Optimization and Rule-Based Models for Hospital Inventory Management

by

Caeley Gaw Harihara

B.S.E., Electrical Engineering
Princeton University, 2019

Submitted to the MIT Sloan School of Management and
Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degrees of

Master of Business Administration

and

Master of Science in Electrical Engineering and Computer Science

in conjunction with the Leaders for Global Operations program

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2024

Authored by: Caeley Gaw Harihara
MIT Sloan School of Management and
Department of Electrical Engineering and Computer Science
May 10, 2024

Certified by: Martha Gray
Professor of Electrical Engineering & Computer Science
Thesis Supervisor

Certified by: Yanchong Karen Zheng
Associate Professor of Operations Management
Thesis Supervisor

Accepted by: Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Studies

Accepted by: Maura Herson
Assistant Dean, MBA Program
MIT Sloan School of Management

Optimization and Rule-Based Models for Hospital Inventory Management

by

Caeley Gaw Harihara

Submitted to the MIT Sloan School of Management and
Department of Electrical Engineering and Computer Science
on May 10, 2024, in partial fulfillment of the
requirements for the degrees of
Master of Business Administration
and
Master of Science in Electrical Engineering and Computer Science

Abstract

This thesis shows how optimization, rule-based models, and operational analytics can be used to help manage hospital surgical inventory. The models were created for AITA™, a team under Johnson & Johnson's Ethicon subsidiary. The AITA™ Smart System is an intelligent inventory management solution that stores, organizes, and distributes products via Kiosk, Smart Shelf, and Mobile Hub devices. Every device requires a planogram, or a visual representation of which products to stock and the location of each product. This project focuses on creating models to automatically build and update these planograms. The models presented in this paper have already been adopted by the AITA™ team and have begun to show accuracy and efficiency gains when compared to the current manual process. Model-designed kiosks cover, on average, 7% more historical procedures than hand-made kiosks. Also, model-generated planograms are free from manual product selection and sorting errors. From an efficiency perspective, automatically creating and updating planograms will save the AITA™ team an average of 145 hours annually for every hospital served. These accuracy and efficiency gains will add value across the entire chain of care. The AITA™ team will have more time to grow their business and to develop new features. Meanwhile, providers will save time when managing and retrieving hospital inventory, which will free up more capacity for direct patient care.

Thesis Supervisor: Martha Gray
Title: Professor of Electrical Engineering & Computer Science

Thesis Supervisor: Yanchong Karen Zheng
Title: Associate Professor of Operations Management

Acknowledgments

This research would not have been possible without the help of so many wonderful people. First, I want to thank my thesis advisors, Professor Martha Gray and Professor Karen Zheng for their invaluable guidance, insightful feedback, and unwavering support throughout this project. I learned so much from working with you both and appreciate all time and effort you have taken to help me grow as a researcher.

Next, I would like to thank Matt Chila, Jude Francis, Natalie Henderson, and the rest of the wonderful AITA™ team for bringing me aboard and for guiding me throughout my project. Your support, expertise, and guidance have greatly contributed to my professional development and to completion of my research. I am so grateful for the opportunity to have worked with and learned from you all.

Last, but certainly not least, I would like to thank my parents, Kris and Jenny, and my boyfriend, Andrew, for being the best cheerleaders that I could ever ask for. I would not be where I am today without your love, encouragement, and support.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

Chapter 1

Introduction

1.1 Project Context and Background

1.1.1 Introduction to Ethicon

Ethicon is a subsidiary of Johnson & Johnson (J&J) that specializes in the development of surgical and interventional technology. While the company's product catalog now spans a wide variety of offerings including medical devices, robotics, and digital solutions, its largest product line is dedicated to surgical sutures [15]. Sutures are sterile surgical threads attached to needles that are used to close (i.e. sew together) skin and other bodily tissues. Suturing promotes wound healing by better aligning tissues, controlling bleeding, and reducing the chance of infection through minimized exposure of internal bodily structures [14].

Sutures are recorded to have been used as early as Ancient Egyptian times, but they were never mass produced until 1887 when J&J began manufacturing sterile sutures made of catgut and silk. J&J officially formed Ethicon Suture Laboratories in 1949 and renamed it to Ethicon Inc. in 1953. Since then, Ethicon has continued to develop innovative suture technologies. For example, in 1969 they created synthetic sutures made from polymer polypropylene, which are still used in cardiac bypass surgeries to this day. In 1982, they introduced sutures made of polydioxanone that are specifically designed to close fascia, connective tissue under the skin. In 2014, Ethicon introduced a needle-less skin closure system called Dermabond® Prineo® that uses a liquid adhesive and self-adhering mesh to non-invasively close surgical incisions. Such innovations have delivered consistent, highly-effective results - leading to Ethicon being the current market leader in surgical sutures [20].

