

Develop an AI-Driven Pharmacy Optimisation Platform for Enhanced Operational Efficiency and Cost Reduction in Healthcare Systems

Aryav Odhrani

Abstract

This project investigates how artificial intelligence (AI) and automation can be applied to optimise inventory and vendor management in hospital pharmacies, with a particular focus on underserved regions of India. It begins with a critical review of existing pharmacy management systems, identifying current solutions in literature and their key inefficiencies. In collaboration with two charitable hospitals in Gujarat and Rajasthan, two complementary systems were developed: a Pharmacy Optimisation Management System (POMS) and a Pharmacy Vendor Management System (PVMS). The POMS prototype uses real-world hospital data to forecast demand, classify medicines by stock-out and expiry risk, and generate procurement quantities, while the PVMS automates vendor data processing, vendor comparison, prescription analysis, and order generation. When deployed in the pilot hospitals, the models are predicted to reduce manual labour by approximately 25,000 hours annually, corresponding to an estimated cost saving of \$35,000. The systems currently manage over 6,000 unique pharmaceutical items, improving essential medicine availability and accessibility for approximately 25,000 patients each month. The findings demonstrate that AI-driven pharmacy management systems can meaningfully improve operational efficiency and cost reduction in healthcare settings.

Table of Contents

1	INTRODUCTION	2
1.1	AIMS	2
1.2	RATIONALE	3
1.3	TERMINOLOGY	4
2	RESEARCH REVIEW	5
2.1	CHALLENGES	9
2.2	SOLUTIONS	11
2.3	RESOURCE REVIEW	18
3	MODEL	23
3.1	METHODOLOGY	23
3.2	PHARMACY OPTIMISATION MANAGEMENT SYSTEM	24
3.3	PHARMACY VENDOR MANAGEMENT SYSTEM	31
4	CONCLUSION	33
5	EVALUATION	34
6	REFERENCES	37

1 Introduction

Hospital pharmacies play an important role in today's healthcare systems, serving as the conduit between the diagnosis and prescription of doctors and the consumption of medication by patients. Pharmacies enable the delivery of medicines, ranging from nonessential and seldom-used medications to lifesaving and vital treatments, to patients. However, today, healthcare facilities face challenges providing adequate services to their patients, including struggles with stockouts, overstocking, and wastage (Akter, 2023). These financial and operational inefficiencies cause significant negative impacts on pharmacies from a business perspective.

In the next section, we review the landscape of healthcare and the involvement of information technology, both globally and specifically within India. We study existing pharmacy management systems and identify major challenges in the sector. We then analyse and evaluate potential solutions to issues in present systems. In the following sections, we contribute a novel model that assists pharmacies enhance their inventory management systems. We pilot the model in two hospitals in India: Monilek Hospital and Research Centre, Jaipur, Rajasthan, and Shrimad Rajchandra Hospital and Research Centre, Dharampur, Gujarat. We discuss the effectiveness of our approach to optimisation, specifically in improving operational efficiency and reducing costs. Finally, we discuss potential for further study in this area.

1.1 Aims

This project begins with an evaluation of existing problems and solutions to pharmacy management systems and the resulting distribution of critical medicines in the Indian healthcare economy. Subsequently in the project, we aim to develop a prototype for a platform-agnostic, AI-driven pharmacy optimization system designed to improve the daily operations of charitable hospitals. In the process, I aim to understand the technological requirements involved in creating such a model and to evaluate how automation and the use of AI in decision-making can improve cost efficiency and accessibility to essential medicines. By coding a functional prototype using real-world data, I will be able to assess the scalability and potential impact of my model. This will be done through two key metrics: cost savings and manual labour savings, which can be assessed with quantitative metrics provided by client hospitals. Additionally, my evaluation of impact would be supported by two further metrics, one determining the number of unique that

my system caters to, hence provides patients with increased and more efficient access to, and another determining the number of patients who benefit from the system each month.

By identifying the key challenges and inefficiencies in existing systems, specifically their inventory and vendor management features, and by then building a working proof-of-concept that solves these problems, I hope to contribute to my final goal, outlined here. My ulterior aim is to improve efficiency in healthcare distribution and increase access to critical medicines in underserved regions of India. This typically refers to rural areas where citizens have lower per capita income, high poverty rates, and inadequate social infrastructure (including a low availability of hospitals). As a result, I have selected two hospitals in different states of India, each facing a slightly different range of patients within the socioeconomic scale. In the next section, I will go over my rationale behind the project, explaining why I have chosen healthcare distribution as my goal and why the two pilot hospitals as my initial clients and collaborators.

1.2 Rationale

The identification of an issue in India's healthcare distribution inefficiency and the struggle faced by the rural population began when I was volunteering at a sprawling medical camp hosted by a charitable hospital in Gujarat. I was serving medicines and treatments free of cost to the five thousand rural poor who had lined up for critical care. Dozens of stalls distributed portable ECG machines, glucometers, blood pressure monitors, and ran check-ups on everything from chronic back injuries and diabetes to minor surgeries and cataract removal. Stationed in the pharmacy section, as I distributed boxes of antibiotics and antihypertensives, I found it ironic and perplexing that despite India being the world's largest producer of generic medicines, critical drugs could run out simply because of inefficiencies in distribution and inventory management. Stockpiles that should have lasted days vanished in hours. This struggle, faced by the thousands in front of me and millions nationwide, was a systemic failure.

Speaking directly to patients made it exponentially more personal. None of them complained; they had normalised struggle, and that bothered me to the core. The drugs and funding existed, but the system of getting medicines from suppliers to shelves to patients was broken. Unlike in developed countries like the United Arab Emirates, inefficiency wasn't an inconvenience, rather, it was the difference between recovery and relapse.

According to the Centers for Disease Control and Prevention, 6 in 10 United States adults have at least one chronic disease, and 4 in 10 have two or more (Market Data Forecast, 2025). Given

that the rate of healthcare issues is a majority of the population in developed countries with relatively widespread healthcare accessibility like the U.S., there would be significantly high levels of disease in India and an extreme inadequacy of services to support the millions struggling.

Thus, I decided to collaborate with the hospital running the medical camp: Shrimad Rajchandra Hospital (henceforth referred to as SRH). Located in the town of Dharampur, in the Valsad district of Gujarat, an Indian state, SRH is a charitable hospital that is operated by the UN-affiliated non-governmental organisation Shrimad Rajchandra Love and Care (SRLC). The multi-speciality hospital was inaugurated by the Prime Minister of India, Shri Narendra Modi, himself, and has treated over 1 million patients since its original inception (Shrimad Rajchandra Hospital, 2025). I was thoroughly impressed by the concept of SRH, a state-of-the-art charitable institute providing healthcare services of the highest quality to the surrounding areas, the underprivileged rural populations of South Gujarat. Having volunteered with SRLC's various other charitable activities since childhood, I deemed SRH to be an effective means to create positive social impact in the region. After personally touring the facility, including their (outpatient) Raj Pharmacy and interacting with their doctors on site, my confidence was reinforced, and I finalised SRH as the first pilot hospital for my project's pharmacy optimisation system.

The second pilot hospital I have chosen is Monilek Hospital and Research Centre, founded in 1986 in Jaipur, Rajasthan. Monilek is similar to SRH in being another charitable, multi-speciality hospital that provides free and subsidized healthcare treatments to patients. With a strong social service mission, Monilek was the first independent hospital in the state of Rajasthan to initiate a long-distance healthcare program through mobile medical and surgical units and have since hosted many medical camps for low-income patients in multiple states (Monilek Hospital, 2025). I have been to Monilek's Jaipur facility several times and was impressed by their Outpatient Clinic (pharmacy), so finalised Monilek as the second pilot hospital for my project.

In the process of building this project, I have developed strong relations with leading members of both SRH's and Monilek's Pharmacy and the charitable foundations running them, and in hindsight, the decision to collaborate with these two hospitals has been a great one.

1.3 Terminology

Active ingredient: The pharmacologically effective component of a medicine responsible for its therapeutic action.

Demand forecasting: The process of estimating future medicine demand based on historical consumption data, trends, seasonality, and contextual factors, in order to inform procurement and stock planning decisions.

Forecast accuracy: A measure of how closely predicted demand aligns with actual observed demand, commonly used to evaluate and refine forecasting models.

Inventory: the stock of pharmaceutical products retained to meet future demand (Ali, 2011).

Inventory control: the process of managing inventory in order to meet customer demand at the lowest possible cost and with a minimum investment (Rachmania & Basri, 2013).

Operational workflow: The sequence of processes through which data and decisions move within a system, from input ingestion to final output.

Optimisation: The application of analytical or computational methods to determine the best possible outcome under given constraints, such as minimising cost while maintaining required service levels.

Overstock: A condition in which inventory levels exceed operational requirements, increasing the risk of expiry, wastage, and unnecessary capital expenditure.

Procurement: The process of sourcing and purchasing medicines from vendors or suppliers based on demand, pricing, availability, contractual terms, and regulatory considerations.

Reorder point: The predefined inventory level at which a new order should be placed to replenish stock before it reaches a critical shortage.

Reorder quantity: The calculated quantity of a medicine that should be ordered when replenishment is triggered, often determined using demand forecasts and safety margins.

Stockout: A situation in which a required medicine is unavailable in inventory at the time of demand, potentially disrupting patient care.

Vendor: An external supplier or distributor that provides pharmaceutical products to a pharmacy or healthcare institution.

Vendor management: The systematic evaluation, selection, monitoring, and optimisation of vendor relationships to ensure reliable supply, competitive pricing, and consistent quality.

