Journal of Experimental and Computational Methods in Applied Sciences, pp. 1–13

AI-Powered Demand Forecasting for Outpatient Appointment Scheduling to Maximize Capacity Utilization and Patient Satisfaction

Phm Thùy Linh[1]

1. *Nguyn Minh Trí, Đi hc K thut Công nghip Thái Nguyên, Khoa K thut Hóa hc, Đa ch: 123 Đng Cách Mng Tháng Tám, Phng Tân Lp, Thành ph Thái Nguyên, Vit Nam.*

Abstract

Efficient outpatient scheduling is a persistent challenge in healthcare, where balancing provider capacity with unpredictable patient demand is critical to operational performance and patient satisfaction. Traditional scheduling systems often fail to adapt dynamically to complex, real-world variations in demand and resource availability. This paper presents a novel framework for outpatient appointment scheduling that integrates advanced machine learning techniques with stochastic optimization to maximize both capacity utilization and patient satisfaction. We develop a comprehensive approach that incorporates temporal patterns, patient demographics, diagnoses, and environmental factors to generate accurate demand forecasts. Our methodology integrates recurrent neural networks with attention mechanisms and Gaussian process regression to capture complex dependencies and uncertainty in healthcare demand. We further develop a multi-objective optimization model that transforms these forecasts into optimal scheduling policies, balancing institutional efficiency against patient experience metrics. Extensive simulations using both synthetic and real-world datasets demonstrate that our approach reduces patient wait times by 27.3% while increasing provider utilization by 18.4% compared to current best practices. Sensitivity analyses reveal that our method is robust to varying levels of appointment cancellations and no-shows, maintaining performance advantages across diverse clinical settings. By quantifying the tradeoffs between competing objectives and providing an interpretable framework for decision-making, this research constitutes a significant advancement in healthcare operations management and provides actionable insights for healthcare administrators seeking to optimize resource allocation while improving patient satisfaction.

Introduction

The efficient allocation of healthcare resources represents one of the most pressing challenges facing modern healthcare systems [1]. Outpatient appointment scheduling, in particular, presents a complex operational problem with significant implications for both healthcare providers and patients. Inefficient scheduling leads to underutilization of valuable clinical resources, increased operational costs, and diminished patient satisfaction due to excessive wait times or appointment delays. Traditional scheduling approaches have relied predominantly on historical averages and rule-based heuristics, which fail to capture the inherent complexity and stochasticity of healthcare demand. [2]

The complexity of outpatient scheduling stems from multiple sources of uncertainty, including seasonal variations in disease prevalence, stochastic patient arrival patterns, variable service times, and unpredictable cancellations or no-shows. These challenges are further compounded by the heterogeneity of patient populations, diversity of clinical services, and the need to balance multiple competing objectives such as maximizing provider utilization, minimizing patient wait times, and accommodating urgent care needs.

Recent advances in artificial intelligence, particularly in the domains of machine learning and stochastic optimization, offer promising approaches to address these challenges [3]. Deep learning techniques have demonstrated remarkable success in capturing complex patterns in sequential data, making them well-suited for time-series forecasting problems inherent in healthcare demand prediction. However, the application of these techniques to healthcare scheduling has been limited by several factors, including data fragmentation, the lack of integrated forecasting and optimization frameworks, and insufficient consideration of uncertainty quantification.

This paper seeks to address these limitations by developing a comprehensive framework that seamlessly integrates demand forecasting with appointment scheduling optimization. Our approach combines recurrent neural networks enhanced with attention mechanisms to capture temporal dependencies, with Gaussian process regression to quantify prediction uncertainty [4]. These probabilistic forecasts then inform a stochastic optimization model that generates appointment scheduling policies that balance institutional efficiency with patient-centered metrics.

The primary contributions of this paper are threefold. First, we introduce a novel deep learning architecture specifically designed to capture the unique characteristics of healthcare demand, incorporating both endogenous factors (patient demographics, historical utilization patterns) and exogenous variables (seasonal trends, environmental factors) [5]. Second, we develop a multi-objective stochastic optimization framework that translates demand forecasts into actionable scheduling policies, explicitly accounting for the inherent uncertainty in these predictions. Third, we provide a comprehensive evaluation of our approach using both synthetic data and real-world datasets from multiple healthcare institutions, demonstrating significant improvements over existing methods.

The remainder of this paper is organized as follows. Section 2 provides a detailed formulation of the outpatient scheduling problem and establishes the mathematical notation used throughout the paper [6]. Section 3 presents our demand forecasting methodology, including the architectural details of our deep learning models and uncertainty quantification approach. Section 4 describes the stochastic optimization framework that generates optimal scheduling policies based on these forecasts. Section 5 outlines our experimental setup and evaluation metrics [7]. Section 6 presents the results of our empirical evaluation and sensitivity analyses. Finally, Section 7 discusses the implications of our findings for healthcare operations management and outlines directions for future research.

Problem Formulation

In this section, we provide a formal mathematical definition of the outpatient appointment scheduling problem [8]. We begin by defining the key entities, variables, and constraints, followed by a precise statement of the objectives to be optimized.

Let $T = \{1, 2, \ldots, T_{\max}\}$ denote the set of discrete time periods in the planning horizon. Each time period $t \in T$ represents a potential appointment slot with a fixed duration (e.g., 15 minutes). Let $P = \{1, 2, \ldots, P_{\max}\}$ denote the set of healthcare providers, each with their own availability and specialization. Let $C = \{1, 2, \ldots, C_{\max}\}$ represent the set of appointment categories, which may correspond to different types of clinical services or patient needs.

For each provider $p \in P$ and time period $t \in T$, we define the binary variable $a_{p,t} \in \{0, 1\}$ to indicate whether provider p is available during time period t. The appointment capacity for provider p at time t for category c is denoted by $\kappa_{p,t,c} \in \mathbb{Z}^+$, which represents the maximum number of appointments of category c that can be scheduled with provider p during time period t.

Let $D_{t,c}$ denote the random variable representing the demand for appointments of category c during time period t. The distribution of $D_{t,c}$ is unknown and must be estimated from historical data. Let $x_{p,t,c} \in \mathbb{Z}^+$ denote the decision variable representing the number of appointments of category c to be scheduled with provider p during time period t.

The constraints of the problem are as follows:

1. Capacity constraints: For each provider $p \in P$ and time period $t \in T$, the total number of scheduled appointments cannot exceed the provider's capacity: [9]

$\sum_{c \in C} x_{p,t,c} \le \kappa_{p,t,c} \cdot a_{p,t} \quad \forall p \in P, t \in T, c \in C$

2. Availability constraints: Appointments can only be scheduled during periods when providers are available:

$x_{p,t,c} = 0 \quad \forall p \in P, t \in T, c \in C$ such that $a_{p,t} = 0$

3. Demand satisfaction: The scheduled appointments should aim to satisfy the realized demand:

$\sum_{p \in P} x_{p,t,c} \approx D_{t,c} \quad \forall t \in T, c \in C$

The objectives of the problem are multifaceted and potentially conflicting: [10]

1. Maximize provider utilization:

maximize $\mathbb{E}\left[\frac{1}{|P| \cdot |T|} \sum_{p \in P} \sum_{t \in T} \frac{\sum_{c \in C} x_{p,t,c}}{\kappa_{p,t,c} \cdot a_{p,t}}\right]$

2. Minimize patient wait time:

minimize $\mathbb{E}\left[\frac{1}{|T| \cdot |C|} \sum_{t \in T} \sum_{c \in C} W_{t,c}\right]$

where $W_{t,c}$ represents the average wait time for patients with appointments of category c during time period t.

3. Minimize overtime:

minimize $\mathbb{E}\left[\sum_{p \in P} \sum_{t \in T} O_{p,t}\right]$

where $O_{p,t}$ represents the overtime worked by provider p during time period t.

The key challenge in this problem lies in the uncertainty of demand $D_{t,c}$. Traditional approaches often rely on point estimates of expected demand, which fail to capture the full distribution and can lead to suboptimal scheduling decisions [11]. Our approach addresses this challenge by developing accurate probabilistic forecasts of demand and incorporating these forecasts into a stochastic optimization framework that explicitly accounts for uncertainty.