1.1.2 Sutures Overview

Ethicon offers over 3,700 types of sutures, which vary across a range of attributes. These individual products can also be referred to as SKUs, or "stock-keeping units". SKUs are commonly referenced in the retail industry, where every SKU is a unique type of product that can be purchased. In this setting, while there are many ways to distinguish between suture SKUs, for the purposes of this project, the most important characteristics are product group, suture size, needle length, needle point, and closure style. Sutures can be split into dozens of different product groups such as Ethibond, Ethilon, Monocryl, Prolene, Stratafix Spiral PDO, and Coated Vicryl. Each product group has a unique combination of materials and filaments (1-1). For example, Ethibond sutures are made of braided polyester, while Ethilon sutures are made of monofilament nylon.

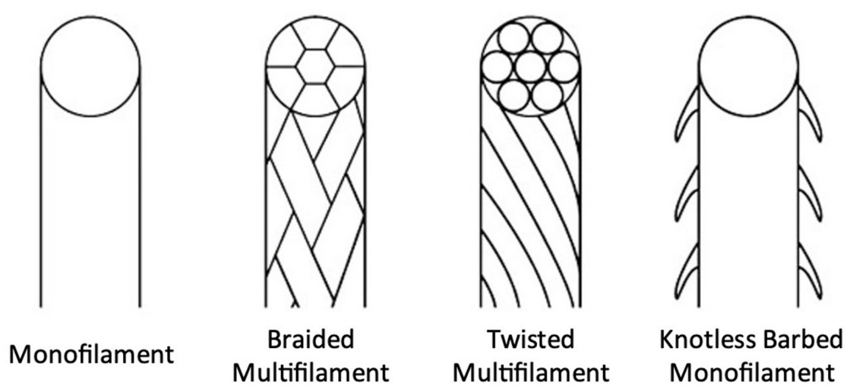

Figure 1-1: Common suture filament structures [16]

Suture size indicates the diameter of the suture needle (1-2). The thickest sutures are numbered from thickest to thinnest starting at 0-10, 0-9, 0-8, etc. The thinnest sutures are numbered from thinnest to thickest starting at 12-0, 11-0, 10-0, etc [14].

Meanwhile, needle length refers to the size of the needle itself, from its point to the location where the needle meets the thread (the swage). The length of suture needles typically range from 3mm to 254mm (10 in).

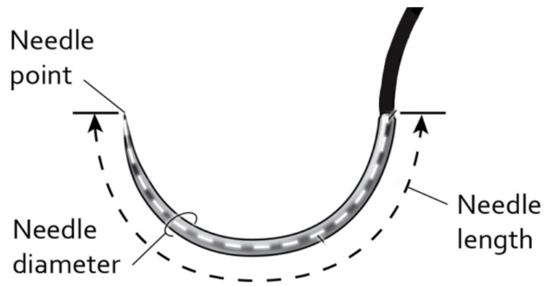

Figure 1-2: Definitions of suture measurements: size/diameter and needle length [1]

The needle point of a suture refers to the shape of the needle's tip. Some common needle points are shown in 1-3. Each point shape enables a suture to better carry out specific tasks. For example, taper point needles are thin enough to pass through delicate tissues with minimal damage. Blunt point needles are less likely to cause tissue injury, meaning they are suited for liver and kidney sutures. Conventional cutting needles have at least two cutting edges to provide the sharpness needed to penetrate dense tissue. Meanwhile, reverse cutting needles carry the sharpest edge on the outer curvature, to minimize the risk of cutting out tissue. In total, Ethicon offers 13 different types of needle points [10].

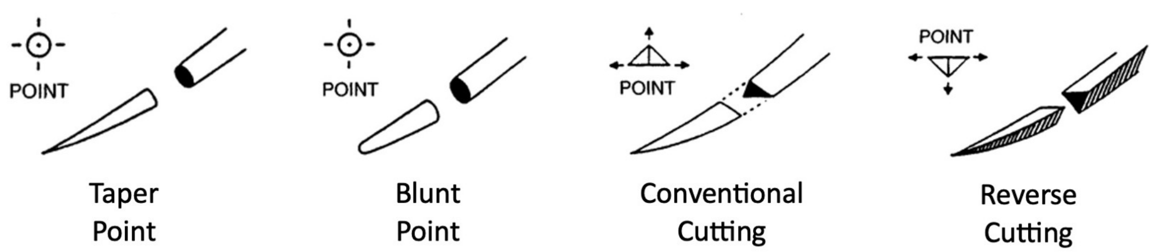

Figure 1-3: Four common types of needle points used in sutures [25]

Finally, closure styles are the last primary characteristic used in this project to delineate sutures. In this case, closure styles refer to how the end of a suture's thread is manufactured. As illustrated in 1-4, there are four main types of closure styles. The most common style, the standard style, has no pre-created closure for the thread.