2 Research Review

There is not a single industry in the world that would survive without the healthcare industry. The global economy is heavily dependent on healthcare markets because the industry provides

services that treat illnesses, furthers research in curing diseases, and develops new technology to enhance the state of living of people worldwide. It is not understatement to say that healthcare holds critical importance to supporting life and enhancing our contemporary lifestyles. Thus, healthcare industries are some of the biggest in their respective economies. In this United States, spending on healthcare exceeded \$4.5 trillion annually as of 2024 (Market Data Forecast, 2025), which is almost a fifth of total GDP, and is projected to reach over \$7.5 trillion within the next decade. The US healthcare industry is a vast ecosystem with over 15 million healthcare professionals and the industry's largest firms generate hundreds of billions of dollars in revenue, demonstrating the scale and importance of healthcare globally.

India, by contrast, has some of the lowest healthcare spending as a percent of GDP, amounting to only 1% of GDP in 2017 (Jaffrelot & Jumle, 2020). The population of India grew by 13% in the previous decade (160 million people), but healthcare spending only rose by 0.29% to be 1.29% of GDP in 2019-2020, according to the National Health Profile 2019, Government of India (Mukherjee, 2025). Despite this, the private sector of healthcare has risen during the same period, largely due to the growth and rising affluence of its 400-million-large middle class. Private infrastructure thus accounts for 62% of the industry, with 43,000 hospitals compared to the 25,000 public hospitals funded by the taxpayer (Jaffrelot & Jumle, 2020). However, it is important to note that in a country as vast and diverse as India, such statistics are not representative of the microcosms of small communities, and even individual states. For example, the states of Bihar and Maharashtra have similar sized populations, with censuses varying being 100 million to 130 million, and a comparable number of hospitals – around 3,000. However, the capacity of Bihar's hospitals stands at 30,000 compared to Maharashtra's 230,000! Out of those 30,000, only 11,000 are public beds, so there are only 0.09 beds per 1,000 people in Bihar (Jaffrelot & Jumle, 2020), demonstrating a severe disparity in the country's healthcare system.

Amongst other initiatives launched to improve Indian healthcare, a standout government scheme is the Ayushman Bharat Digital Mission (ABDM). Previously called the National Digital Health Mission, ABDM aims to create an integrated healthcare system linking India's practitioners with patients by digitizing patient data, prescriptions, and overall pharmacy operations. The Ministry of Health & Family Welfare's Health Management Information System (HMIS) is a web-based monitoring information system that supports this mission effectively (Ministry of Health and Family Welfare, Government of India, 2025). For example, within the first two weeks of its inauguration, the HMIS processed over 89,000 patients in the Surat Municipal

Corporation's 30 urban health centres (Times of India, 2025), significantly helping both hospital managements and patients. The system provides readily available data to support planning on the behalf of hospital management, leading to better health services for patients. In addition to pharmacy management, the HMIS includes digital registration, electronic health records, operation theatre and facility management, and detailed records of diagnostic tests. A Picture Archival and Communication System stores radiological images (like X-rays), and a Queue Management System provides token-based patient handling. On the patient-side, an e-Health mobile app allows patients to access their health records, including blood reports, prescriptions, and more, on their mobile phone. Government schemes such as ABDM and the HMIS have led to tracking and management systems like a Unique Health ID (UHID) (Times of India, 2025) that are unified across hospitals, streamlining services, and increasing efficiency.

There is a wide range of separate, company-provided inventory management tools that are used by hospitals in India. These offer point of sale functions, invoicing, and stock tracking, but adoption is limited due to cost, fragmented ecosystems, and training gaps. A study by Healthray outlines key features in AI-powered pharmacy software systems, including demand forecasting, automated order placement, real-time inventory tracking, and expiration date monitoring (Balar, 2025). These theoretically enhance medicine availability and prevent the running out of essential drugs, but such systems seldom incorporate advanced functionality such as predictive analytics or vendor similarity matching. Added to the challenges of up-front costs, integration with existing systems, and concerns about data security (Balar, 2025), systems are currently not effectively implemented. Another challenge we will discuss later is workforce digital literacy and staff training, and its impact in causing the relatively low presence of such AI-powered systems in the current market.

We have seen that India's healthcare industry is a large and scattered, with thousands of hospitals treating millions of patients in their own ways. The quality of service and care provided varies across state, district, and town, and there is no single, centralised system as of yet, although the government has attempted to initiate some through their schemes. As a result of this status, with Indian healthcare being an organised mess, we wish to create a platform-agnostic pharmacy management system: one that is generic to hospitals nationwide and can thus be easily scaled, increasing its potential impact. To achieve this, we must enable the platform to run on any hospital's dataset, rather than being trained for a specific one. Thus, we will initially train our models on data from our pilot hospitals, before creating a generic system incorporating an ingestion feature to absorb and analyse data from any client.

We now come to review inventory management in pharmacy practice, including related terms and their specific details. Inventory management can ensure patient safety as a result of software systems flagging expired or substandard products. It is also important to consider that stock management is important to hospitals and pharmacies because inventory is usually their largest current asset and liquid asset. This is especially true as the variety and cost of pharmaceutical products rise over time (Ali, 2011). There are two key perspectives that support this viewpoint.

First, from a financial standpoint, efficiently managing inventory will lead to higher gross profits and net profits for firms but reducing the cost of pharmaceutical products and their associated operational expenses. Since profits are defined as revenue less costs and expenses, lower costs as a result of better management will increase profits for firms. For most profit-motivated businesses, this meets their objectives, albeit even for non-profit organisations, including charitable hospitals like SRH and Monilek, higher profits enable the firm to expand their output, meeting their objective to increase the impact of their charitable activity. Since money is saved in the process of purchasing and storing products, better management also lead to better cash flow, which can be used to invest in higher quality services for the firm.

Second, from an operational standpoint, effective inventory management improves the quality of service provided to the customer, i.e., the patient. Having stock of medicines ensures customer happiness and provides them with utility, whereas in contrast, unavailability of a product would adversely affect patients (Ali, 2011). In our scenario, where there is a lack of alternative hospitals in underserved regions of India, this negative affect would be extreme, potentially leading to loss of life if critical medicines are unavailable. Additionally, unavailability of a product causes an inconvenience to the doctors and physicians prescribing the medicine, adding hindrance to their efforts to serve low-income patients.

There are four main costs associated with pharmacy inventories: (1) acquisition costs, which is the amount paid to vendors to buy the product, (2) procurement costs, associated with the purchasing of the product, including the placement of orders, receiving stock, and paying invoice fees, (3) carrying costs, which are costs related to transporting and storing the product, including overhead costs of air-conditioning a warehouse facility to prevent spoiling of drugs, for instance, and costs incurred as a result of crises, such as theft or damage, and finally (4) shortage costs, also known as stock-out costs (Ali, 2011).

The most practical method to manage an inventory is to calculate the inventory turnover rate (ITOR) for various ranges of products, such as an entire firm, a specific department, or even a single product. ITOR is the total cost of products sold, using the above four costs, divided by the average of the beginning and ending inventory values in the desired period of time (Ali, 2011). For example, if a pharmacy stores an average of \$100,000 of inventory in a year, and has sold products costing \$1 million, then their ITOR value is 10, so they have sold their entire inventory 10 times in the year. Higher ITOR values indicate better management and quick rate of purchase, storage, and sale. However, given that inventory setups vary with different pharmacies, these statistics should be considered relative to a firm's own previous values rather than others, concerning how well the firm has improved its management over time. A lower ITOR indicates poor management because a long held stock suggests that the incorrect quantities or products have been ordered compared to the demand received.

Another important factor is the percent net profit (PNP) of a firm, which compares net profit to average inventory as a percentage ratio (as opposed to comparing cost/revenue to average inventory in ITOR). Higher PNPs overall indicate greater profit margins. The relationship between a pharmacy's ITOR and PNP over time can reveal significant information about ongoing trends. For instance, in a situation where ITOR is rising in a given time period, but PNP is falling, one can infer that more products were sold with lower-markups, since profits are lower (Ali, 2011).

2.1 Challenges

Contemporary studies on the healthcare industry have identified four primary challenges related to inventory control: inventory inefficiencies and related high costs, lack of predictive and real-time demand forecasting, fragment vendor and procurement processes, and limited automation and safety measures. These are accompanied by several other, albeit less serious, problems, not limited to barriers to implementation - such as digital literacy and data readiness - and regulatory complexities.

The greatest and most obvious challenge regarding inventory control is inventory inefficiencies and following cost waste. According to the World Health Organization (WHO), 20%-30% of all healthcare resources in the world are lost due to inefficiencies in systems managing inventories and their supply chains (Akter, 2023). Given that one-third of the average hospital's annual budget is spent on buying medicinal materials and supplies (Jaju, et al., 2023), the lack of

effective management consumes significant money. Many pharmacies, which are often the most extensively used ‘therapeutic facilities’ at a hospital, have inventories that are manually managed. This increases inefficiency, since human staff spend considerably more hours tracking stock levels than an automated system. The limits of human ability also result in a lack of real-time visibility on the medicines at all times, compared to computer systems.

Second, a lack of effective predictive and real-time demand forecasting tools worsens the aforementioned challenges. Quantities of ordered medicines are often wrong, so low-demand medications may be overstocked and would eventually expire, resulting in wastage. Globally, studies have shown that around one-quarter of all medicines suffer stockouts (from understocking) and one-sixth suffer expiry (from overstocking) (Jaju, et al., 2023). It is quite easy to understand that, because India’s healthcare economy is very large, with thousands of hospitals and independent pharmacies all over the country, there is no wholistically trained, accurate model to predict trends in nationwide medicine demand. While basic models are commonly used, including EOQ and ABC, which we will discuss later, these aren’t effective enough to reduce stockouts and overstocking on the scale required. Thus, I intend to build a model using real data from my two pilot hospitals, in a platform-agnostic system, before scaling it to be used for any client pharmacy.

