patients and chronic conditions, self-reporting, adherence support, access to references to medical knowledge base and pharmaceutical catalogs, capacity building and procurement of drugs and medical supplies, and inventory and distribution management.

In recent years, Artificial Intelligence (AI) has emerged as a powerful tool in healthcare, capable of utilizing extensive data for decision-making and personalized interventions.<sup>9-12</sup> AI systems can appropriately track and analyze data from digital health platforms to offer insights, predictions, and recommendations into patient and provider behaviors, healthcare quality, and demand, facilitating the development of adaptive interventions

delivered through digital tools.<sup>13-17</sup> AI-enhanced adaptive interventions are not only tailored but also continuously evolve, adapting in real-time to each user's unique choices and circumstances, thereby determining the content and timing of interventions, and identifying individuals who may need additional support. Furthermore, utilizing AI to anticipate demand variations for medications and supplies aids in supply management is crucial to enable timely reminders and guidance systems that ensure consistent availability of vital supplies across all distribution channels<sup>18</sup>.

Digital adaptive strategies are predominantly seen in high-income countries, particularly in entertainment and e-commerce.<sup>19,20</sup> However, the rising mobile health infrastructure in LMICs and surging smartphone adoption present a significant potential for positive global health impacts through adaptive measures.<sup>21,22</sup>

This paper presents an AI platform designed to generate personalized predictions, recommendations, and insights by harnessing and analyzing data from existing digital health tools and other relevant sources. The platform provides the capability to create, test, and deploy adaptive interventions driven by reinforcement learning (RL) to enable optimized user engagement, resource allocation recommendations, and clinical and behavioral guidance for users. These capabilities are possible through the platform's integration with existing digital tools, enhancing their operational efficacy for a spectrum of health-focused entities, including patients, medical facilities, pharmacists, Community Health Workers (CHWs), midwives, and distributors of drugs and medical consumables. Intended for a public health audience, this paper deliberately avoids delving into the technicalities and specificities of the AI and ML algorithms and methods utilized by the platform, prioritizing instead a focus on the practical benefits and outcomes relevant to the public health context.

## **2 DIGITAL HEALTH AND AI**

Artificial Intelligence (AI), and in particular, Machine Learning (ML) are becoming essential tools in the modern scientific landscape and are increasingly being used in public health and healthcare. AI has recently received significant attention due to the popularization of large language models (LLMs).

At their core, AI and ML are a set of techniques that enable systems to learn patterns from large sets of data without being explicitly programmed for specific tasks. These tools are, in most cases, designed to make predictions from large amounts of data without necessarily

having a specific focus on understanding the underlying mechanisms driving the system outputs.<sup>23</sup> There are several domains within AI that have been used for different types of problems, such as natural language processing, deep learning, and computer vision. However, AI is a rapidly evolving field with an active research community exploring new applications underpinned by the increasing computational power available nowadays.

Given the enormous amounts of data generated by digitized health systems, harnessing AI and ML to extract meaningful insights and make evidence-based decisions has become one of the frontiers in public health and healthcare.<sup>24–27</sup> Potential applications specifically for global health have also gathered attention in recent years.<sup>12,9</sup>

Predictive analytics, an important component of ML, allows public health professionals to forecast trends in disease and potential outbreaks. For instance, ML algorithms can be trained on historical epidemiological data to predict future disease hotspots or to identify factors that make populations susceptible to certain health challenges. This forward-looking capability can support optimal resource allocation and the timely deployment of public health measures to protect vulnerable populations more efficiently with targeted interventions.<sup>28–30</sup>

Another transformative application lies in the realm of behavioral nudges and other personalized interventions.<sup>31–34</sup> By analyzing individuals' behaviors, habits, and preferences, ML algorithms can craft personalized messages or recommendations designed to encourage the adoption of best practices, healthier choices, and adherence to treatment and public health interventions, among other behaviors.<sup>14,35</sup> These systems dynamically adapt and allow rapid innovation cycles, supporting the deployment of agile and responsive public health and healthcare systems.

Finally, in the face of digital health fragmentation, where systems often operate in silos with limited interoperability, AI systems are capable of integrating diverse data streams and then offering contextualized insights. CHW job aids, procurement and supply chain management apps, electronic health records, and customer-facing tools serving patients and caregivers all generate vast data. When exploited by AI, these data can significantly enhance the effectiveness and efficiency of health systems.

The following subsections expand on some specific ways AI can be leveraged to improve the health system's efficiency and outcomes.

## 2.1 Frontline healthcare worker engagement and support

Frontline healthcare workers (FHCWs), especially in LMICs, grapple with challenges like fragmented programming, lack of continuous performance assessments, and the overwhelming task of managing diverse patient needs with limited resources. In addition, health systems struggle to ensure healthcare workers are well-trained, supervised, and equipped to provide quality care.

Digital tools can support FHCWs by offering guidance through clinical processes and connecting them to their patients and peers. They can also be used for capacity building, continuous professional development, operational tasks such as appointment management, referrals, and quantifying the required tests and drugs to be requested by medical stores.

The integration of AI and ML into digital tools in the hands of FHCWs offers promising avenues to improve their ability to provide quality healthcare.<sup>36</sup> Personalized interventions can support adaptive learning journeys, timely reminders for key tasks, incentives for best practices, and patient-specific recommendations based on clinical and behavioral data, optimizing the performance and impact of healthcare workers. For example, AI-empowered digital tools can use patient data to assist FLWs in the identification of high-risk patients who could benefit from extra support and the planning of visits and targeted communications. This is particularly important in the case of Community Health Workers (CHWs), who play a crucial role in linking formal healthcare systems with local and often remote communities and often suffer from limited tools and skills.

When it comes to improving the skills of FHCWs, often the challenge is ensuring adherence to digitally enabled training programs. ML tools allow to building of adaptive learning journeys, in which user data are utilized to design individualized nudges and content so we ensure they are engaged in the program and able to pass every certification at an adequate speed for everyone.

Finally, AI systems can also support FHCWs by offering patient-specific recommendations regarding clinical workflows (e.g., on additional testing advised for diagnosis or before prescription or referral to a facility). While this paper is focused on operational and behavioral aspects of healthcare and public health, the AI functionalities presented in the following sections would be equally applicable for purely clinical interventions.

## **2.2 Patient and caregiver engagement and support**

Patients and caregivers in LMICs face daunting challenges when seeking healthcare, including timely access to life-saving commodities and medical advice. The use of digital apps can support them by connecting them to relevant health information, healthcare providers, and pharmacies.

These digital tools generate vast amounts of data on patients' health, the information they seek, and other health-related behaviors. For example, AI has the potential to significantly improve the impact of family planning programs. With increasing mobile penetration in these countries, particularly among the youth, AI-driven apps can ensure women are provided with tailored advice on contraceptive choices, potential side effects, and proper usage. Through the integration with pharmacy chains, women can also access contraceptives online, decreasing the barrier to these products related to social norms.

By centering patient needs and offering AI-driven solutions to address challenges, the healthcare landscape in LMICs can be significantly transformed, ensuring that patients receive the robust support they require.

## **2.3 Pharmacies**

On many occasions, pharmacies are the first line of contact for patients seeking treatment. Pharmacists play an integral role in ensuring patients and caregivers get the right medical advice, adhere to treatment, and responsibly use drugs to counter antibiotic resistance.

In many LMICs, one of the key challenges is ensuring pharmacies stay abreast of the guidelines and offer quality care to patients. Digital tools offer an opportunity to provide training and gather data on pharmacy patients. AI can leverage this data to ensure pharmacists receive the right nudges and incentives to follow guidelines and provide the right care, strengthening the role of pharmacies in the overall health system.

Another dimension is the pharmacist as a business operator. Inventory management, forecasting, and procurement optimization are some of the key functions pharmacists struggle with. This involves proactive collaboration with pharmaceutical warehouses to replenish stock while judiciously balancing the demand against the risk of overstocking, which can lead to wastage through expired medications. AI systems can support all these functions, ensuring pharmacies are well-stocked with life-saving commodities.