Under this construction, a surgeon can use any technique they prefer to close the suture and can use any length of the attached thread. The other three main types of closure styles (loops, reels, and ties) are designed to offer convenience to medical providers. When loops and ties are already attached to the end of a suture, it saves surgeons the time needed to create the closures themselves. Meanwhile, reels compactly contain a long spool of thread, enabling surgeons to use more than a standard length of material without needing to use and close multiple, individual sutures.

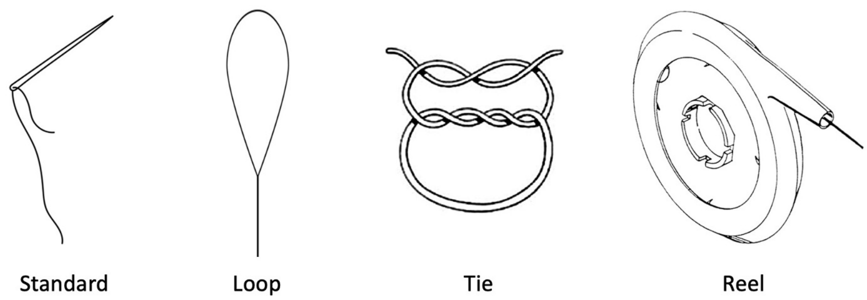

Figure 1-4: Common closure styles that can be pre-installed on sutures [24]

1.1.3 Introduction to AITA™

With thousands of sutures and other surgical products available to medical providers, it is often difficult for hospitals to manage their complete set of inventory. Storage systems can be disorganized, inefficient, and can provide little transparency to hospital administrators on product usage trends. So, Ethicon's AITA™ team was started with the goal of creating a better way to manage hospital inventory. The team now creates and distributes the AITA™ Smart System, an intelligent inventory management solution that uses configurable hardware and cloud-based software to store and distribute surgical inventory while providing data-driven insights.

The AITA™ Smart System uses three types of devices to manage inventory: Kiosks, Smart Shelves, and Mobile Hubs (1-5). Kiosks store up to 180 different types of surgical sutures and products. When a procedure is scheduled, a list of all the products needed is sent to the kiosk. If the kiosk holds more than a threshold percentage of the needed products, it will be assigned to fill the procedure's order. The kiosk fills the order by automatically using its picking arm to grab the necessary products from its slots. It places these items into one of the lockers on the left side of the device. When it is time for the procedure, a member of the hospital staff will open the locker to retrieve the products. Then, when the procedure is completed, the kiosk can help re-stock unused products. All the unopened sutures simply have to be placed in the restocking bay on the right side of the device, and the kiosk will automatically put them back in its bins. This process is significantly faster than typical manual restocking, meaning that providers will save time and fewer sutures will be wasted, because the barrier to restock items is lower.

AITA™ Smart Shelves comprise a series of racks (usually 3-4 racks per shelf), and each rack holds 96 products. Hospitals store all their surgical inventories on these shelves. Smart Shelves manage two main types of inventory: safety stock and 'just-in-case' products. Safety stock is required when a kiosk or a hub runs out of stock for a certain product. Hospital workers can replenish the product by moving units from the shelf to the hub or kiosk. Meanwhile, 'just-in-case' inventory is needed in situations when procedures cannot be filled by the kiosk (e.g. if the percent of a procedure's products stocked in the kiosk is lower than the pre-determined threshold). In these scenarios, the procedures are instead filled by the shelves. This process begins with a hospital staff member selecting the procedure they need to fill on a tablet next to the shelf. The shelf then illuminates the bins containing the necessary products so the hospital staff member can pick up the correct items.

AITA™ Mobile Hubs are smaller, more portable devices that hold 48 products each. The hubs are placed near operating rooms (ORs) so they can enable convenient on-demand picking of add-on sutures. On-demand picking occurs when the products selected from the kiosk or shelf are not enough to complete a procedure. This could happen for multiple reasons including variability in patient anatomy, unexpected complications, and ad-hoc changes to the surgical plan. When additional products are needed, a member of the care team can pick up the items from a nearby hub instead of the farther-away stock room, saving valuable time.