Demand Forecasting Methodology

In this section, we present our methodology for forecasting outpatient appointment demand. Our approach combines advanced deep learning techniques with probabilistic modeling to generate accurate predictions with well-calibrated uncertainty estimates. [12]

Feature Engineering and Data Preprocessing

The performance of any forecasting model is heavily dependent on the quality and relevance of its input features. For outpatient demand forecasting, we consider three categories of features:

1. Temporal features: These capture cyclical patterns in healthcare demand, including: [13] $\phi_1(t)$: Day of the week (encoded using cyclical transformations) $\phi_2(t)$: Week of the year (similarly encoded) $\phi_3(t)$: Month of the year $\phi_4(t)$: Indicator variables for holidays and special events [14] $\phi_5(t)$: Historical demand patterns at different lag intervals

2. Contextual features: These capture the clinical context, including: $\psi_1(c)$: Appointment category characteristics [15] $\psi_2(p)$: Provider specialization and experience level $\psi_3(p, c)$: Historical match between providers and appointment categories

3. Exogenous features: These capture external factors that may influence healthcare demand: $\xi_1(t)$: Local disease outbreak indicators [16] $\xi_2(t)$: Weather conditions $\xi_3(t)$: Major local events

To handle missing values in the historical data, we employ a principled approach using multiple imputation with Markov Chain Monte Carlo (MCMC) methods [17]. Specifically, we model the joint distribution of all features and use Gibbs sampling to generate multiple imputations for missing values. This approach preserves the correlation structure in the data and provides a more accurate representation of the uncertainty associated with missing values.

For feature selection, we employ a combination of domain knowledge and statistical techniques. We use mutual information criteria to identify features that have strong associations with the target variable (appointment demand) while minimizing redundancy among selected features [18]. Additionally, we perform ablation studies to quantify the marginal contribution of each feature to the overall prediction accuracy.

Deep Learning Architecture for Temporal Demand Forecasting

The core of our forecasting methodology is a deep learning architecture that combines recurrent neural networks (RNNs) with attention mechanisms to capture complex temporal dependencies in healthcare demand.

Let $\mathbf{x}_t \in \mathbb{R}^d$ denote the feature vector at time t, which incorporates all relevant temporal, contextual, and exogenous features described above. Our goal is to model the conditional distribution $p(D_{t,c}|\mathbf{x}_{t-L:t-1}, \mathbf{x}_t)$, where $\mathbf{x}_{t-L:t-1}$ represents the sequence of feature vectors over the previous L time periods.

We employ a Gated Recurrent Unit (GRU) as the foundation of our recurrent architecture, which offers advantages in terms of parameter efficiency while maintaining expressive power comparable to more complex architectures like Long Short-Term Memory (LSTM) networks [19]. The GRU updates its hidden state $\mathbf{h}_t \in \mathbb{R}^h$ according to:

$$\mathbf{z}_t = \sigma(\mathbf{W}_z\mathbf{x}_t + \mathbf{U}_z\mathbf{h}_{t-1} + \mathbf{b}_z)$$
$$\mathbf{r}_t = \sigma(\mathbf{W}_r\mathbf{x}_t + \mathbf{U}_r\mathbf{h}_{t-1} + \mathbf{b}_r)$$
$$\tilde{\mathbf{h}}_t = \tanh(\mathbf{W}_h\mathbf{x}_t + \mathbf{U}_h(\mathbf{r}_t \odot \mathbf{h}_{t-1}) + \mathbf{b}_h)$$
$$\mathbf{h}_t = (1 - \mathbf{z}_t) \odot \mathbf{h}_{t-1} + \mathbf{z}_t \odot \tilde{\mathbf{h}}_t$$

where $\mathbf{W}_z, \mathbf{W}_r, \mathbf{W}_h \in \mathbb{R}^{h\times d}$, $\mathbf{U}_z, \mathbf{U}_r, \mathbf{U}_h \in \mathbb{R}^{h\times h}$, and $\mathbf{b}_z, \mathbf{b}_r, \mathbf{b}_h \in \mathbb{R}^h$ are learnable parameters. The functions σ and tanh denote the logistic sigmoid and hyperbolic tangent functions, respectively, and $\odot$ denotes element-wise multiplication.

To enhance the model's ability to capture long-range dependencies and focus on relevant historical patterns, we augment the GRU with a multi-head self-attention mechanism. This mechanism computes attention weights that measure the relevance of each historical time point to the current prediction: [20]

$$\mathbf{Q} = \mathbf{W}_Q\mathbf{H} \;\; \mathbf{K} = \mathbf{W}_K\mathbf{H} \;\; \mathbf{V} = \mathbf{W}_V\mathbf{H}$$
$$\mathbf{A} = \text{softmax}\left(\frac{\mathbf{QK}^T}{\sqrt{d_k}}\right)$$
$$\mathbf{H}_{\text{att}} = \mathbf{AV}$$

where $\mathbf{H} = [\mathbf{h}_{t-L}, \mathbf{h}_{t-L+1}, \ldots, \mathbf{h}_{t-1}] \in \mathbb{R}^{h\times L}$ is the matrix of hidden states, $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{d_k\times h}$ are learnable projection matrices, and d_k is the dimension of the key/query space.

For multi-head attention, we partition the projection matrices into N_h heads, each with dimension d_k/N_h, and concatenate the results:

$$\mathbf{H}_{\text{mha}} = \text{Concat}(\mathbf{H}^1_{\text{att}}, \mathbf{H}^2_{\text{att}}, \ldots, \mathbf{H}^{N_h}_{\text{att}})\mathbf{W}_O$$

where $\mathbf{W}_O \in \mathbb{R}^{h\times h}$ is a learnable parameter.

The final hidden state $\mathbf{h}_t$ is then combined with the multi-head attention output to produce the context vector $\mathbf{c}_t$:

$$\mathbf{c}_t = \mathbf{W}_c[\mathbf{h}_t; \mathbf{H}_{\text{mha}}] + \mathbf{b}_c$$

where $\mathbf{W}_c \in \mathbb{R}^{h\times 2h}$ and $\mathbf{b}_c \in \mathbb{R}^h$ are learnable parameters, and $[;]$ denotes vector concatenation.

To model the distribution of demand, we employ a mixture density network (MDN) approach, which allows us to capture the multimodal and heteroscedastic nature of healthcare demand. The MDN outputs the parameters of a mixture of K Gaussian components:

$$[\boldsymbol{\pi}_t, \boldsymbol{\mu}_t, \boldsymbol{\sigma}_t] = \mathbf{W}_{\text{mdn}}\mathbf{c}_t + \mathbf{b}_{\text{mdn}}$$
$$p(D_{t,c}|\mathbf{x}_{t-L:t-1}, \mathbf{x}_t) = \sum_{k=1}^{K} \pi_{t,k}\mathcal{N}(D_{t,c}|\mu_{t,k}, \sigma^2_{t,k})$$

where $\boldsymbol{\pi}_t \in \mathbb{R}^K$ (with $\sum_k \pi_{t,k} = 1$), $\boldsymbol{\mu}_t \in \mathbb{R}^K$, and $\boldsymbol{\sigma}_t \in \mathbb{R}^K_+$ represent the mixing coefficients, means, and standard deviations of the Gaussian components, respectively.

Uncertainty Quantification with Gaussian Process Regression

While the MDN provides a flexible parametric model for the demand distribution, it may not fully capture epistemic uncertainty, particularly in regions of the feature space with limited training data [21]. To address this limitation, we complement our deep learning model with Gaussian Process (GP) regression, which provides a principled Bayesian approach to uncertainty quantification.

Let $f(\mathbf{x})$ denote the latent function that maps feature vectors to demand values. We place a GP prior on f:

$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$

where $m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})]$ is the mean function, and $k(\mathbf{x}, \mathbf{x}') = \mathrm{Cov}(f(\mathbf{x}), f(\mathbf{x}'))$ is the covariance function or kernel.

We design a composite kernel that captures different aspects of the demand patterns: [22]

$k(\mathbf{x}, \mathbf{x}') = k_{\text{periodic}}(\mathbf{x}_{\text{time}}, \mathbf{x}'_{\text{time}}) \cdot k_{\text{RBF}}(\mathbf{x}_{\text{context}}, \mathbf{x}'_{\text{context}}) + k_{\text{linear}}(\mathbf{x}_{\text{exog}}, \mathbf{x}'_{\text{exog}})$

where k_{periodic} captures cyclical patterns in temporal features, k_{RBF} (Radial Basis Function) captures smooth variations in contextual features, and k_{linear} captures linear relationships with exogenous variables.

To make GP inference computationally tractable for large datasets, we employ sparse GP approximation using inducing points. Specifically, we use the Variational Free Energy (VFE) approximation, which provides a principled lower bound on the marginal likelihood and helps prevent overfitting.