Another frequently observed problem was rooted in the process of pharmacies procuring medicines from wholesale vendors. The process is defined by the fact that multiple vendors offer the same active ingredients in their medicines, with the only differences being in the brand name, price, and quality. This fragmented system has no universally implemented automation system that compares equivalent products or suggests best-value vendors for a client pharmacy. Indeed, a study done on a National Institute in East India stated that key challenges arose from “rural location, communication hurdles, and vendor management” (Jaju, et al., 2023). In some of the pharmacies that I personally visited, staff would manually review offers from different vendors, which were often in varying formats, such as spreadsheets and contract documents, resulting in slow processing time. Resulting from this human effort, delays and errors are common. Another option is for pharmacies to use NLP-enhanced entity parsing, but this is not commonly used and not the most efficient solution.

The fourth main issue revolved around limited automation, especially in low-resource hospitals, which are often those in rural areas, or those set up as charitable institutions. The lack of

barcoding, automated dispensing cabinets, and Computerized Provider Order Entry (CPOE) has led to medication error levels remaining high at such hospitals. There's a clear need for a low-cost, scalable system that offers the same safety gains as the aforementioned automation systems using simple software solutions, rather than expensive, capital-heavy structures.

Other barriers to the implementation of management systems that improve efficiency include inadequate staff training, inconsistent data and processing formats (of vendors, medicines, finances, and more), and overall limited digital infrastructure in pharmacies. Infrastructure limitations prevent digital adoption of existing frameworks due to unreliable internet, electricity issues, and lack of hardware. Small and charitable hospitals that were to invest in such infrastructure would face high upfront and ongoing costs, with no guaranteed ROI, disincentivizing investment. We solve these financial constraints by providing our system free to client hospitals, aiming to maximize social benefit.

Furthermore, AI models require training on data, and rural pharmacies tend to not properly store historical data. Even if that is not the case, such many pharmacies maintain a small scale of operation, with small datasets that are ineffective in training deep learning models. This data availability issue in under-resourced hospitals means that the various frameworks used at top private hospitals cannot be adapted to low-tech environments.

In situations where digital tools have been implemented by hospitals, a significant proportion of healthcare workers struggle to operate them, leading to lower productivity and thus, an ineffective system. Workers often underutilize systems due to inadequate digital literacy and technological adaptability in the rural workforce. The implication of the discussed challenges is that, for true real-world feasibility, any system must be simple, user-friendly for both hospital workers and patients, and able to operate in constrained environments, requiring minimal infrastructure.

2.2 Solutions

As we have seen in the previous section, there are various challenges faced by firms aiming for pharmacy optimization, centred around weaknesses that prevent existing systems from being widespread and effective means of cost reduction and management efficiency. In this section,

we explore existing solutions and their advantages and disadvantages in the context of pharmacy operational optimization.

A popular method to running a business, particularly relating to the manufacturing of goods, is the Lean Six Sigma methodology. This is a combination of two distinct methodologies: namely, Lean and Six Sigma. The former of the two is the lean manufacturing method, developed as the *kaizen* business model in Japan, meaning “improvement”. *Kaizen* revolves around business optimization occurring from the cumulative effect of repetitive, continuous acts of small improvement in production. Thus, the system aims at improving all parts of the company, from the CEO to assembly line workers, and from purchasing to supply chain (Wikipedia, 2025). The roots of the system can be traced to Taiichi Ohno, who created the Toyota Production System in the mid 20th-century, aiming to achieve optimization through just-in-time manufacturing. *Kaizen* has developed over decades to include many sub-types and applications, including the 5S, 7M, 7W, and 3 Mu checklists. However, the main aim is a company the reduction of the seven key *Muda* (wastages). These are losses due to: overproduction, idle waiting time, unnecessary transportation, over-processing the product, excessive storage (of inventory), unnecessary motion, and defects (production of faulty parts) (Chittenden, n.d.). In the 1980s, former Waymo CEO John Krafcik, then a Masters student at MIT, coined the term “Lean” to describe the Toyota Production System’s efficiency.

Interestingly, the Six Sigma set of techniques and tool for process improvement, was created to compete with the Japanese lean method. The name refers to the goal of the strategy’s pioneering company, Motorola, to achieve a level of quality control such that their defect rate was limited to the fraction of the normal curve lying outside of six standard deviations of the mean. i.e., 3.4 defects per million, or an outstanding 99.99966% yield. The method builds upon *Kaizen*’s ideologies of continuous efforts leading to sustained quality improvement but differs by emphasising on identifying and removing causes of defects and reducing variability in production processes (hence the reference to statistical control and the normal curve). It also stands apart for focusing on measurable financial returns and emphasising management leadership and support (Wikipedia, 2025). Various project-specific methodologies include DMAIC (Define, Measure, Analyse, Improve, Control) and DMADV (Define, Measure, Analyse, Design, Verify). Within each phase of a DMAIC or DMADV project, Six Sigma adopters, some of the first of which were Honeywell and General Electric, use many quality-management tools, ranging from

axiomatic design, Ishikawa diagrams, and CTQ trees, to SIPOC analysis, Taguchi methods, and value-stream-mapping (Wikipedia, 2025).

During the 2000s, Lean Six Sigma developed as its own specific process of manufacturing, combining Lean management and the Six Sigma method. The methodology has since expanded into healthcare, finance, and supply chains. Mark George's book, "The Lean Six Sigma Guide to Doing More with Less" expounds how to "cut costs, reduce waste, and lower your overhead" with detailed explanations of key process and enterprise cost reduction methods, as discussed above. Amongst others, the book refers to the *Kaizens* and their seven *Muda*, measuring process efficiency through PCE, Little's Law, and WIP Cap methods, as well as practical means to deploy the Six Sigma method in enterprises (George, 2010).

We now focus on two specific cross-sectional, exploratory studies conducting ABC-VED matrix analysis on pharmacies in Ethiopia and India, as well as investigating various other problems and solutions to inventory control. The first study was conducted across 15 public health facilities in West Shewa region of Ethiopia from September 2015 to September 2018 (Deressa, et al., 2022). The second study was conducted in the All India Institute of Medical Sciences, Deoghar, India, for the financial year of 2022-2023 (Jaju, et al., 2023).

One proposed solution was the adoption of integrated ABC and VED analysis, referring to the "Always Better Control" and "Vital Essential Desirable" classification methods to inventory management, respectively. Separately, these methods may lack depth to provide useful information for making decisions regarding the change of which medicines to purchase and sell. ABC focuses on the monetary value of the drug and fails to consider its essentiality, while considering VED alone, costly non-essential drugs may be treated as a priority, leading to wasteful expenditure (Deressa, et al., 2022). However, integrated them in an ABC-VED matrix has shown promising gains in pharmacy optimization, overcoming these limitations (Jaju, et al., 2023).

First, in ABC analysis, following the Pareto ideology of "separating the vital few from the trivial many" (Jaju, et al., 2023), each drug is given a score relating to its annual value of usage. This is calculated by multiplying its consumption volume by its cost, thus giving it a score impacted by both clinical and economic impact. The drugs are then arranged into three groups based on their scores: the first 70% in the A group, the next 20% in the B group, and the final 10% in the C group.

Second, VED analysis compares the criticality of items to a business's operations, where the Vital group contains high-risk items, such as critical machines and life-saving drugs that are absolutely essential to a firm, the Essential group is for important items where minor disruptions are manageable, and the Desirable group consists of items that improve efficiency but have no immediate impact if are absent (Deressa, et al., 2022). It is worth noting that the Desirable group is sometimes referred to as the Non-essential group, and thus this method may also be called VEN analysis. Similar classifications are SDE (Scarce Difficult Easy) analysis and FSN (Fast Slow Non Moving] analysis for when different inventory features need to be assessed.

After conducting the two analyses, the firm must create an ABC-VED matrix to prioritise different drugs that are ordered and stored at pharmacies: category I consists of all drugs belonging to AV, AE, AD, BV, and CV subcategories, category II consists of BE, CE, and BD subcategories, and all other cheaper and desirable drugs in the CD subcategory form category III. Category I drugs must be monitored and controlled continuously; category II drugs must be controlled periodically; and no control is required for category III drugs, saving costs and time (Deressa, et al., 2022).

The second study also proposed implementing more indirect solutions to operational efficiency, such as integrating electronic health records and bar code technology in inventory management systems and using first-in first-out (FIFO) and just-in-time (JIT) methods (Jaju, et al., 2023). While this second study suggests a potentially effective method of analysis in the ABC-VED matrix, there are several issues that prevent it from being a perfect solution: their participant number was low and the problems their pharmacists faced "might not reflect the actual scenario of a rural setup" (Jaju, et al., 2023).

Another solution is the Economic Order Quantity (EOQ) method. A study conducted in a public hospital in Indonesia (Rachmania & Basri, 2013) discusses the EOQ formula, which was developed by F.W. Harris in 1915. It calculates the optimal lot size of purchasing different medicines for a specific medicine to minimize total operating cost. $EOQ = \sqrt{\frac{2DC}{H}}$, where D is the annual demand (units), C is the cost per order for that medicine, and H is the annual per unit holding cost (Anusha, 2025). Benefits of using this system include a balanced inventory with fewer gaps and stockouts during peak periods, reduced warehousing and storage costs, and reduced product wastage, ultimately leading to better financial outcomes for the pharmacy (Anusha, 2025). However, the EOQ method has many downsides. Most importantly, its classical

model assumes many constants, including demand, lead time, and order cost per order. It also assumes replenishment is instantaneous, neglecting transport time and costs, as well as the possibility of uncertainty in demand and quantity discounts (bulk purchasing) (Rachmania & Basri, 2013).