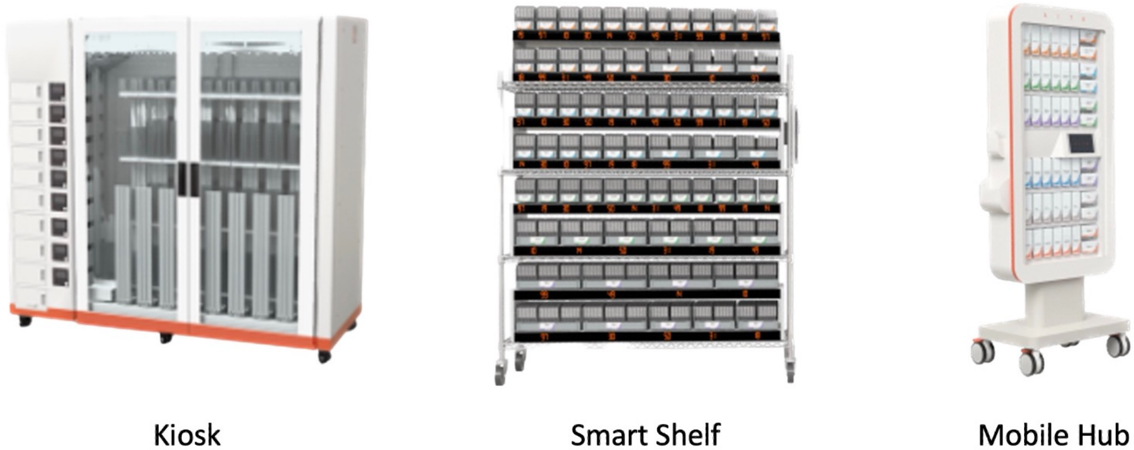

Figure 1-5: Devices used in the AITA™ Smart System

Through these devices and a cloud-based data insights engine, the AITA™ Smart System can greatly improve hospitals' inventory management. Specifically, AITA™ can benefit hospitals by enabling greater efficiency, standardization, simplification, and sustainability. On the efficiency front, AITA™ saves valuable time by making it easier to pick, restock, and replenish sutures. In fact, for a hospital that picks 8 sutures per procedure, the kiosk could pick up to 236 cases per day, saving 900 hours of manual picking per year [9]. AITA™ also helps with standardization, because it can highlight which sutures surgeons are using for the same types of procedures. Leveraging this data to minimize variability in suture usage would allow hospitals to lower costs and ensure better accuracy when ordering inventory. AITA™ also lowers costs and improves ordering accuracy by enabling simplification. Product usage data generated by AITA™ can help hospitals identify opportunities to eliminate redundant and underused SKUs. Finally, AITA™ improves sustainability by reducing the number of unopened sutures that are wasted after a procedure. For a hospital that wastes 18% of its picked sutures, the easy AITA™ returns system can save up to $80,000 worth of products per year [9].

1.2 Problem Statement

Before any AITA™ device (a Kiosk, Smart Shelf or Mobile Hub) can be installed in a hospital, the AITA™ team needs to create a planogram (POG) that dictates its set-up. Planograms show which products to stock, how many of each product to stock, and where within a device each product should be located. 1-6 shows an illustrative example of a planogram for one rack of a Smart Shelf. The rack has 12 columns and 8 rows, and the planogram shows which products go into each of the 96 slots. The planogram is also color coded according to product group (e.g. green could represent Prolene sutures, orange could represent Ethilon, etc.).