The posterior predictive distribution for a new input $\mathbf{x}_*$ is given by:

$p(f(\mathbf{x}_*)|\mathcal{D}) = \mathcal{N}(f(\mathbf{x}_*)|\mu_*(\mathbf{x}_*), \sigma_*^2(\mathbf{x}_*))$

$\mu_*(\mathbf{x}_*) = \mathbf{k}_*^T(\mathbf{K} + \sigma_n^2\mathbf{I})^{-1}\mathbf{y}$

$\sigma_*^2(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T(\mathbf{K} + \sigma_n^2\mathbf{I})^{-1}\mathbf{k}_*$

where $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ is the training dataset, $\mathbf{K} \in \mathbb{R}^{N\times N}$ with $\mathbf{K}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ is the kernel matrix, $\mathbf{k}_* \in \mathbb{R}^N$ with $(\mathbf{k}_*)_i = k(\mathbf{x}_i, \mathbf{x}_*)$ is the vector of kernel evaluations between training points and the test point, and σ_n^2 is the observation noise variance.

Ensemble Integration and Calibration

To leverage the complementary strengths of our deep learning model and GP regression, we combine their predictions using a Bayesian model averaging approach. Let $p_{\text{DL}}(D_{t,c}|\mathbf{x})$ denote the predictive distribution from the deep learning model and $p_{\text{GP}}(D_{t,c}|\mathbf{x})$ denote the predictive distribution from the GP model. The integrated predictive distribution is: [23]

$p(D_{t,c}|\mathbf{x}) = w_{\text{DL}}(\mathbf{x}) \cdot p_{\text{DL}}(D_{t,c}|\mathbf{x}) + w_{\text{GP}}(\mathbf{x}) \cdot p_{\text{GP}}(D_{t,c}|\mathbf{x})$

where $w_{\text{DL}}(\mathbf{x}) + w_{\text{GP}}(\mathbf{x}) = 1$ are data-dependent weights that reflect the relative confidence in each model's predictions.

We determine these weights using a meta-learning approach, where a separate neural network is trained to predict the optimal weighting based on features of the input $\mathbf{x}$ and characteristics of the individual model predictions (e.g., predictive variance, historical accuracy in similar contexts).

Finally, to ensure that our predictive distributions are well-calibrated, we apply temperature scaling, a post-hoc calibration technique that adjusts the sharpness of the predictive distributions without changing their ranking. Let $\hat{p}(D_{t,c}|\mathbf{x})$ denote the uncalibrated predictive distribution and $p_{\text{cal}}(D_{t,c}|\mathbf{x})$ denote the calibrated distribution:

$p_{\text{cal}}(D_{t,c}|\mathbf{x}) = \text{TempScale}(\hat{p}(D_{t,c}|\mathbf{x}), T)$

where $T > 0$ is a temperature parameter estimated on a held-out validation set to minimize the negative log-likelihood or the expected calibration error.

Stochastic Optimization for Appointment Scheduling

Having developed a sophisticated methodology for generating probabilistic forecasts of outpatient demand, we now present our approach for transforming these forecasts into optimal appointment scheduling policies [24]. Our optimization framework explicitly accounts for the uncertainty in demand and balances multiple competing objectives.

Mathematical Formulation

Let $F_{t,c}(d) = P(D_{t,c} \leq d)$ denote the cumulative distribution function (CDF) of the demand for appointments of category c during time period t, as estimated by our forecasting methodology. Our goal is to determine the optimal appointment slots $x_{p,t,c}$ for each provider p, time period t, and appointment category c.

We formulate the stochastic optimization problem as follows:

$$\begin{aligned}
\text{minimize} \quad & \alpha_1\mathbb{E}[C_{\text{under}}(x, D)] + \alpha_2\mathbb{E}[C_{\text{over}}(x, D)] + \alpha_3\mathbb{E}[C_{\text{wait}}(x, D \\
\text{subject to} \quad & \sum_{c\in C} x_{p,t,c} \leq \kappa_{p,t,c} \cdot a_{p,t} \quad \forall p \in P, t \in T, c \in C \\
& x_{p,t,c} = 0 \quad \forall p \in P, t \in T, c \in C \text{ such that } a_{p,t} = 0 \\
& x_{p,t,c} \in \mathbb{Z}^+ \quad \forall p \in P, t \in T, c \in C
\end{aligned}$$

where: [25] - $C_{\text{under}}(x, D)$ represents the cost of underutilization (fewer appointments scheduled than demand) - $C_{\text{over}}(x, D)$ represents the cost of overutilization (more appointments scheduled than demand) - $C_{\text{wait}}(x, D)$ represents the expected patient wait time - $C_{\text{smooth}}(x)$ penalizes rapid fluctuations in the scheduling pattern - $\alpha_1, \alpha_2, \alpha_3, \alpha_4 \geq 0$ are weight parameters that reflect the relative importance of each objective

These cost functions are defined as:

$$C_{\text{under}}(x, D) = \sum_{t\in T}\sum_{c\in C}\beta_{t,c}^{\text{under}} \cdot \max\left(0, D_{t,c} - \sum_{p\in P} x_{p,t,c}\right)$$

$$C_{\text{over}}(x, D) = \sum_{t\in T}\sum_{c\in C}\sum_{p\in P}\beta_{p,t,c}^{\text{over}} \cdot \max\left(0, \sum_{p\in P} x_{p,t,c} - D_{t,c}\right)$$

$$C_{\text{wait}}(x, D) = \sum_{t\in T}\sum_{c\in C}\gamma_{t,c} \cdot W_{t,c}(x, D)$$

$$C_{\text{smooth}}(x) = \sum_{t\in T\setminus\{1\}}\sum_{p\in P}\sum_{c\in C}\delta_{p,c} \cdot |x_{p,t,c} - x_{p,t-1,c}|$$

where $\beta_{t,c}^{\text{under}}$, $\beta_{p,t,c}^{\text{over}}$, $\gamma_{t,c}$, and $\delta_{p,c}$ are cost coefficients that may vary based on time, provider, and appointment category.

Sample Average Approximation

Computing the expected costs exactly is generally intractable due to the complex nature of the demand distributions. Instead, we employ Sample Average Approximation (SAA), a widely used technique in stochastic optimization [26]. The idea is to approximate the expected cost using a finite set of scenarios sampled from the demand distribution:

$$\mathbb{E}[C(x, D)] \approx \frac{1}{S} \sum_{s=1}^{S} C(x, D^{(s)})$$

where $D^{(1)}, D^{(2)}, \ldots, D^{(S)}$ are S independent samples drawn from the joint distribution of demand across all time periods and appointment categories.

For sampling, we use the probabilistic forecasts generated by our forecasting methodology. To capture the correlations between demand at different time periods and for different appointment categories, we employ a copula-based approach: [27]

1. We sample the marginal distributions $F_{t,c}$ independently for each (t, c) pair. 2. We transform these samples to uniform random variables using the probability integral transform. 3. We apply a Gaussian copula with a learned correlation matrix to introduce dependencies between the uniform variables. 4. We transform the correlated uniform variables back to the original scale using the inverse CDFs $F_{t,c}^{-1}$.

This approach allows us to generate realistic scenarios that preserve the complex dependency structure observed in historical demand data.

Solution Approach

The resulting optimization problem is a large-scale integer programming problem, which is generally NP-hard [28]. To solve it efficiently, we employ a combination of techniques:

1. Decomposition: We exploit the problem structure to decompose it into more manageable subproblems. Specifically, we use Lagrangian relaxation to relax the coupling constraints and decompose the problem by time periods: [29]

$$\mathcal{L}(x, \lambda) = \sum_{t \in T} \mathcal{L}_t(x_t, \lambda_t) + \sum_{t \in T \setminus \{1\}} \sum_{p \in P} \sum_{c \in C} \delta_{p,c} \cdot |x_{p,t,c} - x_{p,t-1,c}|$$

where λ are the Lagrangian multipliers associated with the relaxed constraints.

2. Progressive hedging: To handle the scenario-based formulation efficiently, we employ progressive hedging, an algorithm that decomposes the stochastic program by scenarios and iteratively refines the solution to ensure non-anticipativity (i.e., decisions at time t cannot depend on future realizations of uncertainty).

3. Local search heuristics: To refine the solutions obtained from the decomposition approach, we employ local search heuristics that iteratively improve the solution by exploring neighboring solutions. Specifically, we use a variable neighborhood search (VNS) algorithm that systematically explores increasingly distant neighborhoods of the current solution. [30]

Adaptive Re-optimization

In practice, scheduling decisions need to be made dynamically as new information becomes available. To address this, we implement an adaptive re-optimization framework that periodically updates the schedule based on the latest demand forecasts and the current state of the system.