A noteworthy trend in LMICs is the growth of tech companies specializing in B2B e-commerce distribution to avoid dependence on expensive intermediaries to acquire supplies. Simultaneously, many pharmacies are evolving into franchises and chains. Consequently, e-commerce applications for healthcare products are becoming essential to a cheaper way to purchase goods and also allow the tracking and visibility of the demand and drug stock movement. AI provides the environment to leverage the large amounts of data generated through these digital tools, providing clients with tailored and timely recommendations for the right consumption and helping pharmacists manage their stocks to ensure essential medicines are always available and to increase the variability of drugs.

## 2.4 Supply Chains

Supply chains are a key component of health systems. However, the intricacy of these networks, especially in LMICs, often results in stockouts across the supply chain and at the point of care.<sup>37</sup> The prevalence of counterfeit drugs is also an issue in some markets. Efficient supply chains prevent wastage and counterfeiting, safeguard against stock-outs, and ensure that life-saving commodities, such as antimalarials and antibiotics, are available at the point of care.

Digital tools have the potential to make the supply chain more efficient. They can ease procurement (through B2B solutions described in the previous section) and reduce paperwork, enabling the use of different technologies such as computer vision or RFID to improve inventory management and prevent counterfeits. Software for supply chain processes facilitates the storage of reliable data with integrity, enabling its systematic use for knowledge extraction and intervention.

AI can use real-time data from digital applications, including inventory levels, demand patterns, and other factors like climate events or disease outbreaks, to understand their past evolution and to predict future needs. Fig. 1 illustrates a schema of how an adaptive intervention can improve an undesired healthcare outcome. This promotes timely stock replenishment and aids in detecting abnormal behaviors or potential supply disruptions early. These models empower stakeholders to make informed decisions, optimizing the supply chain from procurement to delivery.<sup>38</sup>

Customers, distributors, and other key actors in the supply chain system can receive tailored support through reminders, product recommendations and prioritization, and






Fig. 1. **Adaptive interventions.** The goal is to dynamically adjust the strategies to the individuals' changing needs and context over time, using data-driven ML methods that analyze previous responses to interventions, and predicted and observed performance, to maximize the overall performance.

order and distribution plan proposals. These interventions can leverage demand and user behavior predictions and make the adaptive framework work to optimize distribution operations and reduce stockouts of essential medications at points of care.

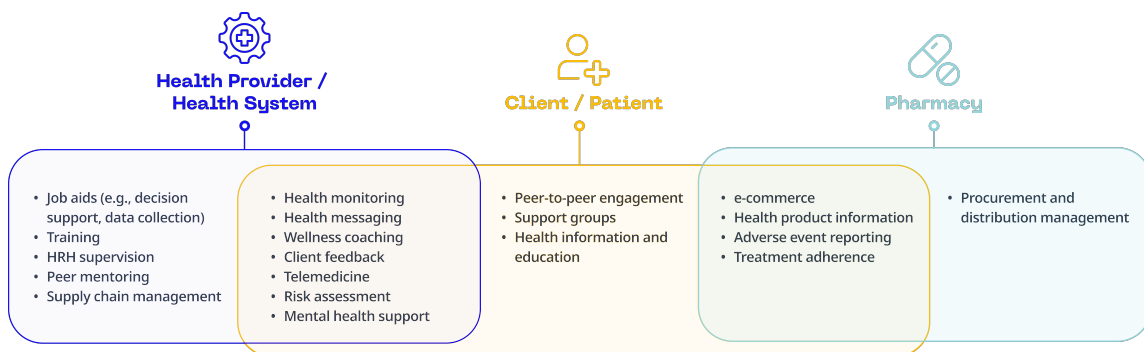


Fig. 2. **Examples of use cases that can benefit from the Reinforcement Learning Platform.** The platform delivers AI-based recommendations and incentives to assist LMIC patients, health systems and health providers (frontline clinicians, nurses, midwives, community health workers, and pharmacists), and pharmacies and drug distributors. It optimizes user engagement and optimization of resources, leveraging the data generated by the ecosystem of applications such as the ones shown in this figure.

### 3 A REINFORCEMENT LEARNING PLATFORM FOR DIGITAL HEALTH

We propose an AI platform designed to support health systems in LMICs (e.g., healthcare facilities, frontline healthcare workers, pharmacies, patients, and supply chains). This platform is the product of our meticulous efforts to ensure the potential of AI can be easily integrated into the growing number of digital health tools to realize the benefits described in the previous section. The result is a platform that allows the use of ML to increase the effectiveness and efficiency of the health system and improve patient outcomes.

Beyond the use of predictive modeling, a unique feature of the platform is that it allows to testing and deploying of adaptive interventions using digital tools. Examples of these interventions include sending messages (SMS, WhatsApp, push notification, in-app messages, etc) to the users to promote behaviors such as adherence to treatment, reading or listening to relevant health information, or taking a training course. By adaptive, we mean that the system tailors the intervention content and timing to each user based on the feedback collected from the ensemble of users on a continuous basis.

As portrayed in Figure 2, the platform allows users to integrate data from different sources, use ML tools to get insights from the data, and create, test, and deploy adaptive interventions using reinforcement learning techniques. These interventions are pushed to the digital tools connected to the platform, and data collected prospectively feeds back into the system, allowing it to learn and adapt the interventions to maximize the intended outcome.

We provide an overview of the platform’s capabilities in the following subsections.

#### 3.1 The platform

The primary aim of our platform is to integrate vast data sets (maximizing their quality, right format, and label), from digital tools, including behavioral, clinical, and contextual data, to provide valuable insights to stakeholders and create, test, and deploy personalized interventions such as nudges, recommendations, and in-app bonuses. The platform ensures that acquired data are well organized and ready for ML predictive models. This emphasis on data collection and the use of data schemas is at the core of the platform functionality.

The digital platform is built upon three foundational components (see Figure 3): the frontend, which provides a comprehensive user interface with data and results visualization; the backend, which serves as the operational core; and the SDK, acting as the crucial bridge between the digital tools and the platform. Each component plays a pivotal role, ensuring users can seamlessly access, analyze, and manage their data, implement ML models, and facilitate real-time interventions. Below, we provide additional details for each of these three components.

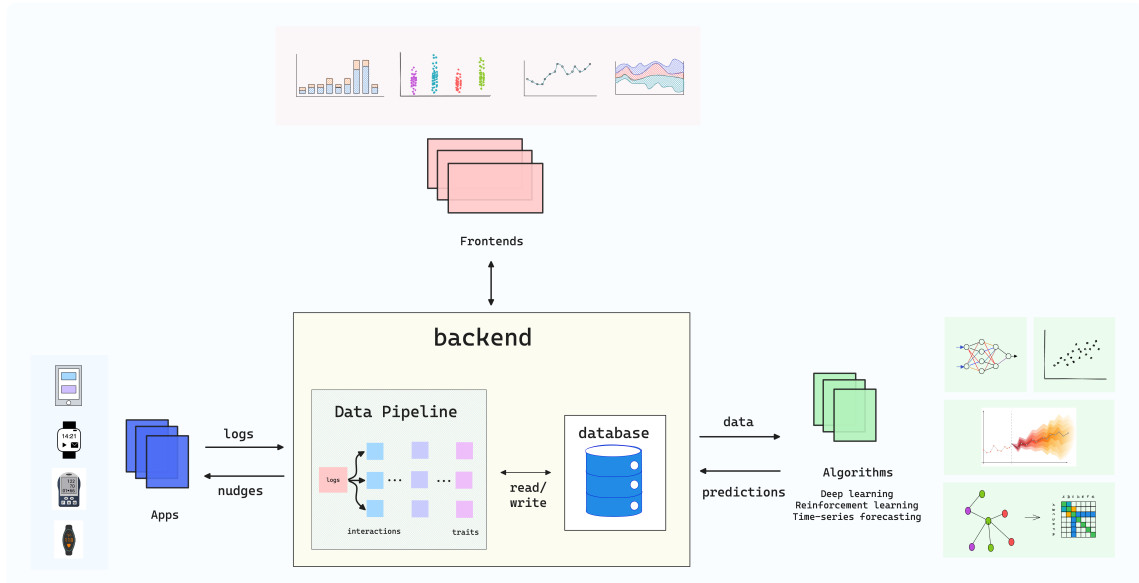


Fig. 3. **The Reinforcement Learning platform structure.** The platform creates a robust data structure that logs user events on digital health devices (e.g., mobile apps), classifying them into dynamic and static traits. The backend system analyzes and processes these traits for statistical learning and machine learning predictions, and to generate adaptive interventions based on reinforcement learning algorithms. User reactions to interventions are registered via the digital health devices and tracked through the same system to feed back into the reinforcement learning algorithms.