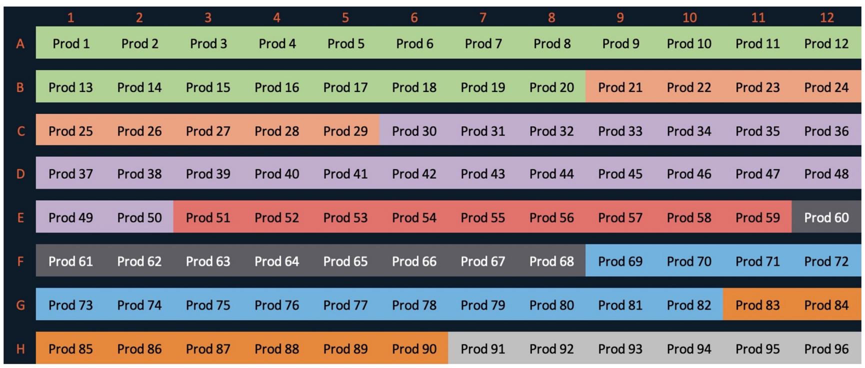

	1	2	3	4	5	6	7	8	9	10	11	12
A	Prod 1	Prod 2	Prod 3	Prod 4	Prod 5	Prod 6	Prod 7	Prod 8	Prod 9	Prod 10	Prod 11	Prod 12
B	Prod 13	Prod 14	Prod 15	Prod 16	Prod 17	Prod 18	Prod 19	Prod 20	Prod 21	Prod 22	Prod 23	Prod 24
C	Prod 25	Prod 26	Prod 27	Prod 28	Prod 29	Prod 30	Prod 31	Prod 32	Prod 33	Prod 34	Prod 35	Prod 36
D	Prod 37	Prod 38	Prod 39	Prod 40	Prod 41	Prod 42	Prod 43	Prod 44	Prod 45	Prod 46	Prod 47	Prod 48
E	Prod 49	Prod 50	Prod 51	Prod 52	Prod 53	Prod 54	Prod 55	Prod 56	Prod 57	Prod 58	Prod 59	Prod 60
F	Prod 61	Prod 62	Prod 63	Prod 64	Prod 65	Prod 66	Prod 67	Prod 68	Prod 69	Prod 70	Prod 71	Prod 72
G	Prod 73	Prod 74	Prod 75	Prod 76	Prod 77	Prod 78	Prod 79	Prod 80	Prod 81	Prod 82	Prod 83	Prod 84
H	Prod 85	Prod 86	Prod 87	Prod 88	Prod 89	Prod 90	Prod 91	Prod 92	Prod 93	Prod 94	Prod 95	Prod 96

Figure 1-6: Illustrative planogram showing one rack of a Smart Shelf

Planograms are created when the AITA™ Smart System is first installed in a hospital, and are updated periodically throughout the year. Unfortunately, the current process to create and update these planograms is manual, time consuming, and imperfect. It can take up to 2 days to create a new planogram, and up to 4 hours to update it. These processes can be repeated multiple times for a single installation, because the AITA™ team collaborates with hospitals to iterate on the POG designs. For example, providers might request product substitutions, provide a list of exclusions, or impose other constraints. Finally, since the planogram creation process is manual, there is no guarantee that the final product will correctly meet the required design rules. Manual errors could result in delays, further iterations, or confusion when the planogram is installed or implemented.

1.3 Project Motivation

This project focuses on creating a new, automated planogram design process that will address the current challenges and provide value for patients, providers, and the AITA™ team. The existing process is time consuming, requires frequent iterations, and can yield imperfect results. Meanwhile, an automated process would enable fast planogram creation, easy iteration with customizable parameters, and a solution that will objectively guarantee that all the design requirements are satisfied.

This automated planogram design process will provide value across the entire chain of care – from Ethicon to providers to patients. At AITA™, the data science team will benefit from the project, because it will save weeks' worth of work every year. The team can reallocate this time to other value-add activities, rather than spending their capacity on the current highly-manual planogram design process.

The automated process will benefit the AITA™ business development on two fronts. First, the team can leverage the models' easily adjustable parameters and fast turnaround times to have smoother interactions with hospitals. Under the current process, when hospital staff members request changes to planograms, it takes over an hour to create an updated result. Under the new process, a sales team member could easily make the changes while in the meeting, and the model would give a revised planogram in less than two minutes, greatly improving customer experience and satisfaction. The automated process will also improve AITA™ business development because it will shorten implementation time at each hospital. This will ease the team's workload, enabling them to grow their revenue by on-boarding more hospitals each year.

Outside of AITA™, this project creates value for both providers and patients. For example, the new model will optimally select kiosk products to ensure that the maximum possible number of surgeries are picked using the kiosk. Since the kiosk automatically picks products for procedures, increasing the number of cases it handles will free up more time for providers. Providers will also save more time, because the automated process will make implementation and re-designs faster. Reducing the time providers spend on these types of manual inventory management tasks will ultimately benefit patients, because providers will have more capacity and energy for patient care.

1.4 Project Overview

This project centers on creating three models that will automatically and quickly create planograms that are error-free and easily adjustable. Each model is specifically developed for a different AITA™ device: the Kiosk, Smart Shelf, or Mobile Hub. For each model, a user will be able to upload hospital data, set design parameters, and receive an output in much less time than it usually takes to manually design a planogram. The output should contain all the information needed for implementation, including product lists and visualizations.