Pharmacies can also implement two types of replenishment processes: continuous and periodic reviews. The former involves a system that places a re-order of a medicine when its stock level declines to meet a certain threshold, the re-order-point (ROP), and the latter places orders at regular periodic intervals (Rachmania & Basri, 2013). Safety stock must also be considered when setting replenishment systems, as this extra stock can be used during replenishment lead time if there is actual demand turns out to be higher than the expected demand, preventing a stockout (Rachmania & Basri, 2013).

We now come to discuss a major solution: demand forecasting. This can refer to a wide range of adaptive forecasting techniques, whereby historical data is analysed to determine future demand. This has its own set of challenges, including varying demand trends for different drugs (including shocks such as the outbreak of a pandemic), and the requirement of adequate data. Various brands that sell drugs with the same active ingredients add another layer of complexity to the mathematical modelling of demand, as it adds client and physician biases, as well as varying quality and costs.

Implementations of demand forecasting can be built upon one another, leading to increasingly expansive models. A pharmacy can begin with analysing simple exponential smoothing. Exponential smoothing of time series data is when older data is given exponentially less weighting and priority than newer, more relevant observations, allowing for short-term forecasts to be made more accurately. The basic formula for this is as follows:

$$S_{t+1} = \alpha y_t + (1 - \alpha)S_t,$$

where α is the smoothing constant, a value from 0 to 1 (Christopher, 2021). There can then be the additions of more equations, leading to double or triple exponential smoothing, which considers trend and seasonality factors.

For instance, a pharmacy can utilise a Holt's model, which is trend-corrected or double exponential smoothing, or subsequently Winter's model, which is a trend-seasonality-corrected or triple exponential smoothing. The study on Indonesian public hospitals compares these four forecasting methods with three measures of forecasting error for a specific set of drugs: the mean average deviation (MAD), mean absolute percentage error (MAPE), and the tracking signal (TS). Tracking signal, calculated as the ratio of cumulative forecast error (RSFE) and MAD, identifies biases in a model, indicating if it consistently over- or under-predicts demand (Top Consulting Firms, 2025). A tracking signal close to 0 is the best. The study showed that the TS range for every product was in the ± 6 range, suggesting that all four models were relatively accurate and thus acceptable. Holt's model was found to have both the smallest MAD and MAPE for 5 out of the 6 drug products studied, leading to the paper concluding that it was the most appropriate model out of the four for medication demand forecasting (Rachmania & Basri, 2013).

We thus look deeper into Holt's model to study why it was successful. Created in 1957, it extends simple exponential smoothing to forecast data with a trend (Hyndman & Athanasopoulos, n.d.). There are three equations involved, for forecasting, smoothing level, and smoothing trend. The equations are as follows:

$$\begin{aligned}\hat{y}_{t+h|t} &= l_t + hb_t, \\ l_t &= \alpha y_t + (1 - \alpha)(l_{t-1} + b_{t-1}), \\ b_t &= \beta(l_t - l_{t-1}) + (1 - \beta)b_{t-1},\end{aligned}$$

where $\hat{y}_{t+h|t}$ denotes the predicted output at h time-steps ahead of a time t , given a historic trend of t , l_t denotes the estimate of the level of the series at time t (which is its actual value), b_t denotes the estimate of the trend of the series as time t (which is the gradient, or rate of change of the actual value), and α, β are the smoothing parameters where $0 \leq \alpha, \beta \leq 1$. The second equation shows that the l_t is a weighted average of the previous forecast at time t added to a weighted $l_{t-1} + b_{t-1}$ term, allowing for a training forecast for the next step. Similarly, b_t is the weighted average of the actual trend added to a weighted b_{t-1} term, which is the previous time-step's estimate of the trend (Hyndman & Athanasopoulos, n.d.). In 1985, Gardner and McKenzie added a damping parameter to the Holt method in order to trend forecasts to a flat line at a large time in the future, since Holt's existing linear method would often result in an increasing trend into the indefinite future, which is unrealistic (Hyndman & Athanasopoulos, n.d.).

An alternative to double and triple exponential smoothing is the double moving average (DMA) demand forecasting method. Similar to exponential smoothing, this builds upon a single moving average. The DMA forecast involves calculating a single moving average S'_t on the original series, and then a single moving average S''_t on the same smoothed series. These are then used to find level parameters a_t and trend parameters b_t to forecast F_{t+m} at a time m periods ahead. The equations are as follows:

$$S'_t = \alpha D_t + (1 - \alpha)S'_{t-1},$$

$$S''_t = \alpha S'_t + (1 - \alpha)S''_{t-1},$$

$$a_t = S'_t + (S'_t - S''_t),$$

$$b_t = \frac{\alpha}{1-\alpha} (S'_t - S''_t),$$

$$F_{t+m} = a_t + b_t$$

(Christopher, 2021).

Finally, we compare traditional manual solutions to a holistic pharmacy management system. A review of literature by Ayad Ali at the University of Florida finds three methods to manage inventory. The visual method is where stock level is visually counted and compared to a list of required stock amounts; the pharmacists orders medicines when the stock level falls below this listed amount. The periodic method is where the pharmacist carries out the visual method on a regular basis, at predetermined periods of time. Finally, the perpetual inventory management method, most commonly used in industrialised countries, is where a computer-based system constantly monitors inventory and stock levels, automatically updating based on carried out sales or purchases. The paper states that hybrids of these methods can be employed by pharmacists (Ali, 2011). In comparison, a different Healthray blog proposes a system with demand forecasting (as we have seen above with the Holt method), automated ordering, real-time inventory tracking (i.e., the perpetual inventory management method), and expiration date monitoring features in a single software-based solution for pharmacies (Balar, 2025). The synthesis of these features provides many potential upsides, which is why this project aims to create such an integrated solution for its pilot hospitals.

2.3 Resource Review

Although my artefact primarily involves the use of programming and creating AI models, much of my research surrounded the application of my artefact to the real world, since that helped me understand the requirements and features to prioritise in the computer-side development. Thus, my research began in the general field of the healthcare industry and given its scale and relevance to the modern economy, there were significant quantities of resources available online detailing overviews of the industry.

Most of this initial overview is centred around a U.S. Healthcare Market Report (Market Data Forecast, 2025). The company that published this report, Market Forecast Data, is a market research firm that specialises in predicting future market trends by analysing large datasets and historical information. Little is available about the firm on the internet, apart from their own self-descriptions claiming to have “rich experience in research,” although I found that they were based in Hyderabad, India. Since the firm has published at least five reports on similar topics in 2024, such as U.S. Healthcare BPO Market and the North America Digital Health Market, the company does come across as an experienced source, increasing the robustness of their publications. The report authors’ Indian background would serve to reduce any potential bias in the report, as they are likely not affiliated with U.S. healthcare firms and would thus have little motive to unjustly promote their cause. This strengthens the credibility of the 90-page report as a reliable and neutral overview of the U.S. healthcare market. The report was last updated in October 2025, indicating that it is extremely up-to-date, maintaining its relevance to my research. However, since access to the full report is paid, I was unable to identify the sources behind the many statistics the report refers to, so this does not rule out potential inaccurate and misinformation in the report.

Next, I sought to understand the healthcare market in India, using an article published online by the Institut Montaigne, which is a think tank based in Paris, France. The institute makes public policy recommendations on social cohesion, public policy, competitiveness, and public finances, so its 2020 report on Private Healthcare in India (Jaffrelot & Jumle, 2020) corroborates the purpose of the think tank. Since the institute is not a typical academic journal, their publications are not peer-reviewed, however the work is a result of rigorous analysis conducted by a group of experts, academics, and professionals working together. The report I use was

written in 2020, and while it is relatively up-to-date, it importantly does not paint a complete picture of the Indian healthcare market today, by not having information on the post-COVID era. COVID-19 was a major pandemic that revolutionised the workings of many of the world's healthcare industries, so an article written before or during COVID, despite the convincing, reliable nature of the authoring institute, lacks a certain relevance. Regardless, the report analysed the specificities and shortcomings of the private healthcare market, providing foundational knowledge that greatly assisted my later exploration of the challenges faced in pharmacy management.

After understanding the wider healthcare industries, both globally and in my target country of India, I read a series of academic research papers, published in reputable journals world-over, discussing specific challenges and solutions relating to pharmacy and inventory optimisation and management. Many of the papers I came across were written in highly technical language, since their target audience may have been fellow PhD's or medical professionals reading such medical journals, however, four papers in particular overcame these lexical limitations and were great pieces of research providing important information in a clear and concise manner.

First, Abdul K. Akter, from the Department of Pharmaceutical Sciences at Green Valley College in Bangladesh, wrote a paper titled "Optimizing hospital pharmacy inventory management systems: Challenges and Solutions," (Akter, 2023) which is extremely closely linked to my intended area of research. While the author's university is not renowned or particularly well established in the field of medicine, the paper appears robust, outlining challenges and solutions in luminous clarity. The fact that paper was published in a peer-reviewed journal specialising in the topic, the Journal of Pharmacist and Hospital Pharmacy, combined with its recent publishing date of 2024, magnifies its reliability.