The *frontend* has an intuitive user interface for data analysis, model management, and intervention management and monitoring. Users can perform behavioral and clinical data analysis, statistical modeling, and create cohorts of users. Furthermore, the interface enables creating, testing, and deploying interventions such as nudges and rewards. It guides the user from the subject cohort (FHCWs, facilities, pharmacists, etc.) to analyze past

behavior, predict future or target interventions, and sample definition through algorithm selection to feature selection and target specification, including nudge alternatives in the case of interventions. The dashboard illustrates various traits, predictions, and metrics in an intuitive format for a clear interpretation of the results and impact of the interventions.

The *backend* is the platform's core, allowing the organization of data and the implementation of ML models and adaptive interventions. Data is collected and organized using domain-specific data schemas that are optimized for ML model implementation. Data security and privacy are key aspects captured in the backend. This standardized design of event logging enables the platform to be a uniform interface for transformation and aggregation across different application domains, e.g., supply chain and e-commerce marketplace, medication tracking, and patient communication. Finally, the backend hosts the predictive modeling engine that handles the ML component and the algorithmic decision-making service. These modules configure, train, host, and manage models for time series forecasting, deep and ensemble survival analysis, and sequential decisions via reinforcement learning.

Finally, an *SDK* is in charge of tracking the data with the right label and format and serving as a bridge between AI-based intervention system allowing the creation of a data pipeline between the digital tools and the platform. This SDK is a lightweight software that is embedded into mobile apps so data is tracked from the source directly. The software is optimized for the variety of mobile devices used in LMICs, such as budget smartphones and tablets, minimizing the impact of platform integration on mobile tool performance. This SDK provides the tools for collecting relevant events and the messaging service that delivers the interventions. The tool does not collect any Personal Identifiable Information (PII), complies with GDPR and HIPAA, and only tracks user actions as time-series data and contextual information such as details about the drugs being prescribed, viewed, or purchased. For interventions, the SDK is responsible for presenting these nudges to the mobile app user and automatically logging their reaction (e.g., if they opened, viewed, discarded, or blocked the nudge), ensuring complete visibility of the intervention lifecycle.

### **3.2 Predictive Modeling**

Predictive modeling is key to defining who should be targeted with specific interventions and optimizing their timing. The platform provides the user with the ability to fully utilize the use of three important prediction models: time-to-event prediction (also known

as survival analysis), demand forecasting, and recommendation systems, such as click-through-rate predictions or causal graphic analysis.

*Time-to-event Prediction* offers a way of characterizing behaviors by modeling the probability of occurrence of different events of interest and the duration between them.<sup>39–41,31</sup> For example, we can use it to predict the risk of non-adherence to treatment and the time for it to happen, client churn rate in an e-commerce platform and a health worker failing to complete a training module. These predictions can be used to identify the target subjects for specific interventions (e.g., reminders to support treatment adherence, promotions to discourage churn, motivational prompts, and incentives to encourage learning).

*Demand forecasting* is crucial in e-commerce, supply chain, and inventory optimization. For example, multivariate forecasting methods that learn either point or probabilistic forecasts from co-dependent time series for predicting the demand for different pharmaceutical products at different sites.<sup>42</sup> Several approaches are available through the platform, including autoregression- and graph-regularized factor analysis and latent state forecasting, state space modeling with LSTM and Gaussian copula for probabilistic forecasts, recurrent networks for multivariate count data with non-uniform scale, and attention-based deep models for multi-horizon interval predictions.<sup>43–47</sup> The observed stocks and previous ordering behavior, together with the predicted evolution of demand, can be used to identify what pharmacies have to receive reminders to restock or suggest products to add to the cart while ordering. Similar approaches can be used for patients using e-commerce services and to support decisions on how to distribute stocks across warehouses.

*Click-through-rate predictions* allow the generation of traits useful for content personalization. Mixed deep learning-based and factorization methods have shown good performance in the literature.<sup>48–52</sup> As described in the next section, the main mechanism in place to learn and exploit user preferences for recommendations and personalization will be that of adaptive delivery functionality. Other recommendation algorithms are available as they can provide valuable insights to be leveraged by the reinforcement learning algorithm.

Other models, such as *causal graph analysis*, help to understand cause-and-effect relationships. Which factors caused an event, e.g., why did a particular user stop using the app, or why a patient is more likely to stop a treatment. The known causes will help drive the design of personalized interventions aimed at reducing the chance of that event in the future.

### 3.3 Adaptive Intervention Delivery and Experimentation

One of the key features of the platform is the ability to design, test, and deploy adaptive interventions. Adaptive intervention delivery aims to leverage individual and contextual data for effective nudging and interventions.

Let us illustrate the main underlying ideas using a concrete example: Imagine we have a digital tool (an app) for pregnant women and mothers through which they report information on their pregnancy and newborn (e.g., overall well-being, health issues they might have, vaccination, use of healthcare services) and receive relevant health information. The tool allows us to detect any health issue the women and their newborns might have and provide them with relevant health information (such as recommendations to go to a health provider). Of course, the impact of this patient engagement tool hinges on women using the tool. However, we know user engagement is challenging to maintain.<sup>53</sup>

Through some formative research, behavioral science analysis, and focus group discussions, we decide that there are interventions we want to test, one that nudges women with some recommendations related to their characteristics (stage of pregnancy or age of their newborn) and another one that sends messages with suggestions based on behaviors similar users exhibited in the past.

We then determine the way we want to measure the success of the intervention. For this case, we use the information submitted by women through the digital tool.

The last piece is to select the characteristics (collected through the digital tool) that could influence how impactful a particular nudge will be at any given time. These could include characteristics of women and their babies (e.g., gestational age, number of previous pregnancies) and their use of the digital tool (e.g., recent exposure and interaction with nudges, the information they have sent through the app, and which content they never used). We refer to these characteristics as contextual traits.

The adaptive intervention delivery then comes into action, deciding each week (or as scheduled when defining the intervention) which of the nudges to send to each user based on the values of these contextual traits, observing how users with different characteristics react (based on how we decided to measure success) in different situations (as defined by the contextual traits). The algorithm will continue to refine (and exploit) the information

of what works best in which situation. It will also detect changes in how users react to the nudges and adapt them correspondingly.

At the end of the experiment, we will be able to select the intervention that was most successful in increasing engagement or discard them if they did not generate a significant change. Note that the selected intervention will continue to learn over time as it collects additional feedback.

This example illustrates how the impact of different intervention strategies can be assessed using the intervention and experimentation capabilities of the platform.

The user can perform rapid cycles of experimentation for intervention refinement using a sample of users, and subsequently roll out of the final intervention to all the users. Digital trials can be easily designed and conducted through the platform, including near real-time monitoring with impact estimation and heterogeneous effects analysis. Both fully randomized and adaptive designs are available. One remark is that adaptive experiments do not need to stop. In randomized trials, such as A/B testing, the researcher essentially pauses the experiment, measures which treatment is statistically significantly better than others, and decides accordingly. In adaptive studies (like the one described in the example above), it is possible to run them indefinitely, making them select the best-responding treatment. The algorithm automatically selects the most conditionally effective treatment and may also adapt to the non-stationarity of the treatment response.