While the general steps needed to run the models will be largely similar, each model has a unique set of objectives, algorithms, and parameters. The kiosk model aims to maximize the kiosk's coverage rate of historical procedures. In other words, the model will choose which products to put in the kiosk so that the device will have the SKUs needed to fill as many procedures as possible. So, the kiosk model begins with an optimization algorithm that chooses the set of products that maximize procedure coverage. Next, the kiosk model will allocate any extra bins to the most popular products. The model will then assign each product to a location within the kiosk in order of usage – where the most popular products are at the bottom, closest to the kiosk's picking arm.

Conversely, the goal of the shelf model is to store a hospital's products in a logical order, so that items are easy to find and restock. Shelves can have multiple racks (usually 3-5), so to minimize confusion, it is best if each product group (e.g. Prolene, Ethilon, etc.) is confined to just one rack. Thus, the first step of the shelf model is an optimization algorithm that minimizes the number of product groups that are split between multiple racks. Next, to make individual products easier to find, the model will sort the products within each group by key characteristics (e.g. decreasing in terms of size).

Finally, to facilitate add-on picking, the hub model should hold the most popular surgical products for the operating rooms (ORs) closest to it. These products should be sorted in a logical order within the hub to make picking and re-stocking easier. So, the model starts by receiving a list of the ORs that each hub will serve and the procedures that took place in those ORs historically. The model will calculate the total usage levels of each product and will select the most popular products to stock in the hub. Finally, like the shelf model, the hub model will sort the products according to product group and product characteristics such as size.

1.5 Thesis Structure

This thesis is organized into 7 chapters:

- Chapter 1 provides background information on surgical sutures, the AITA™ Smart System, and overall project objectives.
- Chapter 2 dives into strategies taken by other researchers to create planograms.
- Chapter 3 summarizes the main data sources used in this project.
- Chapter 4 showcases the results and underlying logic for the kiosk planogram creation model. It explains how the model uses optimization to select products and uses sorting rules to determine the products' placement in the kiosk.
- Chapter 5 details the results and methodology of the shelf planogram creation model, including the optimization formulation used to assign product groups to racks, and the sorting process used to organize products within each group.
- Chapter 6 explains the details of the hub planogram creation model and its corresponding results. This includes the logic used to select popular products, and the process used to sort the chosen products within each hub.
- Chapter 7 concludes with an overview of the models' results and impact, and recommends areas for future work.

Chapter 2

Solution Strategies and Approaches

2.1 Retail Planograms

Planogram (POG) creation is a task that appears in many settings outside of hospital inventory management. Most notably, planograms are widely used in the retail industry to show where merchandise should be displayed in a store. The POGs indicate which shelf each product should be placed in, as well as the location of each product within every shelf.

Retail planograms have a wide variety of goals. Their primary objective is to maximize sales while meeting a set of pre-defined design constraints. For example, large grocery stores will want to maximize their shelf space by displaying as many products as possible. This will drive revenue by giving customers more options and by reducing the likelihood of stock-outs. On top of shelf space optimization, stores must incorporate a set of rules and limitations in their planogram design process. A grocery store, for instance, would need to stock products from the same category in the same area (e.g. placing frozen foods in the same section of the store and placing chips on the same shelf). Other considerations could include agreements with brands to place certain products on eye-level racks, allocating multiple spots to popular products, and stocking items to match the layout of other store locations.

To meet these varied, and sometimes seemingly contradictory goals, retailers have developed a number of techniques to automatically design their planograms. This section delves into some of the most popular approaches, and explores how they can be adapted for the purpose of planogram creation for hospital inventory management. In particular, the knapsack optimization approach can be used to help maximize kiosk coverage, while a modified bin packing optimization approach can be used to help allocate product groups to racks on shelves.

2.2 Optimization Models

Optimization is a technique that is commonly used to design planograms in retail settings. Optimization models work within the confines of a given set of conditions to find the mathematically best solution for a given problem. For example, an optimization model could be used to maximize the number of products placed on a single shelf or to minimize the costs of all the products stocked in a store. While the parameters of every optimization model can be different, each one adheres to the same general structure and has an objective function, decision variables, and constraints.

2.2.1 Objective Functions

Objective functions define the goal of a problem. Mathematically, they represent the value to be maximized or minimized by a model. In their paper "Maximizing Revenue Through Two-Dimensional Shelf-Space Allocation", Geismar, Dawande, Murthi, and Sriskandarajah employ the commonly-used objective function of maximizing revenue [12]. Under their process, every product is given a revenue potential value based on historic demand for the item. The model then chooses to stock the products that maximize total revenue potential, while adhering to constraints such as minimum stocking quantities for certain items.