Second, I came across a review article written by Ayad K. Ali from the University of Florida, titled "Inventory Management in Pharmacy Practice: A Review of Literature" (Ali, 2011). This resource was very similar to the previous paper in that it provides very technical information about the specific topic I am interested in yet remains highly readable to non-experts like myself. As a review of literature, the article itself is balanced and neutral, discussing a range of methods of inventory management without falling into the trap of being biased. However, this resource had several key limitations. The paper is published in Archives of Pharmacy Practice, which infers that it may not be peer-reviewed, raising doubts about its reliability as a source. Also, since it is

a review of literature, the paper can be considered a secondary source, and may distort and poorly explain findings of the original, primary sources such as experiments and studies that it attempts to expound. Most critically, the paper was published in 2011, which was long before the advent of AI and its associated automation systems. However, despite creating tenuity when relating to the intersection of AI and pharmacy management, the paper was uniquely positioned to illuminate traditional methods that predated advanced computational power of today. It also encouraged further exploration of historical methods to improving business efficiency, igniting my research into the Lean Six Sigma method, which was critical to my project.

Many of the research papers I read were authored by academics at universities all over the world, adding weight to my research as providing a holistic, global view of the topic. The third main paper I used was written by Ilma N. Rachmania and Mursyid H. Basri from the School of Business and Management at the Bandung Institute of Technology. This paper is comparable to the previous, since it was published relatively long ago, in 2012. However, unlike the previous paper, the methods discussed are mathematical, so are unaffected by the rise of AI. For instance, they present the Economic Order Quantity model from calculating procurement lot sizes, determined by a single equation that remains timeless regardless of inventions increasing computing power. As another article published in a peer-reviewed journal, the paper, titled “Pharmaceutical Inventory Management Issues in Hospital Supply Chains,” (Rachmania & Basri, 2013) strengthens the credibility of my literature review section.

To provide detailed depth through a case study, I studied a cross-sectional investigative analysis conducted by five authors on the “pharmacy store of a recently established national institute in East India” (Jaju, et al., 2023). Given the fact that public sector organisations and national institutes operate similarly to charitable institutions and given that this study was conducted in my target country of India, this paper was very relevant to my findings. The credibility of the resource is increased significantly by the paper’s availability on the National Library of Medicine’s PubMed Central journal. Associated with the National Center of Biotechnology Information, the journal is highly authoritative and implies that the paper is robust. However, as an investigative analysis of the All India Institute of Medical Sciences, Deoghar, which the authors are affiliated with, the paper may be subject to positive bias where the authors do not wish to denigrate the institute. This was considered when evaluating their findings in my research. Another limitation, which the authors themselves acknowledge, is that their “participant number was low as it is a single-center survey,” so their conclusions may not reflect the “actual scenario

of a rural setup.” While the paper was convincing, it is important to understanding that such limitations could leave its results ambiguous and unapplicable to our desired settings. Furthermore, volunteer bias cannot be ruled out, and the effect of this bias depends on the experimental design.

Another case study I cross-examined alongside the one in India was an “Analysis of Pharmaceuticals Inventory Management” (Deressa, et al., 2022) in health facilities in the West Shewa Zone, Oromia, Ethiopia. The geographical diversity, being the only paper from the continent of Africa, widened the perspectives of my research. The peer-reviewed paper was published in 2022 by Dove Medical Press, and provided a highly readable, understandable description and discussion of their experimentation, identifying ABC-VEN matrix analysis as a suitable method of management. A major limitation in this resource was the fact that, despite the paper being recently written and published, the ABC analysis method discussed was common in the 1970s, a completely different time period with little relevance to pharmacy management nowadays. Despite this downside, the results obtained by the authors in this direct experiment on various health facilities’ recent data holds importance as primary evidence, magnifying its reliability.

I also read several mathematical papers in the process of researching predictive analytics methods to demand forecasting. The Komunitas Dosen Indonesia journal published a paper titled “Information On Pharmacy Inventory Management With Forecasting Method (Double Moving Average & Double Exponential Smoothing)” (Christopher, 2021), providing more objective viewpoints that many of the resourced I had used before. The paper discusses no personal opinions or beliefs of the author, and rather, directly explains multiple methods to set up systems of equations modelling demand over time. One disadvantage that potentially reduces the credibility of this paper is that all 11 of its references appear to be from Indonesian authors, implying that the information provided is heavily biased towards any consistent cultural or habitual beliefs held by academics in the region.

By cross-examining multiple academic papers discussing different aspects of pharmacy and inventory management, I have ensured that I reduce the overall effect of source biases in my research and have thus avoid triangulating my research to a single group of authors or papers that provide the same solutions. The wide variety of standpoints I analysed prevents the distortion of relative weightings and importance given to any single conclusion.

These academic papers were supported by a type of resource that is typically unreliable: news articles. I am aware of the fact that such resources can be exaggerative, so I limited my use of them to serve the purpose of understanding developments on case studies such as the Indian government's Hospital Management and Information System (HMIS) (Times of India, 2025) and the Indian government agency Unique Identification Authority of India (UIDAI) (Mukherjee, 2025). The news articles used are written for a very different audience than academic papers, so prioritise the delivery of clear information to the masses, although they could be biased towards promoting certain government initiatives, like the HMIS, to be favoured by them. One of the newspapers I referred to, the Times of India, is the largest selling English-language daily in the world, so its popularity may imply that many believe it to be neutral and informative. Another newspaper, Mint, is the business- and politics-specialised sister newspaper to the Hindustan Times, the second largest selling English-language daily India, after Times of India. This indicates that Mint article I used stems from expert authors and a reliable media family.

In my research, I also came across corporate resources. A standout article was "How AI-Powered Pharmacy Software Systems Improve Inventory Management" (Balar, 2025), published by Healthray, a firm that sells AI-driven, cloud-based hospital management systems. The author, Yogesh Balar, is a Director of Business Development at the company, suggesting that the article's main purpose is to drive business development through advertising the benefits of AI-powered pharmacy software to inventory management. Considering the objectives of Healthray and the product it sells, it becomes clear that the source bias distorts the information shared in the article, making it unduly positive and supportive of such systems. This is demonstrated with the excessive reference to case studies and success stories and the absence of considerations of or references to downsides of their product in the article's conclusion.

Finally, contrasting the online articles and papers, which are relatively short resources, I also used several books to inform my research. Some of these include "Information Systems for Healthcare Management" by G.L. Glandon, D.H. Smaltz, and D.J. Slovensky, and "Operations Management in Healthcare: Strategy and Practice," by C.M. Karuppan, N.E. Dunlap, and M.R. Waldrum. The most critical book in the context of my literature review was "The Lean Six Sigma Guide to Doing More with Less," by M.O. George (George, 2010). Published in 2010, the book is written for readers who are business leaders interested in cost reduction opportunities and

strategies. Thus, its simple language style and form as a general guide covering many aspects of Lean Six Sigma, makes it a robust piece of literature and a remarkably useful resource.

Some of my most comprehensive resources had restricted access to the public, but I was able to benefit from them using my personal access to Stanford University Libraries' collections of over 12 million physical and digital resources via SearchWorks. Overall, the breadth and depth of resources used indicate the widespread availability of high quality knowledge and information accessible via the internet.

3 Model

Having discussed both the key challenges and existing solutions in the healthcare industry and market, we now proceed to discuss the methodology implemented in this project's pharmacy optimisation management system.

3.1 Methodology

After exploring the background and existing challenges and solutions surrounding pharmacy management, I began a collaboration with the two selected pilot hospitals. Over months, I arranged many calls with the pharmacy teams of both hospitals to understand how their current operational systems worked, and to identify the major challenges they faced. I was also able to conduct physicals visits to both sites, in Gujarat and Rajasthan, on multiple occasions, for the purpose of this project. After these discussions and observations, we identified three main features for a prototype POMS: demand forecasting, stockout risk prediction, and expiry risk prediction. My goal would be to merge these features into a single platform that fit into their existing inventory management system, enhancing it by reducing risks and occurrences of stockouts and expiry. POMS would also generate a regular pharmacy report based on the three models, providing useful information to pharmacists regarding their inventory. The result would be more efficient running of the pharmacy, with patients having desired medicines available and on-hand with less of a waiting time, thus increasing their utility and welfare. The next step was to improve vendor management through a suite of AI automation tools called PVMS.

In summary, we develop two solutions: first, a Pharmacy Optimisation Management System (POMS) prototype model and second, a more advanced Pharmacy Vendor Management System (PVSM) pilot application created in collaboration with the hospital's software engineers.

The POMS has been created on Jupyter Notebook, a web-based interactive development environment. We use this because it provides a simple platform to experiment with live code while receiving immediate outputs from our models, including forecast data tables and visualisations, alongside the ability to write explanatory texts such as markdowns. We use the programming language Python for its versatility, ease of use, and suitability for our needs.

The PVSM is much more extensive, and a wide range of pre-requisites and programming tools. We use VS Code as the source code editor because it enables one to maintain a clear workspace as we switch between dozens of code files in our project folder. VS Code is also compatible with external services. We primarily use JavaScript, especially for the backend programming, API routes, database operations, and frontend logic. We use TypeScript for frontend components, API client typing, and hooks, as this is a full web-stack application with integrated GUI. We use a combination of Typescript with Javascript XML syntax in some places, in the form of TSX coding. We also use, HTML, CSS, and JSON, for styling and data-related programming.