Adaptive designs maximize statistical power and minimize risk, as the algorithm will consider the experimental outcomes already observed and reduce the number of subjects assigned to treatment in the event of an adverse effect on the success measure. Experimental designs that assign each subject multiple times throughout the experiment can be better options when studying the immediate impact of repeated interventions (e.g., messages sent with a specific frequency), while trials with single assignment and a *pure* control group remain the golden standard to understand one-shot interventions and long-term effects.<sup>54–61</sup>

Figure 4 shows a monitoring plot for an experiment conducted with health product recommendations delivered to pharmacies through an app in which they can order drugs from a distributor. The recommendations aimed to increase the weekly variety of ordered products by pharmacists by recommending pairs of items that, while typically ordered




Fig. 4. **Improving the variety of products in stock.** Example of an experiment in production where the goal was to send recommendations to pharmacists to increase the variety of products in stock (the so-called SKUs). The daily mean difference between the weekly purchased variety of pharmacies in treatment (receiving periodic in-app messages with item recommendations) and control (no recommendations received) is plotted together with the 95% confidence intervals. Bar plots indicate the number of users who interacted with the messages daily. The intervention was significant over the first two weeks. On average, those in the treatment group purchased five more unique items in the past week than the control group, including items that were never purchased before.

together, the pharmacy receiving the recommendation purchased only one of them regularly. The plot shows the daily mean difference in the weekly purchased variety (the number of different health products ordered in the previous seven days), and the shaded confidence intervals are related to hypothesis testing to establish differences in means. The bar plots show the number of messages interacted with daily. Even if the sample sizes used were small, the intervention had a significant (95%) positive impact during the first two weeks, with pharmacies receiving the recommendations purchasing, on average, five more different health products than those that didn't. It also shows indications of a smaller positive effect later (impact of the recommendation days and even weeks after receiving the message) and fatigue (decreasing impact of the recommendations as the novelty disappears, leading eventually to negative effects if the frequency of notifications is too high). All these were corroborated in successive rounds of experiments.



The platform offers significant flexibility in designing experiments. For example, for interventions designed for facility health workers, in which interaction between participants belonging to the same site is expected to create cross-contamination, the platform allows to use cluster randomization (i.e., the assignment to either treatment or control to all participants of the same facility). It also allows to use pairwise matching, an assignment mechanism typically used to achieve sample balance, notably in combination with cluster randomization, as it alleviates the need to increase sample sizes (to achieve the same statistical power) due to intracluster correlation. In this case, the underlying idea is to select pairs of facilities that share characteristics that we hypothesize will impact the outcome of interest and assign one to control and one to treatment arms for comparison.<sup>61–63</sup>

### 3.4 Bandit Algorithms

Adaptive interventions are typically modeled as a Markov Decision Process (MDP), which makes *Reinforcement Learning* (RL) the appropriate algorithmic paradigm.<sup>64</sup> RL-based interventions balance between outcome optimization and knowledge extraction and adapt to the feedback and observed covariates of the intervention subjects. Within RL, *bandit algorithms* are particularly well suited for the type of problems described in the previous section.

Bandit algorithms can be thought of as *online optimization methods* with *built-in model identification* from *partial feedback*: the algorithm makes adaptive sequential decisions based on the past observed covariates and interaction history with the goal of eliminating suboptimal choices.<sup>65</sup>

In our setting, the algorithm analyzes the available features of each observed subject (the static and dynamic traits), decides on the allocation or type of intervention, and later receives a feedback score of its prior decision, which it uses to adapt and improve its decision mechanism. Possible applications of adaptive algorithms include sending actionable reminders about treatment adherence and scaling back on them when the notification has been sent too recently or is not being interacted with.

The analysis of bandit algorithms is beyond the scope of this paper, but there are plenty of resources exploring these algorithms, their use, challenges, and solutions.<sup>64–92</sup>

## 4 GENERATIVE AI

*Generative AI* (GAI) technologies use *deep learning* models trained on enormous amounts of data to generate new text and images. *Large Language Models (LLMs)* are part of GAI and are trained on large datasets consisting of text. Once trained, the model can generate coherent, contextually relevant, and grammatically correct texts, answer questions, and perform text-based tasks. LLMs can revolutionize how healthcare workers interact with information systems, enabling more efficient access to medical data and supporting decision-making. The integration of LLMs in frontline health workers' mobile applications can leverage the benefits of a chat-based interface to facilitate user interaction and access to information in an easier way. Within the proposed framework, chatbots are integrated into the AI platform to employ LLM-based functionalities.

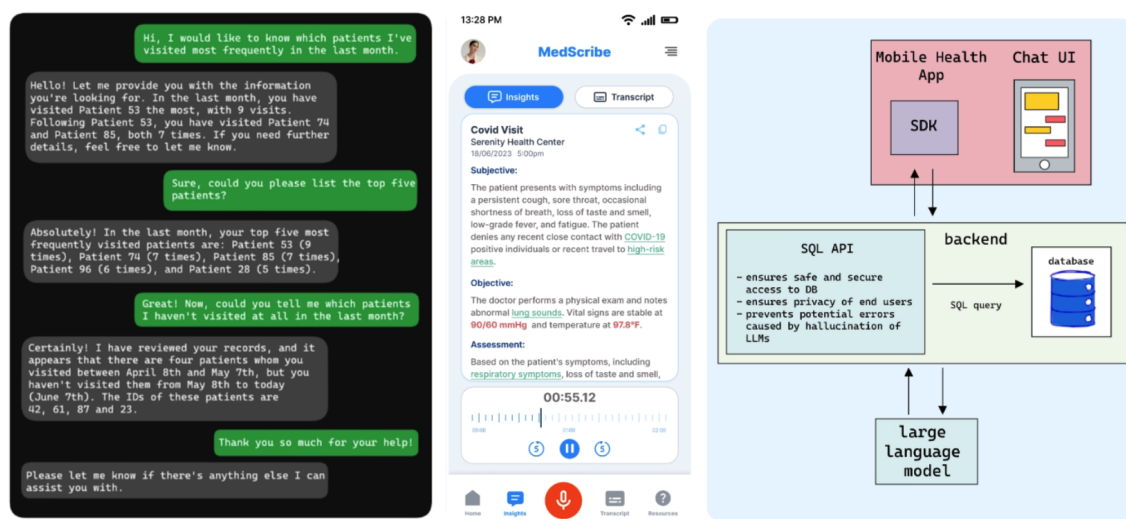


Fig. 5. **Integration of LLMs into the platform.** Left: A healthcare worker interacting with a Large Language Model (LLM); It shows how users can engage with the LLM to obtain relevant information or perform specific tasks, like creating an intervention. We can see that the LLM can handle each scenario, including understanding user intent, generating the queries to the database, and returning the requested information. Right: The proposed system architecture.

The product aims to empower CHWs, midwives, pharmacists, and clinicians by providing them with insights derived from their own data related to their practices or patients. The LLMs are pre-trained with internal data and predictions based on mobile application

usage, which are stored and organized by the AI platform. Frontline health workers can access the chatbot embedded in their digital tools to obtain answers to questions about the past or future to help their supply ordering and patient management. Some examples of questions would be: Which patients haven't I visited for more than three months? Who have I referred to clinical practice in the past two months? How many HIV RDTs should I order for the next four weeks?

LLMs can also be integrated into the SOAP (Subjective, Objective, Assessment, and Plan) framework to create a tool that will assist healthcare workers in their daily patient interactions. The tool records and summarizes provider-client interactions, converting audio to structured data, sending next-step reminders, streamlining data collection, and reducing manual data entry to enable more focused patient interactions. Workers can review session details, read audio transcripts, and add insights anytime. Additionally, GAI-generated personalized care plans and medication recommendations can combine patient-specific data with evidence-based guidelines.

Moreover, GAI and LLMs, trained on specialized datasets, aid both patients and pharmacists by providing comprehensive information about available drugs, their possible side effects, and interactions. In health-tracking and medication procurement apps, LLMs help patients make informed treatment decisions regarding non-prescription or over-the-counter drugs by offering insights into potential side effects and interactions, allowing them to explore detailed information about their prescribed drugs or find suitable alternatives.

## **5 ETHICAL AI**

In the transformative realm of digital healthcare, artificial intelligence (AI) offers notable prospects like enhanced decision-making support and more efficient patient care while also posing ethical challenges, particularly around equitable healthcare resource access, patient privacy, and data security. Addressing these, along with issues of bias prevention, fairness, and reliability, is pivotal, and efforts to establish frameworks and guidelines are ongoing.<sup>93–97</sup>

The lag between AI technological advancements and the development of adequate regulatory and ethical understanding frameworks is evident. Achieving AI tailored to societal needs necessitates development and deployment guided by transparency, accountability,

and engagement with all stakeholders in addressing ethical concerns. Essential participants in these discussions should range from patient groups and community leaders to national healthcare authorities.