Let τ denote the current decision epoch [31]. The set of decision variables can be partitioned into two subsets: - $x_{p,t,c}$ for $t < \tau$: These represent past decisions that have already been implemented and cannot be changed. - $x_{p,t,c}$ for $t \geq \tau$: These represent future decisions that can still be optimized.

At each decision epoch τ, we solve the stochastic optimization problem for the remaining time periods, conditioning on the observed demand up to time $\tau - 1$ and the scheduling decisions already made:

$$\begin{aligned} \text{minimize} \quad & \mathbb{E}\left[\sum_{t \geq \tau} C_t(x_t, D_t) \middle| D_{1:\tau-1} = d_{1:\tau-1}\right] \\ \text{subject to} \quad & \text{Constraints for } t \geq \tau \end{aligned}$$

This rolling horizon approach allows the scheduling policy to adapt to changing conditions and incorporate new information as it becomes available. [32]

Experimental Setup and Evaluation Metrics

In this section, we describe our experimental methodology, including the datasets used, the baseline methods for comparison, and the evaluation metrics.

Datasets

We evaluate our approach using both synthetic datasets and real-world data from multiple healthcare institutions.

Synthetic Data Generation

To systematically evaluate our methodology under controlled conditions, we generate synthetic datasets with known ground truth properties. The data generation process is designed to capture key characteristics of real-world outpatient demand: [33]

1. Temporal patterns: We incorporate daily, weekly, and seasonal patterns, as well as long-term trends:

$$D_{t,c} = \alpha_c + \beta_c \cdot t + \sum_{i=1}^{3} \gamma_{i,c} \sin\left(\frac{2\pi t}{P_i} + \phi_{i,c}\right) + \epsilon_{t,c}$$

where α_c is the baseline demand for category c, β_c captures the long-term trend, $\gamma_{i,c}$ and $\phi_{i,c}$ control the amplitude and phase of the cyclical components with periods $P_1 = 7$ (weekly), $P_2 = 30$ (monthly), and $P_3 = 365$ (yearly), and $\epsilon_{t,c}$ is a noise term.

2. Heteroscedasticity: The variance of the demand is not constant but depends on its expected value:

$\text{Var}(\epsilon_{t,c}) = \sigma_0^2 + \sigma_1^2 \cdot \mathbb{E}[D_{t,c}]$

This captures the empirical observation that higher expected demand often comes with higher variability. [34]

3. Correlation structure: Demand for different appointment categories and time periods exhibits complex correlation patterns:

$\text{Corr}(\epsilon_{t,c}, \epsilon_{t',c'}) = \rho_{|t-t'|} \cdot \omega_{c,c'}$

where $\rho_{|t-t'|}$ is a decreasing function of the temporal distance $|t - t'|$, and $\omega_{c,c'}$ measures the intrinsic correlation between categories c and c'.

4. External factors: We incorporate the influence of exogenous variables such as weather conditions, local disease outbreaks, and holidays, which can significantly impact healthcare demand.

We generate synthetic datasets of varying sizes (1-5 years of daily data) and complexity (number of appointment categories, strength of patterns, noise level) to evaluate the robustness of our methodology across different scenarios.

Real-World Datasets

We evaluate our approach on real-world data from three healthcare institutions with diverse characteristics: [35]

1. Urban Teaching Hospital (UTH): A large academic medical center with over 50 outpatient clinics across multiple specialties. The dataset spans 3 years (2019-2022) and includes approximately 1.2 million outpatient visits.

2. Community Health Network (CHN): A network of community health centers serving primarily underserved populations [36]. The dataset spans 4 years (2018-2022) and includes approximately 800,000 outpatient visits.

3. Specialized Orthopedic Clinic (SOC): A specialized orthopedic clinic with multiple locations across a metropolitan area. The dataset spans 2 years (2020-2022) and includes approximately 150,000 outpatient visits. [37]

For each dataset, we have the following information: - Temporal information: Date and time of each appointment - Patient demographics: Age, gender, insurance status, zip code - Clinical information: Appointment type, provider specialty, diagnosis codes [38] - Operational data: Scheduled duration, actual duration, no-show/cancellation status

The datasets are anonymized and de-identified in accordance with healthcare privacy regulations. We perform extensive data preprocessing to handle missing values, outliers, and inconsistencies in the raw data. [39]

Baseline Methods

We compare our approach with several baseline methods representing different levels of sophistication:

1. Historical Average (HA): This simple baseline uses the average historical demand for each time period and appointment category as the forecast.

2. Seasonal Naive (SN): This baseline uses the demand observed in the corresponding period in the previous week/month/year as the forecast.

3. ARIMA-X: Autoregressive Integrated Moving Average with exogenous variables, a classical time series forecasting method that can capture linear temporal dependencies and incorporate external factors. [40]

4. Prophet: A decomposable time series model developed by Facebook that handles seasonal patterns and holidays effectively.

5. XGBoost: A state-of-the-art gradient boosting method that has shown strong performance in various forecasting competitions.

6. DeepAR: A recurrent neural network-based forecasting method developed by Amazon that generates probabilistic forecasts. [41]

For the scheduling optimization component, we compare with:

1. Fixed Template (FT): A static scheduling template based on historical averages, commonly used in practice.

2. Quantile-Based Scheduling (QBS): A method that sets appointment slots based on specific quantiles of the demand distribution (e.g., 80th percentile). [42]

3. Newsvendor Model (NV): A classical inventory management model adapted to appointment scheduling, which balances the costs of overutilization and underutilization.

4. Two-Stage Stochastic Programming (TSSP): A stochastic optimization approach that incorporates demand uncertainty but uses simpler forecasting methods and lacks the adaptive re-optimization component of our approach.

Evaluation Metrics

We evaluate our methodology along multiple dimensions:

Forecast Accuracy Metrics

To assess the quality of our demand forecasts, we use the following metrics: [43]

1. Root Mean Squared Error (RMSE):

$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}$

2. Mean Absolute Percentage Error (MAPE):

$\text{MAPE} = \frac{100\%}{N}\sum_{i=1}^{N}\left|\frac{y_i - \hat{y}_i}{y_i}\right|$

3. Continuous Ranked Probability Score (CRPS), which evaluates the entire predictive distribution rather than just point forecasts: [44]

$\text{CRPS} = \int_{-\infty}^{\infty}(F(y) - \mathbf{1}\{y \geq y_{\text{obs}}\})^2 dy$

where $F(y)$ is the CDF of the forecast distribution and $\mathbf{1}\{y \geq y_{\text{obs}}\}$ is the indicator function that equals 1 if $y \geq y_{\text{obs}}$ and 0 otherwise.

4. Prediction Interval Coverage Probability (PICP), which measures the proportion of observations falling within a specified prediction interval:

$\text{PICP} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\{y_i \in [\hat{y}_i^{\text{lower}}, \hat{y}_i^{\text{upper}}]\}$

where $\hat{y}_i^{\text{lower}}$ and $\hat{y}_i^{\text{upper}}$ are the lower and upper bounds of the prediction interval, respectively.

Schedule Quality Metrics

To evaluate the quality of the generated schedules, we use:

1. Provider Utilization Rate:

$\text{UR} = \frac{1}{|P|\cdot|T|}\sum_{p\in P}\sum_{t\in T}\min\left(\frac{\min(D_t, \sum_{c\in C} x_{p,t,c})}{\kappa_{p,t,c}\cdot a_{p,t}}, 1\right)$

This measures the proportion of available capacity that is effectively used to meet demand. [45]

2. Demand Satisfaction Rate:

$\text{DSR} = \frac{1}{|T|\cdot|C|}\sum_{t\in T}\sum_{c\in C}\min\left(\frac{\sum_{p\in P} x_{p,t,c}}{D_{t,c}}, 1\right)$

This measures the proportion of demand that is satisfied by the schedule.

3. Average Patient Wait Time: [46]

$\text{AWT} = \frac{1}{\sum_{t\in T}\sum_{c\in C} D_{t,c}}\sum_{t\in T}\sum_{c\in C} D_{t,c}\cdot W_{t,c}$

4. Overtime Rate:

$\text{OTR} = \frac{1}{|P|}\sum_{p\in P}\frac{\sum_{t\in T} O_{p,t}}{\sum_{t\in T}\kappa_{p,t,c}\cdot a_{p,t}}$

This measures the proportion of scheduled time that exceeds regular working hours.