The PVSM requires a modern web development stack comprising Node.js (v18+) with npm for dependency management; a MongoDB database (we use a local Community edition) accessed via Mongoose; and a Unix-based development environment (we use macOS). The backend is implemented in JavaScript using Express.js and depends on libraries for file ingestion and parsing (CSV/XLS/XLSX) and RESTful API handling, alongside integration with the Google Gemini 2.0 Flash API for AI-driven prescription analysis. The frontend is built using Next.js 15 with the App Router, React, TypeScript/TSX, TailwindCSS 4, and shadcn/ui, requiring a modern browser for execution. Additional requirements include environment variable configuration for API endpoints and keys, Git for version control, and a terminal capable of running Node/npm workflows.

After cleaning through the data, we were ready to build the demand forecasting model. In the next section, we outline both our POMS and PVMS models, including their design and structure, and logic behind their AI features.

3.2 Pharmacy Optimisation Management System

The application has both a front-end and a back-end, the latter of which has four key components: data cleaning, demand forecasting, stockout and risk classification, and automated report generation.

To conduct predictive analytics demand forecasting, we required a very standardised format of data.

Shrimad Rajchandra Hospital has two branches within its pharmacy: an Inpatient (IP) Pharmacy for patients living in the hospital for treatment, and an outpatient pharmacy for visitors, called the Raj Pharmacy. We received thorough data from the seven months prior to December 2025, in the form of 14 spreadsheets that shared the two pharmacies' data for each month from May to November 2025. The quantity of data was sufficient since each spreadsheet consistently listed 13 units of information for over 1,200 pharmaceutical items. A sample of a row from the spreadsheets is as follows:

S.No	Item code	Item Name	Item Description	Vendor Name	Manufacturer Name	UOM	Item Type	Sales Qty	Unit Price	MRP	Net Rate	Current stock
81	ORD2 20085 39	CAP DEPLATT CV 20/75/75 MG 15's	CLOPIDOGREL BISULPHATE 75MG + ATORVASTATI N 20MG + ASPIRIN 75MG	FLORIDA MEDICO (SURAT)	TORRENT PHARMACEUTICAL LTD	STRIP OF 15	CAP	73	137.8	200	137.8	227

Table 1. Extract from SRST-IP Pharmacy November 2025 Sales Report.

However, hospital pharmacy data is not collected for machine learning use. The spreadsheets used in this project varied across months and branches, with headers beginning several rows down, inconsistent month naming, blank rows and merged cells, and separate files for in-patient and out-patient pharmacies. Therefore, a robust ingestion pipeline was required before any modelling could be attempted.

Using pandas, all spreadsheets from a directory were load programmatically. The header row was explicitly set to row 5 (index 4), based on inspection of the data structure. Only three columns of particular interest were retained for the initial models: Item Name, Sales Qty, and Current Stock. The first step was to extra data from all 14 spreadsheets, and clean it by removing the remaining, unnecessary columns. Metadata columns for month and pharmacy branch were extracted directly from file names.

We then renamed the key columns to be Medicine, Sales, and Stock, and concatenated all the spreadsheets into a single dataset, called Merged_Data. Then, we rearranged the 13,500+ rows to produce 7x4 arrays for each medicine, with the Medicine, Sales, and Stock listed for each of seven months. Month strings were normalised and converted into a datetime format, allowing the creation of a continuous time index (MonthIndex) to support time-series modelling.

```
# load sheet (header row is 5 → index 4)
df = pd.read_excel(os.path.join(data_dir, fname), header=4)

# standardize column names (strip spaces)
df.columns = [str(c).strip() for c in df.columns]

# required columns
try:
    item_name_col = [c for c in df.columns if "Item Name" in c or "ITEM NAME" in c][0]
    sales_col     = [c for c in df.columns if "Sales Qty" in c or "SALES QTY" in c][0]
    stock_col     = [c for c in df.columns if "Current stock" in c or "CURRENT STOCK" in c][0]
except:
    print("ERROR: columns not found in", fname)
    print(df.columns)
    raise

df_clean = df[[item_name_col, sales_col, stock_col]].copy()
df_clean.columns = ["Medicine", "SalesQty", "CurrentStock"]

def normalize_month(m):
    parts = m.split()
    if len(parts) == 2:
        mon, year = parts
        mon = month_fix.get(mon[:3].capitalize(), mon[:3].capitalize())
        return f"{mon} {year}"
    return m
```

All of this is done via Python code on Jupyter Notebook, specifically through our custom-built functions `load_month_files(data_dir)` and `normalize_month(m)`.

Our model's sample had to be a single medicine with both sales and stock history for a given period of time in history and a predicted sale and stock for the next window of time. Our input had to be a sequence of data points for the past N months, where N was an integer to be decided in conjunction with the pilot hospitals, balancing what was feasibly accessible and what was required for an accurate model. The output would be suggested procurement quantities for the next month, based on the difference between predicted sale and stock. Thus, the demand forecasting feature was a supervised machine learning task.

Each medicine was treated as an independent forecasting problem, because different medicines exhibit fundamentally different demand patterns, and pooling medicines into a single global model would introduce noise, reducing overall accuracy. Initially, a was chosen. It captures temporal patterns and works well to support the small size of the dataset (seven months per medicine).

First, we had to import several library requirements, including pandas, numpy, sklearn, and tensorflow.

```
import pandas as pd
import numpy as np
from sklearn.preprocessing import MinMaxScaler
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import LSTM, Dense
from tensorflow.keras.callbacks import EarlyStopping
```

Sales values were scaled using a MinMaxScaler on a per-medicine basis to stabilise training and avoid domination by high-volume medicines.

```
scalers = {} # dictionary to keep scaler per medicine
medicines = df["Medicine"].unique()

for med in medicines:
    df_med = df[df["Medicine"] == med]
    scaler = MinMaxScaler()
    df.loc[df["Medicine"] == med, ["Sales", "Stock"]] = scaler.fit_transform(df_med[["Sales", "Stock"]])
    scalers[med] = scaler
```

We then created sequences of inputs and outputs to train the model. An example is as follows:

Example input sequence:

```
[[0.46453534 0.44518246]
 [0.3775563 1.    ]
 [0.71628268 0.61451903]]
```

Example output: 1.0

Next, we built the model and trained it on a single medicine with 50 epochs.

```
model = Sequential()
model.add(LSTM(50, activation='relu', input_shape=(X.shape[1], X.shape[2])))
model.add(Dense(1))
model.compile(optimizer='adam', loss='mse')
es = EarlyStopping(monitor='loss', patience=5, restore_best_weights=True)
model.fit(X, y, epochs=50, batch_size=1, verbose=2, callbacks=[es])
```

Finally, we used the trained LSTM to make a prediction of sales for the next month, which was unscaled before being outputted.

```
pred_scaled = model.predict(X[-1].reshape(1, X.shape[1], X.shape[2]))
```

The result was a successful, realistic prediction for December 2025. So, we attempted to loop the LSTM model training for every medicine and produce a complete forecast spreadsheet with all predictions. However, we quickly noticed that an LSTM model was too heavy to work on the vast quantities of medicines the model must run through. Training the model and running a separate forecast for each of 2,500 unique medicines was incredibly time-consuming, resource-consuming, and impractical.

The result was a major change of direction: we decided to switch to a lightweight, linear regression model that could run quickly for all medicines without compromising accuracy. The primary code we implemented is as follows:

for med in medicines:

```
for med in medicines:
    df_med = df[df["Medicine"] == med].sort_values("Month")
    sales = df_med["Sales"].values

    # Need at least 2 points to forecast anything
    if len(sales) < 2:
        continue

    # Prepare training data
    X = np.arange(len(sales)).reshape(-1, 1)
    y = sales

    model = LinearRegression()
    model.fit(X, y)

    # Predict next 1 month
    next_t = np.array([[len(sales)]])
    forecast = float(model.predict(next_t)[0])

    results.append({
        "Medicine": med,
        "LastMonthSales": float(sales[-1]),
        "PredictedNextMonthSales": forecast
    })
```

The resulting forecasts were successfully exported to a CSV called Forecast_Results, completing our first model. Below are five randomly selected pharmaceutical items, with their predicted sales December 2025. Note that this is the unprocessed output, without headers neatened or predictions rounded to integers.

Code	Item	LastMonthSales	PredictedNextMonthSales
98	BETAKIND GARGLE 2%	135	196.8571429
751	INJ T T (BETT) 0.5ML	109	93.57142857

977	NEEDLE NO 18 X 1.5"	1220	1426.428571
1180	SUCTIONPRO72 NO 14FR TT (REFZ216-14)	13	13.42857143
1841	TAB LINOKEM 600MG	33	28.71428571

Table 2. Forecast_Results extract of predicted sales for December 2025.

We then classify each medicine into one of five categories: High Risk (no stock), High Risk, Moderate Risk, Safe, and Overstock / Expiry Risk. This is calculated based on ratio of the stock of the medicine at the end of the previous month compared to the predicted sales for next month, which we call demand. We define High Risk (no stock) to be where stock is zero, High Risk to be where stock is less than 1.2 times demand, Moderate Risk to be where stock is between 1.2 and 1.5 times demand, Safe to be where stock is between 1.5 and 4 times demand, and Overstock / Expiry Risk to be where stock is greater than 4 times demand.