Concerning *privacy and data protection*, adherence to regulatory standards like the General Data Protection Regulation (GDPR) and the Health Insurance Portability and Accountability Act (HIPAA) is imperative, ensuring that data, such as health worker and patient information, is anonymized and devoid of direct identifiers. Additionally, mechanisms should facilitate user data deletion requests and accommodate different sensitivities of data through varied anonymization levels.<sup>98,77,78</sup>

Ensuring machine learning models are *reliable* and *safe* involves scrutinizing performance, understanding biases, and guaranteeing reasonable accuracy for the intended use and target demographics. Strategies like pursuing representative datasets and implementing bias mitigation strategies are essential to ensure fairness and accuracy in model recommendations. Explicit fairness integration into models and decision algorithms is achievable through approaches like multi-objective optimization<sup>99,78,79</sup> and utilizing interpretable models and causal inference in certain use cases.<sup>100,101,61</sup>

Facilitating experimentation with various intervention strategies is crucial, allowing systematic testing of algorithmic-based feature impacts in real-world scenarios, including exploring effects among diverse social groups. Although the system leans on data and models for user support, the ultimate decision-maker remains the human-in-the-loop, with user feedback and validation being integral to enhancing recommendation accuracy and fairness.

The development and deployment of AI in healthcare necessitate establishing clear ethical guidelines that stringently address principles such as privacy, reliability, fairness, and user agency, underscored by transparency and accountability. Engaging all stakeholders, especially those typically underrepresented, in constructive discussions will be pivotal in shaping future directions and standards.

## 6 SUMMARY AND CONCLUSIONS

Pharmacies, healthcare facilities, and health workers are essential in accessing healthcare everywhere. However, in low- and middle-income countries, the fragmented ecosystems dominated by informal retailers, the circulation of counterfeit medicines, and the scarcity

of resources do not allow providers and healthcare systems to perform to their optimal capabilities. Different digital tools are emerging to improve this. Such technologies include capacity-building tools, apps that provide healthcare workers with needed supplies (e.g., drug delivery services) or medical resources (e.g., arrangement of test appointments), tools that connect them to patients (e.g., for a remote follow-up) or physicians (e.g., to get their opinion on test results), and even apps to assist them in clinical triage and diagnosis. They bring efficiency, and their users generate robust sequential data into the healthcare delivery model. The data generated contains rich information on the behavior of their users, preferences, biases and needs, and the evolution of the demand for different supplies, medications, blood, oxygen, etc.

While data collection and organization of the information is the key to any data-driven developments, it is in its ability to transform that data into adaptive interventions to support users in a personalized manner where its largest potential to impact health outcomes and patients' lives is. They are adaptive in that they adapt in time and to each user's context at that moment. Besides personalization, adaptive algorithms can be used for resource allocation to help distribute material and human means fairly and efficiently.

In this work, we presented an operational reinforcement learning platform that uses data, predictive modeling, adaptive experimentation, and bandits to enhance mobile applications through personalized content, incentives, reminders, and recommendations, including resource allocation. The platform allows the definition of controlled trials following different experimental designs to establish the effectiveness of the interventions, including (stochastic-bandit-based) adaptive treatment arm assignment.

Experiments have already shown improvements in broader operational efficiency and reduced stock deficiencies by optimizing information and encouraging all actors toward better-informed decisions in inventory management and patient care.

## **ACKNOWLEDGMENTS**

This work was supported, in whole or in part, by the Bill & Melinda Gates Foundation INV-053824. Under the grant conditions of the Foundation, a Creative Commons Attribution 4.0 Generic License has been assigned to the Author Accepted Manuscript version that might arise from this submission.

## REFERENCES

- [1] World Health Organization. *Working Together for Health: The World Health Report 2006*. World Health Organization, Geneva, 2006.
- [2] Theo Lippeveld, Rainer Sauerborn, and Claude Bodart. *Design and implementation of health information systems*. WHO, 2000.
- [3] Commission on Social Determinants of Health. *Closing the Gap in a Generation: Health Equity through Action on the Social Determinants of Health. Final Report of the Commission on Social Determinants of Health*. World Health Organization, Geneva, 2008.
- [4] Sophie Harman. *Global health governance*. Routledge, 2012.
- [5] Claudia Allemani, Hannah K. Weir, Helena Carreira, Rhea Harewood, Devon Spika, Xiao-Si Wang, Finian Bannon, Jane V. Ahn, Christopher J. Johnson, Audrey Bonaventure, Rafael Marcos-Gragera, Charles Stiller, Gulnar Azevedo E Silva, Wan-Qing Chen, Olufemi J. Ogunbiyi, Bernard Rachet, Matthew J. Soeberg, Hui You, Tomohiro Matsuda, Magdalena Bielska-Lasota, Hans Storm, Thomas C. Tucker, and Michel P. Coleman. Global surveillance of cancer survival 1995-2009: analysis of individual data for 25 676 887 patients from 279 population-based registries in 67 countries (concord-2). *The Lancet*, 385(9972):977–1010, 2015.
- [6] OECD. *Ready for the Next Crisis? Investing in Health System Resilience*. 2023. doi: <https://doi.org/https://doi.org/10.1787/1e53cf80-en>. URL <https://www.oecd-ilibrary.org/content/publication/1e53cf80-en>.
- [7] Amy Abernethy, Laura Adams, Meredith Barrett, Christine Bechtel, Patricia Brennan, Atul Butte, Judith Faulkner, Elaine Fontaine, Stephen Friedhoff, John Halamka, Michael Howell, Kevin Johnson, Peter Long, Deven McGraw, Redonda Miller, Peter Lee, Jonathan Perlin, Donald Rucker, Lew Sandy, Lucia Savage, Lisa Stump, Paul Tang, Eric Topol, Reed Tuckson, and Kristen Valdes. The promise of digital health: Then, now, and the future. *NAM Perspectives*, 2022, 2022. doi: 10.31478/202206e. URL <https://doi.org/10.31478/202206e>. PMID: 36177208.
- [8] World Health Organization. Classification of digital health interventions v1.0: a shared language to describe the uses of digital technology for health. Technical report, World Health Organization, 2018.
- [9] Brian Wahl, Aline Cossy-Gantner, Stefan Germann, and Nina R Schwalbe. Artificial intelligence (ai) and global health: how can ai contribute to health in resource-poor settings? *BMJ global health*, 3(4): e000798, 2018.
- [10] Mark McCarthy and Ewan Birney. Personalized profiles for disease risk must capture all facets of health. *Nature*, 597, 2021. doi: 10.1038/d41586-021-02401-0.
- [11] James A Hanley and Barbara J McNeil. The next generation of evidence-based medicine. *Nature Medicine*, 29(1):29–36, 2023. doi: <https://doi.org/10.1038/s41591-022-02160-z>.
- [12] Ahmed Hosny and Hugo JWL Aerts. Artificial intelligence for global health. *Science*, 366(6468): 955–956, 2019.