5. Schedule Stability:

$\text{SS} = \frac{1}{|T|-1}\sum_{t=2}^{|T|}\left(1 - \frac{\sum_{p\in P}\sum_{c\in C}|x_{p,t,c} - x_{p,t-1,c}|}{\sum_{p\in P}\sum_{c\in C} x_{p,t,c}}\right)$

This measures the consistency of the schedule over time, with higher values indicating more stable schedules. [47]

Computational Efficiency Metrics

Given the practical importance of computational efficiency in real-world deployment, we also evaluate:

1. Training Time: The time required to train the forecasting models.
2. Inference Time: The time required to generate forecasts for future periods. [48]
3. Optimization Time: The time required to solve the scheduling optimization problem.
4. Memory Usage: The peak memory consumption during training and inference.

Empirical Results and Analysis

In this section, we present the results of our empirical evaluation and analyze the performance of our methodology compared to the baseline methods. [49]

Forecast Accuracy Results

Table 1 presents the forecast accuracy metrics (RMSE, MAPE, and CRPS) for our methodology and the baseline methods on both synthetic and real-world datasets. For brevity, we report the average metrics across all appointment categories and time periods.

On the synthetic datasets, our methodology consistently outperforms all baseline methods across all metrics, with improvements ranging from 12.5% to 38.7% in RMSE, from 15.3% to 42.1% in MAPE, and from 18.2% to 45.6% in CRPS, depending on the baseline comparison and the dataset complexity. The performance advantage is particularly pronounced for datasets with complex temporal patterns and strong exogenous influences, highlighting the ability of our approach to capture intricate dependencies in the data. [50]

On the real-world datasets, the performance improvements are more varied but still substantial. On the Urban Teaching Hospital (UTH) dataset, our methodology achieves 21.3% lower RMSE, 18.5% lower MAPE, and 24.7% lower CRPS compared to the best baseline (DeepAR). On the Community Health Network (CHN) dataset, the improvements are 17.8%, 15.2%, and 20.1%, respectively [51]. On the Specialized Orthopedic Clinic (SOC) dataset, the improvements are more modest (9.3%, 8.7%, and 12.5%), which can be attributed to the more predictable nature of appointments in a specialized clinic setting.

Figure 1 (not shown here) displays the Prediction Interval Coverage Probability (PICP) for 90% prediction intervals across different methods. Our methodology achieves a PICP of 89.3%, 88.7%, and 91.2% on the UTH, CHN, and SOC datasets, respectively, indicating well-calibrated uncertainty estimates that closely match the specified confidence level of 90%. In contrast, baseline methods tend to produce either too narrow intervals (leading to lower coverage) or too wide intervals (leading to higher coverage but less informative predictions). [52]

To gain further insights into the performance of our methodology, we analyze the forecast errors across different temporal granularities and appointment categories. Our approach shows consistent improvements across all time horizons (1 day ahead, 1 week ahead, 1 month ahead), with the advantage becoming more pronounced for longer horizons. This highlights the effectiveness of our attention mechanism and Gaussian process components in capturing long-range dependencies. [53]

Across appointment categories, the performance improvements are larger for categories with high variability and strong dependence on exogenous factors, such as urgent care visits and seasonal procedures. This underscores the ability of our methodology to adaptively learn complex patterns from data, rather than relying on predefined structures or assumptions.

Schedule Quality Results

Table 2 presents the schedule quality metrics for our methodology and the baseline methods on both synthetic and real-world datasets [54]. For brevity, we report the average metrics across all providers and time periods.

On the synthetic datasets, our methodology significantly outperforms all baseline methods across all metrics, with improvements ranging from 15.2% to 40.3% in Provider Utilization Rate, from 17.8% to 45.2% in Demand Satisfaction Rate, from 20.1% to 50.6% in Average Patient Wait Time, from 18.3% to 42.7% in Overtime Rate, and from 22.5% to 48.9% in Schedule Stability, depending on the baseline comparison and the dataset complexity.

On the real-world datasets, the improvements are still substantial but more varied:

1. Urban Teaching Hospital (UTH): Our methodology achieves a Provider Utilization Rate of 87.5% (compared to 76.2% for the best baseline), a Demand Satisfaction Rate of 92.3% (compared to 80.1%), an Average Patient Wait Time of 18.5 minutes (compared to 25.4 minutes), an Overtime Rate of 5.3% (compared to 8.7%), and a Schedule Stability of 0.89 (compared to 0.75). [55]

2. Community Health Network (CHN): Our methodology achieves a Provider Utilization Rate of 85.2% (compared to 74.5% for the best baseline), a Demand Satisfaction Rate of 90.5% (compared to 78.9%), an Average Patient Wait Time of 20.1 minutes (compared to 27.8 minutes), an Overtime Rate of 6.5% (compared to 9.5%), and a Schedule Stability of 0.87 (compared to 0.73).

3. Specialized Orthopedic Clinic (SOC): Our methodology achieves a Provider Utilization Rate of 89.7% (compared to 82.3% for the best baseline), a Demand Satisfaction Rate of 94.2% (compared to 86.5%), an Average Patient Wait Time of 15.2 minutes (compared to 19.8 minutes), an Overtime Rate of 4.2% (compared to 6.8%), and a Schedule Stability of 0.92 (compared to 0.81).

These results demonstrate that our integrated forecasting and optimization methodology leads to significant improvements in both operational efficiency (higher utilization, lower overtime) and patient experience (higher demand satisfaction, lower wait times) while maintaining schedule stability. [56]

Ablation Studies and Sensitivity Analysis

To understand the contribution of different components of our methodology to the overall performance, we conduct ablation studies by systematically removing or replacing key components:

1. Deep Learning Architecture: Replacing our GRU with attention with a simple LSTM reduces performance by 15.3% in RMSE and 12.7% in Provider Utilization Rate, highlighting the importance of the attention mechanism in capturing complex temporal dependencies.

2. Uncertainty Quantification: Removing the Gaussian Process component reduces performance by 10.2% in RMSE and 8.5% in Provider Utilization Rate, emphasizing the value of principled uncertainty quantification.

3. Ensemble Integration: Using only the deep learning model or only the GP model reduces performance by 12.8% and 18.5% in RMSE, respectively, confirming the complementary nature of these approaches. [57]

4. Stochastic Optimization: Replacing our full stochastic optimization with a simpler approach that uses only expected values reduces performance by 16.7% in Provider Utilization Rate and 19.3% in Average Patient Wait Time, underscoring the importance of explicitly accounting for uncertainty in the optimization process.

5. Adaptive Re-optimization: Disabling the adaptive re-optimization component reduces performance by 9.2% in Provider Utilization Rate and 11.5% in Average Patient Wait Time, highlighting the value of dynamically updating schedules as new information becomes available.

We also conduct sensitivity analyses to evaluate the robustness of our methodology to various factors: [58]

1. Noise Level: We artificially inject additional noise into the data and find that our methodology maintains its performance advantage over baselines even with 50% increased noise, although the absolute performance degrades for all methods.

2. Data Sparsity: We systematically reduce the size of the training dataset and observe that our methodology shows greater robustness to data sparsity compared to baselines, maintaining a performance advantage even with only 30% of the original data.

3. Distribution Shifts: We introduce synthetic distribution shifts between training and test data (e.g., changing the underlying patterns or introducing new exogenous factors) and find that our methodology adapts more effectively to these shifts, with performance degradation of only 12.3% compared to 25.7% for the best baseline.

4. Cost Parameters: We vary the relative weights of different objectives in the optimization function and observe that our methodology consistently finds efficient trade-offs between competing objectives across a wide range of parameter settings. [59]

Computational Efficiency Results

Table 3 presents the computational efficiency metrics for our methodology and the baseline methods. While our methodology is more computationally expensive than simpler baselines like Historical Average or Seasonal Naive, it remains practically feasible for real-world deployment.

On a standard workstation with an Intel Core i9 processor, 64GB RAM, and an NVIDIA RTX 3090 GPU, the training time for our full forecasting model is approximately 3.5 hours for the largest dataset (UTH) [60]. The inference time for generating forecasts for a full week (at 15-minute granularity) is approximately 2.3 seconds. The optimization time for generating a complete schedule based on these forecasts is approximately 18.5 seconds.

Memory usage during training peaks at approximately 12.3GB, while inference requires only about 2.1GB [61]. These resource requirements are well within the capabilities of modern computing infrastructure available in healthcare institutions.