The classify(row) function executes this comparison and adds a risk category column to the Forecast_Results, exporting to a new spreadsheet called Classified_Inventory.

```
def classify(row):
    stock = row["Stock"]
    demand = row["PredictedNextMonthSales"]

    if stock <= 0:
        return "HIGH RISK: No Stock"

    if stock < demand * 1.2:
        return "HIGH RISK"

    if stock < demand * 1.5:
        return "MODERATE RISK"

    if stock > demand * 4:
        return "OVERSTOCK / EXPIRY RISK"

    return "SAFE"
```

We then generate a suggested procurement order for the pharmacy, taking into account the stockout and expiry risk classification. This model is part of the POMS, so does not compare vendors and identify an optimal vendor. Vendor management will be a feature in the PVMS. The formula we use is as follows:

```
#FORMULA:
RequiredStock = PredictedDemand × SafetyMultiplier
SafetyMultiplier = 1.5 → 2.0
ProcurementQty = max( RequiredStock - CurrentStock , 0 )
```

The model outputs a Procurement_Plan spreadsheet incorporating all the above features: demand forecasting, stockout and expiry risk classification, and suggested procurement quantities. Below is the final output for the same items as Table 2.

Item	LastMonthSales	PredictedNextMonthSales	Stock	RiskCategory	SuggestedOrderQty
BETAKIND GARGLE 2%	135	196.8571429	187	HIGH RISK	206.7142857
INJ T T (BETT) 0.5ML	109	93.57142857	92	HIGH RISK	95.14285714
NEEDLE NO 18 X 1.5"	1220	1426.428571	1268	HIGH RISK	1584.857143
SUCTIONPRO72 NO 14FR TT (REFZ216-14)	13	13.42857143	36	SAFE	0
TAB LINOKEM 600MG	33	28.71428571	36	MODERATE RISK	21.42857143

Table 3. Procurement_Plans extract of suggested order quantities for December 2025.

The table below shows two items at extreme opposite ends of the risk category classification, where one is at risk of overstock, leading to excess storage costs (and expiry if it is a consumable that can become expired), and the other has no stock, and is at high risk of a stockout.

Item	LastMonthSales	PredictedNextMonthSales	Stock	RiskCategory	SuggestedOrderQty
BC2435-20 NEONATAL OXYGEN THERAPY NASAL CANNULA	1	1	5	OVERSTOCK / EXPIRY RISK	0
ZOCON EYE DROPS 5ML	1	1.5	0	HIGH RISK: No Stock	3

Table 4. Procurement_Plans extract of two extreme risk items.

Out of the 2,471 medicines that the procurement plan consisted of for December 2025, 305 had no stock, 850 were at high risk with some stock, 204 were at moderate risk, 644 were safe, and 468 were overstocked or at expiry risk. These statistics produce useful insight for pharmacists at our pilot hospitals.

Finally, it generates a plain text PDF pharmacy report summarising its findings, in a clear, readable format for the pharmacist. This is done with a Gemini API analysing the Procurement_Plan spreadsheet.

3.3 Pharmacy Vendor Management System

Next, we come to the Pharmacy Vendor Management System (PVMS), a comprehensive digital platform that optimises procurement, inventory decision-making, and vendor selection in hospital pharmacy operations. The system was developed using modern web technologies and built in collaboration with software engineers working at Shrimad Rajchandra Hospital. The system addresses, amongst others, two fundamental operational challenges faced by hospital pharmacies: vendor data arrives in highly inconsistent formats, and manual scanning of this data to find best vendors for a certain procurement is error-prone, time-consuming, and non-scalable.

PVMS follows a modular architecture consisting of a frontend dashboard, a backend API server, and a central database layer. The backend runs on Port 3000, using Node.js with Express.js. Its database is managed through MongoDB with Mongoose ODM, and AI Integration occurs through Google Gemini 2.0 Flash. File Processing happens with Excel (.xlsx) and CSV parsers. The frontend runs on Port 3001, with a framework of Next.js 15 with App Router. Styling is done through TailwindCSS 4 with shadcn/ui components, and the language is TypeScript.

The project follows a clearly defined directory structure:

```
pharma-vendor/  
├── Backend/  
│   ├── models/      # Database schemas  
│   ├── routes/      # API endpoints  
│   ├── utils/       # AI and file processing  
│   ├── data/        # Vendor Excel files  
│   └── api-server.js # Main server file  
├── pharma-vendor-frontend/  
│   ├── src/  
│   │   ├── app/     # Next.js pages  
│   │   ├── components/ # UI components  
│   │   ├── lib/      # API client and utils  
│   │   ├── hooks/    # Custom React hooks  
│   └── package.json  
└── README.md
```

PVMS's backend is structured around four core modules that form a continuous operational pipeline: vendor data ingestion, order processing, stock analysis, and prescription analysis. With

the use of AI, each module abstracts away unnecessary and excessive manual labour and subjective human judgement.

In real hospital settings, vendor catalogues arrive in highly inconsistent formats, including spreadsheets, PDFs, and even WhatsApp messages containing free-text descriptions of active ingredients. The vendor data ingestion module accepts Excel or CSV files and routes them through an AI-driven normalisation function, where Google Gemini semantically interprets the data. Gemini transforms heterogeneous formats into a consistent structure, and parses active ingredient strings are then parsed into key components, including detecting potential synonyms across brands and generics, to ensure that medicines with identical pharmacological content are linked together. Once normalised, the data is added to a master database optimised for ingredient-based search and efficient vendor comparison. This module enables all subsequent decision-making by converting unstructured, vendor-specific data into a reliable medicine catalogue.

The order processing module processes stock balance sheets and daily sales reports from multiple departments, namely the in-patient, out-patient, and operation theatre departments. Consumption patterns and purchase requirements are then analysed against current stock levels, with AI-powered reorder quantity calculations generating an order sheet. Each medicine is then matched against the vendor database, and optimal vendor recommendations are generated by evaluating prices, discounts, margins, and availability simultaneously. The result is a procurement output that is both cost-optimised and operationally efficient.

Next, the stock analysis module extends the system beyond short-term ordering into continuous inventory optimisation. It aggregates data from goods receipt notes, stock balance sheets, sales reports, and inter-department transfer records. Initial stock, sold quantities, transferred quantities, usage, and remaining balances are used to classify inventory states such as critical shortage, low stock, optimal stock, or overstock. The output of this module is a reorder analysis that provides stock insight and informs future purchasing action.

Finally, the prescription analysis module processes prescription images using AI-based text extraction, allowing the system to identify prescribed medicines, match them against the internal medicine database, and evaluate real-time availability and pricing across vendors, enabling immediate cost and supply comparison.

The frontend is designed as an interactive, operational dashboard rather than a static reporting interface, allowing pharmacy staff to interact with the system continuously throughout the day. It is a modern web application that communicates through typed API calls. Thus, both our frontend and backend design choices establish PVMS as an effective, deployable model for modern pharmacy operations.

4 Conclusion

My artefact's success lies in its own construction: I have been able to review the market in depth and identify gaps and challenges faced in existing solutions. Following this research, I designed and successfully programmed two independent models for implementation in pilot hospitals. The POMS is able to utilise real-world data from hospitals to predict sales for the next month, classify medicines into stockout and expiry risk categories, and compare predicted demand to current stock to suggest procurement quantities for the pharmacy. The ability of the model to execute these tasks indicates a success outcome, as predictive analytics features, when implemented in the pilot hospitals, will reduce costs and improve operational efficiency, meeting my original objectives.

Furthermore, the PVMS makes use of AI automation in the workflow of managing a hospital inventory. By processing and analysing vendor data, stocks, and prescriptions, as well as comparing medicine offers and generating orders, the system is well positioned to be implemented in Shrimad Rajchandra Hospital in Gujarat and Monilek Hospital in Rajasthan and is predicted to reduce manual labour by 25,000 hours annually, translating to \$35,000 of cost saving. Thus, it will improve the hospitals' ability to care for the 25,000 patients they cater to monthly, increasing access to over 6,000 unique pharmaceutical items (Stock Keeping Units). The social and communal benefit predicted to be created by these systems demonstrates to its successful design and development, surpassing my initial criteria and targets.

A comparative analysis of approaches found that predictive analytics reduce stockouts by 25%, reduce wastage by 15%, and improve efficiency by 20% (Akter, 2023). This suggests that the POMS will minimize stockouts and overstocking, leading to more efficient inventory management, as desired. Furthermore, the study found that automated inventory systems, such as the PVMS, reduce stockouts by 30%, reduce wastage by 20%, and improve efficiency by 25%

(Aker, 2023). Additional benefits of real-time tracking are enhanced accuracy and reduced medication delays, which indicate that the PVMS will reduce the time it takes for critical medicines to be made available for low-income patients in rural communities, benefitting them.

5 Evaluation

I have learned a wide variety of skills and lessons in the process of completing my Extended Project Qualification, both overall as executer of a project from end-to-end and specifically through technical development of my AI-driven pharmacy optimisation and vendor management systems.

Relating to the former area of improvement, my time management skills were tested thoroughly. In the Project Proposal Form, I created a project timeline (months in advance of the actual execution), constructing a well thought-out plan that balanced my school workload, travel schedule, and other extracurricular commitments, such as sports fixtures and music exams, alongside the demanding nature of the EPQ. This strengthened my ability to organise my time in an efficient manner to ensure that heavy quantities of work were completed by their respective deadlines.

One major challenge that developed this skill was having to push forward much of my plans for the summer due to my summer program being more time-consuming than I had expected. Initially, the summer was a time when I expected to be very productive with the EPQ, given that we had no school. However, upon arriving at Stanford University and settling into my two month summer course, I realised that studying two undergraduate-level courses in parallel involved 40 hours of lectures and homework. Adding that to the program's many opportunities, including day-trips, workshops, clubs, and social events, there was very little time left for the EPQ. Regardless, I strove to work through as much as possible given the situation and aimed to restructure my plan for the rest of the project, which ultimately worked successfully, given that my paper and artefact were both complete during the winter break.