Objective functions for planograms can vary depending on the needs of a business. Kurtz's paper, "Planogram Optimization in Support of Small Format Retail Inventory Management" shows how inventory decisions can be specially optimized at urban Target stores with less than half the square footage of traditional locations [17]. To handle these space-limited cases, Kurtz and his team created a metric called "Fit" which numerically represents how well items fit on a shelf. The metric is calculated using a combination of an item's likelihood to generate backroom inventory, and how many days a standard shipment of the item can stay in stock on a shelf. Kurtz then creates a model with the objective function of maximizing Fit within the size limits of each shelf.

2.2.2 Decision Variables

Decision variables are the components that can be changed when solving an optimization model. For example, Ostermeier, Dusterhöft, and Hübner aim to maximize total store profit with the decision variable x_i that specifies how many facings x each item i should have [22]. This means the model will select the number of shelf slots for each product, in such a way that the store's total profit will be maximized. For instance, if item i is a television at an electronics store, the model might decide to give it just one facing ($x_{television} = 1$), because it is large and would take up a lot of shelf space. On the other hand, if item i is Coca-Cola at a grocery store, the model might decide to give it five facings ($x_{Coca-Cola} = 5$), because it is a popular item that would sell out faster.

2.2.3 Constraints

It is important to impose constraints in an optimization model to ensure that planograms meet the necessary design requirements. For example, in the paper "A practical approach to the shelf-space allocation and replenishment problem with heterogeneously sized shelves", Dusterhöft, Hübner, and Schaal want to choose which products should be stocked on a shelf to maximize profit [8]. However, they must consider key constraints to ensure that their solution is implementable in a real retail setting. Shelf constraints stipulate that items cannot be taller or deeper than the shelf they are placed on and that the combined width of all the products cannot be wider than the length of the shelf itself. Meanwhile, product constraints specify the maximum number of spots on a shelf that a product can be given and the upper and lower bounds of the number of shelves a product can occupy.

Planogram constraints can be further adjusted to meet businesses' different design requirements. Czerniachowaska and Hernes also created a model to maximize retail profits, but they added in additional constraints to ensure adaptability [7]. Their multi-shelf constraints restrict products' orientations so that only one side is showing for every product (e.g. all cereal boxes are facing the front, or all soda cases have their sides facing the customer). Then, their shelf segment constraints allow retailers to specify which product groups should be placed towards the center of each shelf, and which should be placed nearest to the aisles.

2.2.4 Optimization Approaches

While all optimization models can be customized by adjusting their objective functions, decision variables, and constraints, there are certain types of problem formulations and solutions that show up more often in planogram creation processes.

One of the most popular frameworks is the knapsack approach. In general, the knapsack problem stipulates that there is a fixed-sized space in which you can fit items. The goal is to choose the set of items that maximize or minimize a given metric. For example, say that a hiker has 10 items they could take on their trip, each with their own weight and utility value. However, the hiking backpack can only hold 20 pounds of items. Under the knapsack formulation, the hiker must pack products into their hiking backpack such that their utility is maximized, but the total weight of their packed products does not exceed 20 pounds.

In the retail industry, planogram creation is often framed around the knapsack problem, where the shelf represents the limited-capacity knapsack and the retailer's products are the items to place in the constrained space. Binguler, Bulkan, and Agaoğlu created a planogram optimization model based on the knapsack problem that considers both capacity constraints and the cross-price elasticity of products placed next to one another [5]. Meanwhile, Yang created a heuristic to simulate a solution for the knapsack problem that allocates shelf space to products with the greatest profit to display length ratio [27]. In this project, knapsack optimization is used to select the kiosk products that will maximize the procedure coverage rate.

The bin packing problem is another common framework applied to the planogram creation process. In the bin packing problem, the challenge is to fit a collection of differently sized items into as few fixed-capacity bins as possible. For example, say there are bins with a capacity of 10 units each and three items (items A, B, and C) that are 4, 5, and 8 units respectively. One potential solution would be to give each item a bin, leading to 3 bins being used in total. However, bin packing optimization would find the solution of putting items A and B into one bin (using 9 of 10 total units) and item C into a second bin (using 8 of 10 total units). This reduces the number of bins needed from 3 to 2.