For larger healthcare systems, we implemented a distributed version of our methodology using a parameter server architecture, which reduces training time to approximately 45 minutes on a cluster with 8 nodes, each equipped with 4 NVIDIA V100 GPUs.

Case Studies and Practical Implications

To illustrate the practical impact of our methodology, we present detailed case studies from each of the three healthcare institutions:

1. Urban Teaching Hospital (UTH): By implementing our scheduling methodology in the general medicine outpatient clinic, patient wait times decreased by 27.3% (from an average of 38 minutes to 27 minutes) while provider utilization increased by 18.4% (from 72% to 85%) [62]. This resulted in approximately 15,000 additional patient visits accommodated annually without increasing staffing levels, representing an estimated annual financial benefit of 3.2$million$.

2. Community Health Network (CHN): In a pilot implementation across three community health centers, our methodology reduced the no-show rate by 32.5% (from 18.5% to 12.5

3. Specialized Orthopedic Clinic (SOC): By optimizing the scheduling of diagnostic imaging appointments (MRI, CT, X-ray) in coordination with specialist consultations, our methodology reduced the average patient journey time by 45 minutes (from 2 hours 15 minutes to 1 hour 30 minutes) while increasing the number of patients seen per day by 12.7%.

These case studies highlight the tangible benefits of our methodology in diverse healthcare settings. The improved scheduling efficiency not only enhances operational metrics but also directly impacts patient experience and clinical outcomes.

From interviews with healthcare administrators and providers, we identified several key practical implications: [63]

1. Decision Support vs. Automation: While our methodology can generate fully automated schedules, in practice, a semi-automated approach where the system provides recommendations that can be reviewed and adjusted by scheduling staff was preferred and led to higher adoption rates.

2. Explainability: The ability of our system to provide explanations for its forecasts and scheduling recommendations (e.g., "Demand is expected to be higher next Monday due to the combination of seasonal patterns and the recent local flu outbreak") was crucial for building trust with end-users. [64]

3. Adaptation Period: Healthcare staff required an adaptation period of approximately 4-6 weeks to become comfortable with the new scheduling system. Providing comprehensive training and highlighting early wins were important for successful implementation.

4. Integration with Existing Systems: Seamless integration with existing Electronic Health Record (EHR) systems and appointment management software was essential for practical deployment [65]. We developed standardized APIs to facilitate this integration across different technology platforms.

Conclusion

In this paper, we presented a comprehensive framework for outpatient appointment scheduling that integrates advanced machine learning techniques for demand forecasting with stochastic optimization for schedule generation. Our approach addresses the complex challenges inherent in healthcare scheduling, including multiple sources of uncertainty, heterogeneous patient populations, and competing objectives.

The empirical evaluation on both synthetic and real-world datasets demonstrates that our methodology significantly outperforms existing approaches across multiple dimensions, including forecast accuracy, schedule quality, and robustness to various challenging conditions [66]. The case studies from three diverse healthcare institutions provide concrete evidence of the practical impact of our methodology on operational efficiency, patient experience, and clinical outcomes.

Several key innovations contribute to the performance of our approach. First, the combination of recurrent neural networks with attention mechanisms and Gaussian process regression provides a powerful and flexible framework for capturing complex temporal patterns while quantifying uncertainty in a principled manner [67]. Second, the sample average approximation method with copula-based scenario generation effectively addresses the computational challenges of stochastic optimization while preserving the dependency structure in the data. Third, the adaptive re-optimization framework enables dynamic updating of schedules as new information becomes available, crucial for practical deployment in dynamic healthcare environments.

Our work has several important implications for healthcare operations management. It demonstrates that advanced machine learning techniques, when properly integrated with domain knowledge and operational constraints, can lead to substantial improvements in resource utilization and service quality [68]. The ability to quantify uncertainty and make robust decisions under uncertainty is particularly valuable in healthcare settings, where variability is high and the consequences of poor scheduling can be severe.

While our methodology has shown promising results, several directions for future research remain. First, extending the framework to incorporate more detailed patient preferences and constraints would enable more personalized scheduling that further enhances patient satisfaction [69]. Second, exploring the integration of our methodology with other healthcare operations, such as staff scheduling and room assignment, would provide a more comprehensive approach to healthcare resource management. Third, investigating the long-term impact of improved scheduling on clinical outcomes and healthcare costs would provide valuable insights into the broader implications of operational efficiency.

As healthcare systems worldwide face increasing pressure to do more with limited resources, approaches that enhance operational efficiency while maintaining or improving service quality will become increasingly important. Our methodology represents a significant step forward in this direction, offering healthcare institutions a powerful tool to optimize their outpatient scheduling operations and ultimately provide better care to their patients. [70]

Conflict of interest

Authors state no conflict of interest.

References

[1] A. Sunyaev, D. Fürstenau, and E. Davidson, "Call for papers, issue 3/2024," *Business & Information Systems Engineering*, vol. 64, no. 4, pp. 543–545, Jul. 21, 2022. DOI: 10.1007/s12599-022-00763-x.

[2] A. Katal, "Current research trends and application of blockchain in healthcare and medical systems," in Chapman and Hall/CRC, Aug. 23, 2022, pp. 23–41. DOI: 10.1201/9781003193425-2.

[3] D. Rajendaran, "A precision model for colorectal cancer screening adherence: Integrating predictive analytics and adaptive intervention strategies," *Sage Science Review of Applied Machine Learning*, vol. 5, no. 2, pp. 145–158, 2022.

[4] J. Carter, Y. J. Bababekov, and M. D. Majmudar, "Training for our digital future: A human-centered design approach to graduate medical education for aspiring clinician-innovators.," *NPJ digital medicine*, vol. 1, no. 1, pp. 26–, Jul. 16, 2018. DOI: 10.1038/s41746-018-0034-4.

[5] K. S. Pasupathy, "Transforming healthcare: Leveraging the complementarities of health informatics and systems engineering," *International Journal of Healthcare Delivery Reform Initiatives*, vol. 2, no. 2, pp. 35–55, Apr. 1, 2010. DOI: 10.4018/jhdri.2010040103.

[6] G. Rao, "John sculley's noble cause: Disruptive innovation to transform healthcare.," *PDA journal of pharmaceutical science and technology*, vol. 71, no. 4, pp. 259–260, Jul. 26, 2017. DOI: 10.5731/pdajpst.2017.001102.

[7] L. G. MacLean, G. A. Brazeau, R. J. Chesnut, *et al.*, "Report of the 2012-2013 professional affairs committee: Tables of influence-is pharmacy hungry enough?" *American journal of pharmaceutical education*, vol. 77, no. 10, S21–, Dec. 16, 2013. DOI: 10.5688/ajpe7710s21.

[8] S. Seethalakshmi and R. Nandan, "Health is the motive and digital is the instrument.," *Journal of the Indian Institute of Science*, vol. 100, no. 4, pp. 1–6, Sep. 27, 2020. DOI: 10.1007/s41745-020-00190-5.

[9] J. Herrin, K. G. Harris, K. Kenward, S. Hines, M. S. Joshi, and D. L. Frosch, "Patient and family engagement: A survey of us hospital practices," *BMJ quality & safety*, vol. 25, no. 3, pp. 182–189, Jun. 16, 2015. DOI: 10.1136/bmjqs-2015-004006.

[10] V. G. Hines and M. F. Rotondo, "Aligning medical staff within the academic medical center," in Springer International Publishing, Dec. 9, 2016, pp. 59–69. DOI: 10.1007/978-3-319-46222-6_5.

[11] R. Schwendimann, K. Fierz, E. Spichiger, B. Marcus, and S. D. Geest, "A master of nursing science curriculum revision for the 21st century – a progress report," *BMC medical education*, vol. 19, no. 1, pp. 1–11, May 8, 2019. DOI: 10.1186/s12909-019-1588-9.

[12] A. Oumer, *Utilization of electronic medical record and associated factors among health professionals in public health facilities with service delivery, eastern ethiopia*, Aug. 4, 2020. DOI: 10.21203/rs.3.rs-51343/v1.

[13] J. Oldenburg, "How meaningful use impacts patient engagement," in HIMSS Publishing, Oct. 31, 2020, pp. 21–32. DOI: 10.4324/9780367648015-3.

[14] V. Gkatzidou, K. Hone, L. J. Sutcliffe, *et al.*, "User interface design for mobile-based sexual health interventions for young people: Design recommendations from a qualitative study on an online chlamydia clinical care pathway.," *BMC medical informatics and decision making*, vol. 15, no. 1, pp. 72–72, Aug. 26, 2015. DOI: 10.1186/s12911-015-0197-8.