This project has been a masterclass in daily adaptability, and I have learned that it is very important to be able to bounce back after setbacks, persevering with relentless momentum towards one's goal. Strength and resilience are paramount to progress. Otherwise, any significant project undertaken would never see the light of day. In future projects, I would

definitely keep this in mind, remaining motivated throughout the many challenges I would doubtless face. This conviction arises from my personal experiences in this project. My EPQ has been unique in that I partnered with the actual industry, working with real-world hospitals to obtain data and even programming my second solution, the Pharmacy Vendor Management System, in collaboration with software engineers from one of the hospitals. This involved many hours of communication and coordination across countries and time-zones.

It was particularly challenging to get in touch with the senior pharmacy leaders, whom I had to liaise with to receive permissions, such as asking for extensive private data from the hospitals. Oftentimes, after many failed attempts at messaging and calling them, I would have to make the most of a sudden returned call, no matter where I was at the time, to ensure that I achieved to obtain whatever help or support I needed from them.

Overall, though, working with the hospitals has been a great learning experience. Importantly, it has given me critical interpersonal skills, which I did not expect to gain through a coding project and through writing its associated paper. I have also mastered the art of corporate communication, a skill that will carry forward into my future career. I truly enjoyed working with these hospitals, and after visiting them multiple times in-person have understood the importance of having strong connections with one's partners to ensure the success of a project. In future projects, I will definitely strive to build such positive relationships with my collaborators.

Although the EPQ has been full of challenges and setbacks, there have been countless smaller lessons gained while working through them. From a programming aspect, I have learned to work with many new technical features, including using the terminal to set up frontend localhosts, installing and using MongoDB to work with databases, and using Node.js and Gemini APIs. As my first full project run on Jupyter Notebook and my largest on VS Code, I have learned to set up a structured directory for both programming files and extensive data spreadsheets. It was also my first time working with Tensorflow for deep learning methods in such an extensive way.

Through hours of planning code logic, debugging through syntax errors, matching kernel version for libraries, and fixing broken server connections on Uvicorn API apps, I have undoubtedly improved my coding skill and learned to build ever more complex, useful, and impressive programs.

Writing the paper also taught me a lot. In the process of research, I have learned to effectively use academic resources, including reading research papers at depth and skimming books to glean important information. By engaging with research while writing the literature review, I have developed a clear perspective on the use of AI in business optimisation. I have gained considerable knowledge on the history and features of pharmacy management, with solutions ranging from Lean Six Sigma and the Japanese *Kaizen* to double moving average demand forecasting and Holt's model for trend-corrected exponential smoothing. Writing a paper without reliance on artificial intelligence has strengthened my ability to communicate in an era that is heavily dependent on generative AI. I have critically thought about what I have intended to share with the reader and have effectively conveyed that information, concisely imparting my learnings from reading. The paper has no doubt achieved what my initial objectives were, as I have identified the "key challenges and inefficiencies in existing systems in inventory and vendor management" and have and built "a working proof-of-concept that solves these problems.

My presentation was another section of my EPQ that went very well. I am fortunate to have had significant experience in presenting projects to groups of students in the past, including at summer programs and courses. As a result of my accumulation of leadership positions over the years, I have also gained experience in public speaking, including giving speeches to audience of 1,000+, addressing important individuals at external events, and delivering a TEDx talk. The combination of these influences from my past positioned me well to deliver a clear and logical presentation, with consistent speaking pace, conveying a summary of my entire project to a class of students. My content, from the slides and their balanced text (not too heavy or sparse) to my speaking points and gestures, were well suited to the audience, resulting in an enjoyable and informative presentation. I am also very pleased with my performance relating to responding to questions, such as on data privacy, biggest challenges faced, and the ability of the models to function without AI, demonstrating my insights of the topic and ability to answer on the spot.

Completing this research project has also opened up many avenues for potential further research. Both of my models, the Pharmacy Optimisation Management System (POMS) and the Pharmacy Vendor Management System (PVMS), can be extended in several ways. To the POMS, I can add new deep learning models to increase the accuracy of demand forecasting, and a better frontend interface to improve user experience when working with the system. The PVMS is already very thoroughly built, as a full web-stack application with many robust backend features and an impressive user interface. It is already poised to be implemented, so my goal to extend

this project would be to implement the PVMS in both my pilot hospitals, and from there, expand it to benefit charitable hospitals and their patients in communities all over India. Given the foresight of our planning, the PVMS is design to be scalable and platform-agnostic, as was an objective in my PPF, so there is a strong possibility the implementation and expansion will be successful.

From an academic perspective, while conducting research on pharmacy optimisation, I was especially drawn to the extensive depth of existing research on general business optimisation, including on Lean Management and the Six Sigma model. I am very keen to explore this area further, learning more about these business philosophies through reading books such as Mark George's "The Lean Six Sigma Guide to Doing More with Less" and understanding the concepts of business management better in general. These will serve to benefit me in the future, given plans to eventually join my family's business.

Overall, I am very grateful to have gone through this experience of developing my own EPQ project. I have gained many personal and technical skills that I can carry forward and have drawn perceptive conclusions on the process of producing artefacts in similar projects in the future, as discussed above.

6 References

Akter, A. K., 2023. Optimizing hospital pharmacy inventory management systems: Challenges and Solutions. *Journal of Pharmacist and Hospital Pharmacy*.

Ali, A. K., 2011. Inventory management in pharmacy practice: a review of literature. *Archives of Pharmacy Practice*, 2(4), pp. 151-156.

Anusha, 2025. *Economic order quantity (EOQ) in inventory control*. [Online]

Available at: <https://velexp.com/blog/economic-order-quantity-eoq-in-inventory-control> (Accessed: 18 December 2025)

Balar, Y., 2025. *How AI-powered pharmacy software systems improve inventory management*. [Online]

Available at: <https://healthray.com/blog/pharmacy-management-system/ai-powered-pharmacy-software-systems-improve-inventory-management/> (Accessed: 15 December 2025)

Chittenden, R., n.d. *THE 7 WASTES OF LEAN AND HOW TO REDUCE THEM IN YOUR ORGANISATION*. [Online]

Available at: <https://txm.com/lean-manufacturing-tip-7-wastes/> (Accessed: 18 December 2025)

Christopher, J., 2021. Information On Pharmacy Inventory Management With Forecasting Method (Double Moving Average & Double Exponential Smoothing). *Komunitas Dosen Indonesia*, 4(1).

Deressa, M. B., Beressa, T. B. & Jemal, A., 2022. Analysis of Pharmaceuticals Inventory Management Using ABC-VEN Matrix Analysis in Selected Health Facilities of West Shewa Zone, Oromia Regional State, Ethiopia. *Dovepress*, Volume 2022:11, pp. 47-59.

George, M. O., 2010. *The Lean Six Sigma Guide to Doing More With Less*. New Jersey: John Wiley & Sons, Inc..

Hyndman, R. J. & Athanasopoulos, G., n.d. Exponential Smoothing. In: *Forecasting: Principles and Practice*. Melbourne: s.n.

Jaffrelot, C. & Jumle, V., 2020. *Private healthcare in India: boons and banes*. [Online] Available at: <https://www.institutmontaigne.org/en/expressions/private-healthcare-india-boons-and-banes> (Accessed: 13 September 2025)

Jaju, R. et al., 2023. Inventory Control Mechanism of the Pharmacy Store of a Recently Established National Institute in Eastern India: A Cross-Sectional, Investigative Analysis. *PubMed Central*, 15(11).

Market Data Forecast, 2025. *U.S. healthcare market size, share, trends and forecast*. [Online] Available at: <https://www.marketdataforecast.com/market-reports/us-healthcare-market> (Accessed: 19 September 2025)

Ministry of Health and Family Welfare, Government of India, 2025. *Health Management Information System (HMIS): About Us*. [Online]

Available at: <https://hmis.mohfw.gov.in/#!/aboutus> (Accessed: 15 September 2025)

Monilek Hospital, 2025. *About Us*. [Online]

Available at: <https://monilekhospital.com/about> (Accessed: 14 December 2025)

Mukherjee, A., 2025. *UIDAI registration mandate to bar hotels from making Aadhaar photocopies - here's what you need to know as customers*. [Online]

Available at: <https://www.livemint.com/news/india/uidai-registration-mandate-to-bar-hotels-from-making-aadhaar-photocopies-heres-what-you-need-to-know-as-customers-11765193528777.html> (Accessed: 15 September 2025)

Rachmania, I. N. & Basri, M. H., 2013. Pharmaceutical Inventory Management Issues in Hospital Supply Chains. *Scientific & Academic Publishing*, 3(1), pp. 1-5.

Shrimad Rajchandra Hospital, 2025. *About*. [Online]

Available at: <https://hospital.srmd.org/about/> (Accessed: 14 December 2025)

Times of India, 2025. *SMC's new HMIS has handled 89,119 patients so far*. [Online]

Available at: <https://timesofindia.indiatimes.com/city/surat/smcs-new-hmis-has-handled-89119-patients-so-far/articleshow/121736651.cms> (Accessed: 15 September 2025)

Top Consulting Firms, 2025. *How Tracking Signal Improves Forecasts*. [Online]

Available at: <https://www.allconsultingfirms.com/blog/how-tracking-signal-improves-forecasts/> (Accessed: 16 December 2025)

Wikipedia, 2025. *Kaizen*. [Online]

Available at: <https://en.wikipedia.org/wiki/Kaizen> (Accessed: 13 November 2025)

Wikipedia, 2025. *Six Sigma*. [Online]

Available at: https://en.wikipedia.org/wiki/Six_Sigma (Accessed: 13 November 2025)