- [13] Inbal Nahum-Shani, Shawna N Smith, Bonnie J Spring, Linda M Collins, Katie Witkiewitz, Ambuj Tewari, and Susan A Murphy. Just-in-time adaptive interventions (jitais) in mobile health: key components and design principles for ongoing health behavior support. *Annals of Behavioral Medicine*, 52(6):446–462, 2018. doi: 10.1007/s12160-016-9830-8.
- [14] Marianne Menictas, Mashfiqui Rabbi, Predrag Klasnja, and Susan Murphy. Artificial intelligence decision-making in mobile health. *The Biochemist*, 41(5):20–24, 2019. doi: 10.1042/BIO04105020.
- [15] Stephanie M Carpenter, Marianne Menictas, Inbal Nahum-Shani, David W Wetter, and Susan A Murphy. Developments in mobile health just-in-time adaptive interventions for addiction science. *Current addiction reports*, 7(3):280–290, 2020. doi: 10.1007/s40429-020-00322-y.
- [16] G Bidargaddi.N amd Schrader, P Klasnja, J Licinio, and S Murphy. Designing m-health interventions for precision mental health support. *Translational Psychiatry*, 10(222), 2020. doi: 10.1038/s41398-020-00895-2.
- [17] Daniel DL Coppersmith, Walter Dempsey, Evan Kleiman, Kate Bentley, Susan Murphy, and Matthew Nock. Just-in-time adaptive interventions for suicide prevention: Promise, challenges, and future directions, 2021.
- [18] Emir Hrnjic and Nikodem Tomczak. Machine learning and behavioral economics for personalized choice architecture, 2019.
- [19] Tong Geng, Xiliang Lin, Harikesh S. Nair, Jun Hao, Bin Xiang, and Shurui Fan. Comparison lift: Bandit-based experimentation system for online advertising. *CoRR*, abs/2009.07899, 2020. URL <https://arxiv.org/abs/2009.07899>.
- [20] Yi Liu and Lihong Li. A map of bandits for e-commerce. In *KDD 2021 Workshop on Multi-Armed Bandits and Reinforcement Learning (MARBLE)*, 2021. URL <https://www.amazon.science/publications/a-map-of-bandits-for-e-commerce>.
- [21] Sarah Gimbel, Nami Kawakyu, Hallie Dau, and Jennifer A. Unger. A missing link: HIV-/AIDS-related mHealth interventions for health workers in low- and middle-income countries. *Current HIV/AIDS Reports*, 15(6):414–422, 2018. doi: 10.1007/s11904-018-0416-x.
- [22] África Perriáñez, Andrew Trister, Madhav Nekkar, Ana Fernández del Río, and Pedro L. Alonso. Adaptive interventions for global health: A case study of malaria, 2023.
- [23] Danilo Bzdok, Naomi Altman, and Martin Krzywinski. Statistics versus machine learning. *Nature Methods*, 15:233–234, 2018. ISSN 1548-7105. doi: 10.1038/nmeth.4642. URL <https://doi.org/10.1038/nmeth.4642>.
- [24] Kun-Hsing Yu, Andrew L. Beam, and Isaac S. Kohane. Artificial intelligence in healthcare. *Nature Biomedical Engineering*, 2018.
- [25] Chao Yu, Jiming Liu, Shamim Nemati, and Guosheng Yin. Reinforcement learning in healthcare: A survey. *ACM Comput. Surv.*, 55(1), nov 2021. ISSN 0360-0300. doi: 10.1145/3477600. URL <https://doi.org/10.1145/3477600>.



- [26] Siqi Liu, Kay Choong, Kee Yuan, Leo Anthony, Xingzhi Sun, and Mengling Feng. Reinforcement learning for clinical decision support in critical care: Comprehensive review. *Journal of Medical Internet Research*, 2020.
- [27] Mingyang Liu, Xiaotong Shen, and Wei Pan. Deep reinforcement learning for personalized treatment recommendation. *Statistics in Medicine*, 41(20):4034–4056, jun 2022. doi: 10.1002/sim.9491.
- [28] Xiangjun Du, Aaron A. King, Ryan J. Woods, and Mercedes Pascual. Predictive modeling of influenza shows the promise of applied evolutionary biology. *Science Translational Medicine*, 9(eaan5325), 2017. doi: 10.1126/scitranslmed.aan5325. URL <https://stm.sciencemag.org/content/9/413/eaan5325>.
- [29] Benjamin M. Althouse, Yee Y. Ng, and Derek A.T. Cummings. Early prediction of dengue fever outbreaks using machine learning algorithms. *Scientific Reports*, 8(1), 2018. doi: 10.1038/s41598-018-34107-w. URL <https://www.nature.com/articles/s41598-018-34107-w>.
- [30] Yue Liu, Wei Wang, Jin-Qiu Jiang, Ming Zhang, and Kang-Lin Wan. Use of a machine learning algorithm to predict individuals with increased risk of first episode of tuberculosis: A population-based cohort study. *The Lancet Digital Health*, 2(6):e295–e303, 2020. doi: 10.1016/S2589-7500(20)30077-6. URL [https://www.thelancet.com/journals/landig/article/PIIS2589-7500\(20\)30077-6/fulltext](https://www.thelancet.com/journals/landig/article/PIIS2589-7500(20)30077-6/fulltext).
- [31] Babaniyi Yusuf Olaniyi, Ana Fernández del Río, África Periañez, and Lauren Bellhouse. User engagement and churn in mobile health applications. In *Proceedings of 2022 ACM KDD Workshop on Applied Data Science for Healthcare (DSHealth 2022)*, 2022.
- [32] Anna Guitart, Ana Fernández del Río, África Periañez, and Lauren Bellhouse. Midwifery Learning and Forecasting: Predicting Content Demand with User-Generated Logs, 2021.
- [33] Anna Guitart, Afsaneh Heydari, Eniola Olaleye, Jelena Ljubicic, Ana Fernández del Río, África Periañez, and Lauren Bellhouse. A recommendation system to enhance midwives’ capacities in low-income countries. *arXiv preprint arXiv:2111.01786*, 2021.
- [34] Iraklis Katsaris, , and Nikolas Vidakis and. Adaptive e-learning systems through learning styles: A review of the literature. *Advances in Mobile Learning Educational Research*, 1(2):124–145, 2021. doi: 10.25082/amler.2021.02.007.
- [35] Wendy Hardeman, Julie Houghton, Kathleen Lane, Andy Jones, and Felix Naughton. A systematic review of just-in-time adaptive interventions (jitais) to promote physical activity. *International Journal of Behavioral Nutrition and Physical Activity*, 16:31, 2019. doi: 10.1186/s12966-019-0792-7.
- [36] Anne Caroline Benski, Nicole C Schmidt, Manuela Viviano, Giovanna Stancanelli, Adelia Soaroby, and Michael R Reich. Improving the quality of antenatal care using mobile health in madagascar: Five-year cross-sectional study. *JMIR Mhealth Uhealth*, 8:8(07), 2020. doi: <https://doi.org/10.1186/s12961-021-00761-7>.
- [37] Monique Vledder, Jed Friedman, Mirja Sjöblom, Thomas Brown, and Prashant Yadav. Improving supply chain for essential drugs in low-income countries: Results from a large scale randomized experiment in zambia. *Health Systems & Reform*, 5(2):158–177, 2019. doi: 10.1080/23288604.2019.1596050. URL <https://doi.org/10.1080/23288604.2019.1596050>. PMID: 31194645.



- [38] Ruth Levine, Jessica Pickett, Neelam Sekhri, and Prashant Yadav. Demand forecasting for essential medical technologies. *American Journal of Law & Medicine*, 34(2-3):225–255, 2008. doi: 10.1177/009885880803400206.
- [39] Marvin N. Wright, Theresa Dankowski, and Andreas Ziegler. Unbiased split variable selection for random survival forests using maximally selected rank statistics. *Statistics in Medicine*, 36(8):1272–1284, 2017. doi: <https://doi.org/10.1002/sim.7212>.
- [40] Wei Fu and Jeffrey S. Simonoff. Survival trees for left-truncated and right-censored data, with application to time-varying covariate data. *Biostatistics*, 18(2):352–369, 2016. ISSN 1468–4357. doi: 10.1093/biostatistics/kxw047.
- [41] Changhee Lee, Jinsung Yoon, and Mihaela van der Schaar. Dynamic-DeepHit: A deep learning approach for dynamic survival analysis with competing risks based on longitudinal data. *IEEE Transactions on Biomedical Engineering*, 67(1):122–133, 2020. doi: 10.1109/tbme.2019.2909027.
- [42] Konstantinos Benidis, Syama Sundar Rangapuram, Valentin Flunkert, Yuyang Wang, Danielle Maddix, Caner Turkmen, Jan Gasthaus, Michael Bohlke-Schneider, David Salinas, Lorenzo Stella, François-Xavier Aubet, Laurent Callot, and Tim Januschowski. Deep Learning for Time Series Forecasting: Tutorial and Literature Survey. *ACM Computing Surveys*, 55(6):121:1–121:36, December 2022. ISSN 0360-0300. doi: 10.1145/3533382. URL <https://doi.org/10.1145/3533382>.
- [43] Hsiang-Fu Yu, Nikhil Rao, and Inderjit S Dhillon. Temporal Regularized Matrix Factorization for High-dimensional Time Series Prediction. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL [https://papers.nips.cc/paper\\_files/paper/2016/hash/85422afb467e9456013a2a51d4dff702-Abstract.html](https://papers.nips.cc/paper_files/paper/2016/hash/85422afb467e9456013a2a51d4dff702-Abstract.html).
- [44] Matthias Seeger, Syama Rangapuram, Yuyang Wang, David Salinas, Jan Gasthaus, Tim Januschowski, and Valentin Flunkert. Approximate Bayesian Inference in Linear State Space Models for Intermittent Demand Forecasting at Scale, September 2017. URL <http://arxiv.org/abs/1709.07638>. arXiv:1709.07638 [cs, stat].
- [45] David Salinas, Michael Bohlke-Schneider, Laurent Callot, Roberto Medico, and Jan Gasthaus. High-dimensional multivariate forecasting with low-rank Gaussian Copula Processes. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL [https://proceedings.neurips.cc/paper\\_files/paper/2019/hash/0b105cf1504c4e241fcc6d519ea962fb-Abstract.html](https://proceedings.neurips.cc/paper_files/paper/2019/hash/0b105cf1504c4e241fcc6d519ea962fb-Abstract.html).
- [46] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3):1181–1191, July 2020. ISSN 0169-2070. doi: 10.1016/j.ijforecast.2019.07.001. URL <https://www.sciencedirect.com/science/article/pii/S0169207019301888>.
- [47] Bryan Lim, Sercan O. Arık, Nicolas Loeff, and Tomas Pfister. Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4):1748–1764, October 2021. ISSN 0169-2070. doi: 10.1016/j.ijforecast.2021.03.012. URL <https://www.sciencedirect.com/science/article/pii/S0169207021000637>.