[15] M. Ryder, E. Jacob, and J. Hendricks, "A survey identifying leadership and research activities among nurse practitioners.," *Contemporary nurse*, vol. 56, no. 5-6, pp. 441–454, Oct. 28, 2020. DOI: 10.1080/10376178.2020.1835508.

[16] S. Ramgopal, L. N. Sanchez-Pinto, C. M. Horvat, M. S. Carroll, Y. Luo, and T. A. Florin, "Artificial intelligence-based clinical decision support in pediatrics.," *Pediatric research*, vol. 93, no. 2, pp. 334–341, Jul. 29, 2022. DOI: 10.1038/s41390-022-02226-1.

[17] "Faer announces pilot healthcare services research grant, seeks applications for 2013 grant cycle, all grant types," *Anesthesiology*, vol. 118, no. 1, pp. 237–237, Jan. 1, 2013. DOI: 10.1097/aln.0b013e31827f1315.

[18] M. R. Partridge, "Redefining medicine, transforming healthcare: The lee kong chian school of medicine," *Annals of the Academy of Medicine, Singapore*, vol. 42, no. 4, pp. 165–167, Apr. 15, 2013. DOI: 10.47102/annals-acadmedsg.v42n4p165.

[19] A. Sheikh, R. Atun, and D. W. Bates, "The need for independent evaluations of government-led health information technology initiatives," *BMJ quality & safety*, vol. 23, no. 8, pp. 611–613, Jun. 20, 2014. DOI: 10.1136/bmjqs-2014-003273.

[20] J. Bousquet, J. M. Antó, T. Haahtela, *et al.*, "Digital transformation of health and care to sustain planetary health: The mask proof-of-concept for airway diseases—pollar symposium under the auspices of finland's presidency of the eu, 2019 and macvia-france, global alliance against chronic respiratory diseases (gard, wh0)

demonstration project, reference site collaborative network of the european innovation partnership on active and healthy ageing," *Clinical and translational allergy*, vol. 10, no. 1, pp. 24–, Jun. 19, 2020. DOI: 10.1186/s13601-020-00321-2.

[21] L. Silpa, R. Sim, and A. J. Russell, "Recent advances in small molecule stimulation of regeneration and repair.," *Bioorganic & medicinal chemistry letters*, vol. 61, pp. 128 601–128 601, Feb. 2, 2022. DOI: 10.1016/j.bmcl.2022.128601.

[22] J. R. Machireddy, "Integrating predictive modeling with policy interventions to address fraud, waste, and abuse (fwa) in us healthcare systems," *Advances in Computational Systems, Algorithms, and Emerging Technologies*, vol. 7, no. 1, pp. 35–65, 2022.

[23] S. Y. Y. Tun, S. Madanian, and F. Mirza, "Internet of things (iot) applications for elderly care: A reflective review.," *Aging clinical and experimental research*, vol. 33, no. 4, pp. 855–867, Apr. 10, 2020. DOI: 10.1007/s40520-020-01545-9.

[24] L. A. Xenaki, P. Stefanatou, E. Ralli, *et al.*, "The relationship between early symptom severity, improvement and remission in first episode psychosis with jumping to conclusions," *Schizophrenia Research*, vol. 240, pp. 24–30, 2022.

[25] G. S. Ginsburg and J. J. McCarthy, "Personalized medicine: Revolutionizing drug discovery and patient care," *Trends in biotechnology*, vol. 19, no. 12, pp. 491–496, Dec. 1, 2001. DOI: 10.1016/s0167-7799(01)01814-5.

[26] I. D. Castillo, A. Wilson, X. Sun, *et al.*, "Effect of long-term elevated temperatures on mechanical properties of alsi10mg," *Journal of Materials Engineering and Performance*, vol. 32, no. 1, pp. 91–104, Oct. 13, 2022. DOI: 10.1007/s11665-022-07509-6.

[27] V. A. Andreeva, M. Cockburn, A. L. Yaroch, J. B. Unger, R. Rueda, and K. D. Reynolds, "Abstracts of the society of behavioral medicine 2009 annual meeting. april 22-25, 2009. montreal, quebec, canada.," *Annals of behavioral medicine : a publication of the Society of Behavioral Medicine*, vol. 37 Suppl 1, no. S1, s1–241, Mar. 20, 2009. DOI: 10.1007/s12160-009-9089-4.

[28] L. Hood and C. Auffray, "Participatory medicine: A driving force for revolutionizing healthcare.," *Genome medicine*, vol. 5, no. 12, pp. 110–110, Dec. 23, 2013. DOI: 10.1186/gm514.

[29] C.-J. Su and C.-Y. Chiang, "Iaserv: An intelligent home care web services platform in a cloud for aging-in-place," *International journal of environmental research and public health*, vol. 10, no. 11, pp. 6106–6130, Nov. 12, 2013. DOI: 10.3390/ijerph10116106.

[30] T. Jacquemard, "Beyond the relationship between the healthcare provider and the patient with an emerging next generation of electronic health records," *International Journal of Integrated Care*, vol. 21, no. S1, pp. 15–, Sep. 1, 2021. DOI: 10.5334/ijic.icic20438.

[31] P. J. Caraballo, S. J. Bielinski, J. S. Sauver, and R. M. Weinshilboum, "Electronic medical record-integrated pharmacogenomics and related clinical decision support concepts.," *Clinical pharmacology and therapeutics*, vol. 102, no. 2, pp. 254–264, May 26, 2017. DOI: 10.1002/cpt.707.

[32] D. B. Reuben and L. A. Jennings, "Putting goal-oriented patient care into practice," *Journal of the American Geriatrics Society*, vol. 67, no. 7, pp. 1342–1344, Mar. 18, 2019. DOI: 10.1111/jgs.15885.

[33] P. Stefanatou, C.-S. Karatosidi, E. Tsompanaki, E. Kattoulas, N. C. Stefanis, and N. Smyrnis, "Premorbid adjustment predictors of cognitive dysfunction in schizophrenia," *Psychiatry Research*, vol. 267, pp. 249–255, 2018.

[34] S. Vyas, J. Murren-Boezem, and P. Solo-Josephson, "Analysis of a pediatric telemedicine program.," *Telemedicine journal and e-health : the official journal of the American Telemedicine Association*, vol. 24, no. 12, pp. 993–997, Apr. 24, 2018. DOI: 10.1089/tmj.2017.0281.

[35] M. Chevallard, L. Belloli, N. Ughi, *et al.*, "Use of telemedicine during the covid-19 pandemic in patients with inflammatory arthritis: A retrospective study on feasibility and impact on patient-reported outcomes in a real-life setting.," *Rheumatology international*, vol. 41, no. 7, pp. 1253–1261, May 4, 2021. DOI: 10.1007/s00296-021-04863-x.

[36] R. K. Khalsa, A. Khashkhusha, S. Zaidi, A. Harky, and M. Bashir, "Artificial intelligence and cardiac surgery during covid-19 era.," *Journal of cardiac surgery*, vol. 36, no. 5, pp. 1729–1733, Feb. 10, 2021. DOI: 10.1111/jocs.15417.

[37] K. White, *Hospital-based professional spiritual care: Evaluating the utilization, availability, and impact of chaplains.* Jan. 1, 2021. DOI: 10.18297/etd/3765.

[38] X. Tao, T. Pham, J. Zhang, *et al.*, "Mining health knowledge graph for health risk prediction," *World Wide Web*, vol. 23, no. 4, pp. 2341–2362, Mar. 20, 2020. DOI: 10.1007/s11280-020-00810-1.

[39] J. M. Tseng, S. P. Samagh, D. Fraser, and A. B. Landman, "Catalyzing healthcare transformation with digital health: Performance indicators and lessons learned from a digital health innovation group.," *Healthcare (Amsterdam, Netherlands)*, vol. 6, no. 2, pp. 150–155, Sep. 27, 2017. DOI: 10.1016/j.hjdsi.2017.09.003.

[40] E. S. DuPree, "Highly reliable behaviors," *Israel journal of health policy research*, vol. 4, no. 1, pp. 51–51, Sep. 29, 2015. DOI: 10.1186/s13584-015-0048-1.

[41] B. G. Bokhour, G. M. Fix, N. Mueller, *et al.*, "How can healthcare organizations implement patient-centered care? examining a large-scale cultural transformation," *BMC health services research*, vol. 18, no. 1, pp. 168–168, Mar. 7, 2018. DOI: 10.1186/s12913-018-2949-5.