- [48] Yanru Qu, Han Cai, Kan Ren, Weinan Zhang, Yong Yu, Ying Wen, and Jun Wang. Product-based neural networks for user response prediction. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pages 1149–1154. IEEE, 2016. doi: 10.1109/ICDM.2016.0151.
- [49] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. Deepfm: A factorization-machine based neural network for ctr prediction, 2017.
- [50] Jianxun Lian, Xiaohuan Zhou, Fuzheng Zhang, Zhongxia Chen, Xing Xie, and Guangzhong Sun. xdeepfm: Combining explicit and implicit feature interactions for recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1754–1763. ACM, 2018. doi: 10.1145/3219819.3220023.
- [51] Wantong Lu, Yantao Yu, Yongzhe Chang, Zhen Wang, Chenhui Li, and Bo Yuan. A dual input-aware factorization machine for ctr prediction. In *IJCAI’20: Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pages 3139–3145, 2020. URL <https://www.ijcai.org/proceedings/2020/0434.pdf>.
- [52] Anna Guitart, Ana Fernández del Río, África Periañez, and Lauren Bellhouse. Midwifery Learning and Forecasting: Predicting Content Demand with User-Generated Logs, 2021.
- [53] Siobhan O’Connor, Peter Hanlon, Catherine A. O’Donnell, Sonia Garcia, Julie Glanville, and Frances S. Mair. Understanding factors affecting patient and public engagement and recruitment to digital health interventions: a systematic review of qualitative studies. *BMC Medical Informatics and Decision Making*, 16:120, 2016. ISSN 1472-6947. doi: 10.1186/s12911-016-0359-3. URL <https://doi.org/10.1186/s12911-016-0359-3>.
- [54] Predrag Klasnja, Eric B. Hekler, Saul Shiffman, Audrey Boruvka, Daniel Almirall, Ambuj Tewari, and Susan A. Murphy. Microrandomized trials: An experimental design for developing just-in-time adaptive interventions. *Health Psychology*, 34:1220–1228, 2015. ISSN 1930-7810. doi: 10.1037/hea0000305. Place: US Publisher: American Psychological Association.
- [55] Tianchen Qian, Ashley E. Walton, Linda M. Collins, Predrag Klasnja, Stephanie T. Lanza, Inbal Nahum-Shani, Mashfiqui Rabbi, Michael A. Russell, Maureen A. Walton, Hyesun Yoo, and Susan A. Murphy. The microrandomized trial for developing digital interventions: Experimental design and data analysis considerations. *Psychological Methods*, 27(5):874–894, October 2022. ISSN 1939-1463. doi: 10.1037/met0000283.
- [56] Kelly W Zhang, Lucas Janson, and Susan A Murphy. Statistical inference after adaptive sampling in non-markovian environments, 2022.
- [57] Giuseppe Burtini, Jason Loepky, and Ramon Lawrence. A survey of online experiment design with the stochastic multi-armed bandit, 2015.
- [58] Jiayu Yao, Emma Brunskill, Weiwei Pan, Susan Murphy, and Finale Doshi-Velez. Power constrained bandits. In Ken Jung, Serena Yeung, Mark Sendak, Michael Sjoding, and Rajesh Ranganath, editors, *Proceedings of the 6th Machine Learning for Healthcare Conference*, volume 149 of *Proceedings of Machine Learning Research*, pages 209–259. PMLR, 2021. URL <https://proceedings.mlr.press/v149/yao21a.html>.

- [59] Raaz Dwivedi, Susan Murphy, and Devavrat Shah. Counterfactual inference for sequential experimental design. 2022.
- [60] Ding Xiang, Rebecca West, Jiaqi Wang, Xiquan Cui, and Jinzhou Huang. Multi armed bandit vs. a/b tests in e-commerce - confidence interval and hypothesis test power perspectives. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '22, page 4204–4214, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393850. doi: 10.1145/3534678.3539144. URL <https://doi.org/10.1145/3534678.3539144>.
- [61] Guido W. Imbens and Donald B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge University Press, apr 2015. doi: 10.1017/cbo9781139025751.
- [62] Kosuke Imai, Gary King, and Clayton Nall. The essential role of pair matching in cluster-randomized experiments, with application to the mexican universal health insurance evaluation. *Statistical Science*, 24(1), feb 2009. doi: 10.1214/08-sts274.
- [63] Yuehao Bai. Optimality of matched-pair designs in randomized controlled trials. *American Economic Review*, 112(12):3911–3940, dec 2022. doi: 10.1257/aer.20201856.
- [64] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA, October 2018. ISBN 978-0-262-03924-6.
- [65] Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, Cambridge, 2020. ISBN 978-1-108-48682-8. doi: 10.1017/9781108571401. URL <https://www.cambridge.org/core/books/bandit-algorithms/8E39FD004E6CE036680F90DD0C6F09FC>.
- [66] Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, WWW '10, pages 661–670, New York, NY, USA, April 2010. Association for Computing Machinery. ISBN 978-1-60558-799-8. doi: 10.1145/1772690.1772758. URL <https://doi.org/10.1145/1772690.1772758>.
- [67] Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual Bandits with Linear Payoff Functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 208–214. JMLR Workshop and Conference Proceedings, June 2011. URL <https://proceedings.mlr.press/v15/chu11a.html>. ISSN: 1938-7228.
- [68] Shipra Agrawal and Navin Goyal. Thompson Sampling for Contextual Bandits with Linear Payoffs. In *Proceedings of the 30th International Conference on Machine Learning*, pages 127–135. PMLR, May 2013. URL <https://proceedings.mlr.press/v28/agrawal13.html>. ISSN: 1938-7228.
- [69] Carlos Riquelme, George Tucker, and Jasper Snoek. Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=SyYe6k-CW>.
- [70] Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural Contextual Bandits with UCB-based Exploration. In *Proceedings of the 37th International Conference on Machine Learning*, pages 11492–11502. PMLR, November 2020. URL <https://proceedings.mlr.press/v119/zhou20a.html>. ISSN: 2640-3498.