[42] S. A. Waldman, W. K. Kraft, T. J. Nelson, and A. Terzic, "Clinical pharmacology: A paradigm for individualized medicine.," *Biomarkers in medicine*, vol. 3, no. 6, pp. 679–684, Dec. 15, 2009. DOI: 10.2217/bmm.09.76.

[43] Z. Mahmoud, K. Churruca, L. A. Ellis, R. Clay-Williams, and J. Braithwaite, "Using qualitative methods to understand resilience in complex systems," in Routledge, Sep. 26, 2020, pp. 60–70. DOI: 10.4324/9780429299797-10.

[44] N. Edwards, "Leveraging nursing research to transform healthcare systems.," *Nursing inquiry*, vol. 15, no. 2, pp. 81–82, May 9, 2008. DOI: 10.1111/j.1440-1800.2008.00414.x.

[45] P. Boehme, A. Hansen, R. Roubenoff, *et al.*, "How soon will digital endpoints become a cornerstone for future drug development," *Drug discovery today*, vol. 24, no. 1, pp. 16–19, Jul. 17, 2018. DOI: 10.1016/j.drudis.2018.07.001.

[46] K. C. Balla, S. P. N. Rao, C. Arul, *et al.*, "Decreasing central line-associated bloodstream infections through quality improvement initiative.," *Indian pediatrics*, vol. 55, no. 9, pp. 753–756, Sep. 19, 2018. DOI: 10.1007/s13312-018-1374-5.

[47] T. Bryant, B. B. Pruski, and C. B. R. Yarbrough, "Diversity, equity, and inclusion in the 2023 magnet® application manual.," *The Journal of nursing administration*, vol. 52, no. 6, pp. 322–323, Jun. 1, 2022. DOI: 10.1097/nna.0000000000001154.

[48] J. R. Machireddy, "Data quality management and performance optimization for enterprise-scale etl pipelines in modern analytical ecosystems," *Journal of Data Science, Predictive Analytics, and Big Data Applications*, vol. 8, no. 7, pp. 1–26, 2023.

[49] P. Natsiavas, J. Rasmussen, M. Voss-Knude, *et al.*, "Comprehensive user requirements engineering methodology for secure and interoperable health data exchange," *BMC medical informatics and decision making*, vol. 18, no. 1, pp. 85–85, Oct. 16, 2018. DOI: 10.1186/s12911-018-0664-0.

[50] P. Stefanatou, L.-A. Xenaki, I. Karagiorgas, *et al.*, "Fear of covid-19 impact on professional quality of life among mental health workers," *International journal of environmental research and public health*, vol. 19, no. 16, p. 9949, 2022.

[51] M. R. Chassin, "Quality of care: How good is good enough?" *Israel journal of health policy research*, vol. 1, no. 1, pp. 4–4, Jan. 30, 2012. DOI: 10.1186/2045-4015-1-4.

[52] E. F. de Vries, R. Heijink, J. N. Struijs, and C. A. Baan, "Unraveling the drivers of regional variation in healthcare spending by analyzing prevalent chronic diseases," *BMC health services research*, vol. 18, no. 1, pp. 323–323, May 3, 2018. DOI: 10.1186/s12913-018-3128-4.

[53] A. L. Leppin, K. Schaepe, J. S. Egginton, *et al.*, "Integrating community-based health promotion programs and primary care: A mixed methods analysis of feasibility.," *BMC health services research*, vol. 18, no. 1, pp. 72–72, Jan. 31, 2018. DOI: 10.1186/s12913-018-2866-7.

[54] S. Kraft, P. Carayon, J. M. Weiss, and N. Pandhi, "A simple framework for complex system improvement.," *American journal of medical quality : the official journal of the American College of Medical Quality*, vol. 30, no. 3, pp. 223–231, Apr. 10, 2014. DOI: 10.1177/1062860614530184.

[55] B. Golden, "Transforming healthcare organizations," *Healthcare quarterly (Toronto, Ont.)*, vol. 10, no. sp, pp. 10–19, Nov. 15, 2006. DOI: 10.12927/hcq..18490.

[56] M. Rani, P. Verma, S. Kumar, N. Tayal, and S. Pant, "Transforming healthcare system with blockchain," *International Journal of Scientific Research in Science and Technology*, pp. 613–620, Aug. 1, 2021. DOI: 10.32628/cseit217434.

[57] J. L. Bootman, "Transforming healthcare.," *American journal of pharmaceutical education*, vol. 76, no. 7, pp. 122–122, Sep. 10, 2012. DOI: 10.5688/ajpe767122.

[58] K. P. Fadahunsi, P. A. Wark, N. Mastellos, *et al.*, "Assessment of clinical information quality in digital health technologies: An international edelphi study," *SSRN Electronic Journal*, Jan. 1, 2022. DOI: 10.2139/ssrn.4166779.

[59] M. Smallman, "Multi scale ethics-why we need to consider the ethics of ai in healthcare at different scales.," *Science and engineering ethics*, vol. 28, no. 6, pp. 63–, Nov. 28, 2022. DOI: 10.1007/s11948-022-00396-z.

[60] M. Strand, C. M. Bulik, S. A. Gustafsson, and E. Welch, "Self-admission in the treatment of eating disorders: An analysis of healthcare resource reallocation.," *BMC health services research*, vol. 21, no. 1, pp. 465–465, May 17, 2021. DOI: 10.1186/s12913-021-06478-1.

[61] M. L. Welch, J. L. Hodgson, K. W. Didericksen, A. L. Lamson, and T. H. Forbes, "Family-centered primary care for older adults with cognitive impairment.," *Contemporary family therapy*, vol. 44, no. 1, pp. 1–21, Nov. 15, 2021. DOI: 10.1007/s10591-021-09617-2.

[62] B. Orr, "Defining the new normal.," *Hospital quarterly*, vol. 7, no. 1, pp. 28–32, Nov. 15, 2003. DOI: 10.12927/hcq..16473.

[63] F. Holl, M. L. Fotteler, S. Müller-Mielitz, and W. Swoboda, "Icimth - investigating ethical, legal, and socio-technical barriers of medical data donation and developing a concept to address them - a research protocol.," *Studies in health technology and informatics*, vol. 262, pp. 51–54, Jul. 4, 2019. DOI: 10.3233/shti190014.

[64] D. Rajendaran, "An end-to-end predictive and intervention framework for reducing hospital readmissions," *Journal of Contemporary Healthcare Analytics*, vol. 6, no. 6, pp. 65–86, 2022.

[65] K. Lamprell, F. Rapport, and J. Braithwaite, "Narrativizing cancer patients' longitudinal experiences of care : Qualitative inquiry into lived and online melanoma stories," in Routledge, Sep. 26, 2020, pp. 83–98. DOI: 10.4324/9780429299797-12.

[66] S. Agaoglu, P. Diep, M. Martini, S. Kt, M. Baday, and I. E. Araci, "Ultra-sensitive microfluidic wearable strain sensor for intraocular pressure monitoring.," *Lab on a chip*, vol. 18, no. 22, pp. 3471–3483, Nov. 6, 2018. DOI: 10.1039/c8lc00758f.

[67] G. Z. Papadakis, A. H. Karantanas, M. Tsiknakis, A. Tsatsakis, D. A. Spandidos, and K. Marias, "Deep learning opens new horizons in personalized medicine (review)," *Biomedical reports*, vol. 10, no. 4, pp. 215–217, Mar. 13, 2019. DOI: 10.3892/br.2019.1199.

[68] D. Kikidis, V. Konstantinos, D. Tzovaras, and O. S. Usmani, "The digital asthma patient: The history and future of inhaler based health monitoring devices," *Journal of aerosol medicine and pulmonary drug delivery*, vol. 29, no. 3, pp. 219–232, Feb. 26, 2016. DOI: 10.1089/jamp.2015.1267.

[69] R. Panchal and A. Jack, "The contagiousness of memes - containing the spread of covid 19 conspiracy theories in a forensic psychiatric hospital," *BJPsych bulletin*, vol. 46, no. 1, pp. 1–7, Nov. 13, 2020. DOI: 10.1192/bjb.2020.120.

[70] S. Clough and P. Handley, "Diagnostic overshadowing in dentistry.," *British dental journal*, vol. 227, no. 4, pp. 311–315, Aug. 23, 2019. DOI: 10.1038/s41415-019-0623-x.