- [71] Pan Xu, Zheng Wen, Handong Zhao, and Quanquan Gu. Neural contextual bandits with deep representation and shallow exploration. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=xnYACQQuaGV>.
- [72] Ofir Nabati, Tom Zahavy, and Shie Mannor. Online Limited Memory Neural-Linear Bandits with Likelihood Matching. In *Proceedings of the 38th International Conference on Machine Learning*, pages 7905–7915. PMLR, July 2021. URL <https://proceedings.mlr.press/v139/nabati21a.html>. ISSN: 2640-3498.
- [73] William Fedus, Prajit Ramachandran, Rishabh Agarwal, Yoshua Bengio, Hugo Larochelle, Mark Rowland, and Will Dabney. Revisiting Fundamentals of Experience Replay. In *Proceedings of the 37th International Conference on Machine Learning*, pages 3061–3071. PMLR, November 2020. URL <https://proceedings.mlr.press/v119/fedus20a.html>. ISSN: 2640-3498.
- [74] Gerardo Duran-Martin, Aleya Kara, and Kevin Murphy. Efficient Online Bayesian Inference for Neural Bandits. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, pages 6002–6021. PMLR, May 2022. URL <https://proceedings.mlr.press/v151/duran-martin22a.html>. ISSN: 2640-3498.
- [75] Gerald L. Smith, Stanley F. Schmidt, and Leonar A. McGee. Application of statistical filter theory to the optimal estimation of position and velocity on board a circumlunar vehicle. Technical report, NASA, 1962.
- [76] Frederick E. Daum. Extended kalman filters. In *Encyclopedia of Systems and Control*, pages 411–413. Springer London, 2015. doi: 10.1007/978-1-4471-5058-9\_62.
- [77] Madalina M. Dragan and Ann Nowe. Designing multi-objective multi-armed bandits algorithms: A study. In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, August 2013. doi: 10.1109/IJCNN.2013.6707036. ISSN: 2161-4407.
- [78] Eralp Turgay, Doruk Oner, and Cem Tekin. Multi-objective Contextual Bandit Problem with Similarity Information. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, pages 1673–1681. PMLR, March 2018. URL <https://proceedings.mlr.press/v84/turgay18a.html>. ISSN: 2640-3498.
- [79] Shiyin Lu, Guanghui Wang, Yao Hu, and Lijun Zhang. Multi-Objective Generalized Linear Bandits. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19, IJCAI’19*, pages 3080–3086, Macao, China, July 2019. International Joint Conferences on Artificial Intelligence Organization. ISBN 978-0-9992411-4-1. doi: 10.24963/ijcai.2019/427. URL <https://doi.org/10.24963/ijcai.2019/427>.
- [80] Romain Lopez, Inderjit S. Dhillon, and Michael I. Jordan. Learning from eXtreme Bandit Feedback. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(10):8732–8740, May 2021. ISSN 2374-3468. doi: 10.1609/aaai.v35i10.17058. URL <https://ojs.aaai.org/index.php/AAAI/article/view/17058>. Number: 10.
- [81] Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in finitely-armed and continuous-armed bandits. *Theoretical Computer Science*, 412(19):1832–1852, April 2011. ISSN 0304-3975. doi:

- 10.1016/j.tcs.2010.12.059. URL <https://www.sciencedirect.com/science/article/pii/S030439751000767X>.
- [82] P. Whittle. Restless Bandits: Activity Allocation in a Changing World. *Journal of Applied Probability*, 25:287–298, 1988. ISSN 0021-9002. doi: 10.2307/3214163. URL <https://www.jstor.org/stable/3214163>. Publisher: Applied Probability Trust.
- [83] Yundi Qian, Chao Zhang, Bhaskar Krishnamachari, and Milind Tambe. Restless Poachers: Handling Exploration-Exploitation Tradeoffs in Security Domains. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, AAMAS '16, pages 123–131, Richland, SC, May 2016. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 978-1-4503-4239-1. URL <https://dl.acm.org/doi/10.5555/2936924.2936946>.
- [84] Saeed Bagheri and Anna Scaglione. The Restless Multi-Armed Bandit Formulation of the Cognitive Compressive Sensing Problem. *IEEE Transactions on Signal Processing*, 63(5):1183–1198, March 2015. ISSN 1941-0476. doi: 10.1109/TSP.2015.2389620. Conference Name: IEEE Transactions on Signal Processing.
- [85] Aditya Mate, Lovish Madaan, Aparna Taneja, Neha Madhiwalla, Shresth Verma, Gargi Singh, Aparna Hegde, Pradeep Varakantham, and Milind Tambe. Field Study in Deploying Restless Multi-Armed Bandits: Assisting Non-profits in Improving Maternal and Child Health. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(11):12017–12025, June 2022. ISSN 2374-3468. doi: 10.1609/aaai.v36i11.21460. URL <https://ojs.aaai.org/index.php/AAAI/article/view/21460>. Number: 11.
- [86] Arpita Biswas, Gaurav Aggarwal, Pradeep Varakantham, and Milind Tambe. Learn to intervene: An adaptive learning policy for restless bandits in application to preventive healthcare. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4039–4046. International Joint Conferences on Artificial Intelligence Organization, 8 2021. doi: 10.24963/ijcai.2021/556. URL <https://doi.org/10.24963/ijcai.2021/556>. Main Track.
- [87] Siddharth Nishtala, Lovish Madaan, Aditya Mate, Harshavardhan Kamarthi, Anirudh Grama, Divy Thakkar, Dhyanesh Narayanan, Suresh Chaudhary, Neha Madhiwalla, Ramesh Padmanabhan, Aparna Hegde, Pradeep Varakantham, Balaraman Ravindran, and Milind Tambe. Selective Intervention Planning using Restless Multi-Armed Bandits to Improve Maternal and Child Health Outcomes, October 2021. URL <http://arxiv.org/abs/2103.09052>. arXiv:2103.09052 [cs].
- [88] Elliot Lee, Mariel S. Lavieri, and Michael Volk. Optimal Screening for Hepatocellular Carcinoma: A Restless Bandit Model. *Manufacturing & Service Operations Management*, 21(1):198–212, January 2019. ISSN 1523-4614. doi: 10.1287/msom.2017.0697. URL <https://pubsonline.informs.org/doi/abs/10.1287/msom.2017.0697>. Publisher: INFORMS.
- [89] Biswarup Bhattacharya. Restless Bandits visiting Villages: A Preliminary Study on distributing Public Health Services. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, COMPASS '18, pages 1–8, New York, NY, USA, June 2018. Association for Computing Machinery. ISBN 978-1-4503-5816-3. doi: 10.1145/3209811.3209865. URL <https://doi.org/10.1145/3209811.3209865>.

- [90] Jackie Baek, Justin J. Boutilier, Vivek F. Farias, Jonas Oddur Jonasson, and Erez Yoeli. Policy Optimization for Personalized Interventions in Behavioral Health, March 2023. URL <http://arxiv.org/abs/2303.12206>. arXiv:2303.12206 [cs].
- [91] Jackson A Killian, Andrew Perrault, and Milind Tambe. Beyond "to act or not to act": Fast lagrangian approaches to general multi-action restless bandits. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, pages 710–718, 2021.
- [92] Han-Ching Ou, Christoph Siebenbrunner, Jackson Killian, Meredith B. Brooks, David Kempe, Yevgeniy Vorobeychik, and Milind Tambe. Networked Restless Multi-Armed Bandits for Mobile Interventions, January 2022. URL <http://arxiv.org/abs/2201.12408>. arXiv:2201.12408 [cs].
- [93] Inioluwa Deborah Raji, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. Closing the ai accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT\* '20*, page 33–44, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450369367. doi: 10.1145/3351095.3372873. URL <https://doi.org/10.1145/3351095.3372873>.
- [94] Vinodkumar Prabhakaran, Margaret Mitchell, Timnit Gebru, and Iason Gabriel. A human rights-based approach to responsible ai, 2022.
- [95] Markus D. Dubber, Frank Pasquale, and Sunit Das, editors. *The Oxford Handbook of Ethics of AI*. Oxford University Press, jul 2020. doi: 10.1093/oxfordhb/9780190067397.001.0001.
- [96] Anna Jobin, Marcello Ienca, and Effy Vayena. The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9):389–399, sep 2019. doi: 10.1038/s42256-019-0088-2.
- [97] Irene Y Chen, Emma Pierson, Sherri Rose, Shalmali Joshi, Kadija Ferryman, and Marzyeh Ghassemi. Ethical machine learning in healthcare. *Annual review of biomedical data science*, 4:123–144, 2021.
- [98] Latanya Sweeney. K-anonymity: A model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10(5):557–570, oct 2002. ISSN 0218-4885. doi: 10.1142/S0218488502001648. URL <https://doi.org/10.1142/S0218488502001648>.
- [99] Vishwali Mhasawade, Yuan Zhao, and Rumi Chunara. Machine learning and algorithmic fairness in public and population health. *Nature Machine Intelligence*, 3(8):659–666, jul 2021. doi: 10.1038/s42256-021-00373-4.
- [100] Ben Glocker, Mirco Musolesi, Jonathan Richens, and Caroline Uhler. Causality in digital medicine. *Nature Communications*, 12(1), 2021.
- [101] Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, may 2019. doi: 10.1038/s42256-019-0048-x.