In planogram creation, the bin packing problem is analogous to fitting a store's products into as few shelves as possible. This would minimize wasted shelf space, enabling retailers to fit more products within their footprint. This was the approach taken by Bai in the paper "An Investigation of Novel Approaches for Optimizing Retail Shelf Space Allocation" [2]. Bai used the one-dimensional bin packing problem as a proxy for fitting products into a limited amount of shelf space, and proposed a solution that leverages a hyper-heuristic optimization algorithm. In this project, a bin packing optimization algorithm is used to fit product groups into as few shelf racks as possible, without splitting up single product groups into multiple racks.

While the knapsack problem and bin packing problems are the most common frameworks used to conceptualize planogram optimization, they are not the only options available. For example, Gecili and Parikh partitioned planograms into different product families using a Particle Swarm Optimization approach that iteratively searches for the best candidate solution [11]. Meanwhile, Hübner, Schäfer, and Schaal created a planogram model that maximizes retailers' profits using a genetic algorithm that modifies its candidates in search of an optimal solution [13]. Finally, Gencosman and Bengen created a shelf space allocation model that uses logic-based Benders decomposition to break down the linear programming problem into smaller, more efficiently solved subsets [6].

2.3 Rule-Based Models

While optimization is a commonly employed technique in the planogram creation process, it is not the only option. In fact, there is a school of research centered around rule-based methods for creating shelf layout designs. These rules dictate where products should be placed within a store and the items' ideal location within any given shelf. Retailers must pay special attention to these decisions, because they have a meaningful and tangible impact on consumers. Mehta and Chugan found that customers' purchase behavior is positively correlated to well-structured stores that are laid out in a logical way and that have "appropriately occupied" shelves (i.e. not too crowded nor too sparse) [18].

The book, *Assortment and Merchandising Strategy* by Constant Berkhout lays out a set of universal merchandising guidelines that retailers can use to drive revenue [3]. One of the main guidelines is that stores should be structured around shopper decision trees. A shopper decision tree demonstrates the trade-offs consumers make when purchasing a product. For example, when a consumer wants to buy pasta, they might start by looking for a certain brand and then, within that brand, they might look for a particular pasta shape. Berkhout explains that retailers should align their shelf layouts with shoppers' decision trees to minimize confusion and drive revenue. So, the pasta aisle should be split into different vertical sections, where each section corresponds to a different pasta brand. Then, within each vertical section, the pasta products should be arranged such that all the products with the same shape are placed together. For example, within the Barilla section, the whole wheat penne, white penne, and gluten-free penne would be placed next to each other.

Rooderkerk and Lehmann studied shopper behavior and determined that consumers are less confused when store layouts match the decision tree described by Berkhout [23]. Matching shelves to shoppers' internal categorizations improved choice satisfaction, increased perceived assortment variety, and bolstered shoppers' opinion of the store by reducing perceived complexity.

In the article "The Role of Store Layout and Visual Merchandising in Food

Retailing", Štulec, Petljak, and Kukor offer additional, practical, rules to improve planogram design [26]. For example, they recommend placing everyday necessities (e.g. oil, sugar, and flour) on the bottom shelves, because consumers will go out of their way to find them. Top shelves, on the other hand, should store more expensive products, because they will be less likely to be damaged. Finally, the most popular products, or products that retailers have been paid to promote, should be placed in the middle of the rack to meet shoppers' eye level.

These rules are created by retailers to codify and respond to consumer behavior. Once stores know how shoppers think, they can tailor their planograms to drive profits. Štulec, Petljak, and Kukor show that well-designed planograms can encourage shoppers to explore the store and make unplanned purchases, without feeling like they have wasted their time.

2.4 Hybrid Approach

Given the benefits of both optimization and rule-based models, it is no surprise that retailers often use a hybrid approach to design their planograms. One commonly used methodology involves quantifying consumer behavior rules and then optimizing shelf layouts based on the behavioral values. For example, Zhao, Zhou, and Wahab created a function that quantifies consumer purchasing behavior by estimating demand for a particular item based on its shelf space, display location, and spatial relationship with other items. They then use optimization to maximize total profits given each item's estimated demand [28].

Other hybrid approaches identify different ways to combine optimization and rule-based methodologies. Miller, Smith, McIntyre, and Achabal model customer behavior rules by creating consideration sets that predict what products a shopper will look at, based on their product attribute preferences (e.g. size or brand) [19]. They then create an optimization model to maximize profit based on which products would be the most considered. Meanwhile, Nierop, Fok, and Franses created a statistical model to measure the impact of shelf layout on sales [21]. This implicitly incorporates customer behavior rules, because it quantifies behaviors such as purchasing more items if they are at eye level. Then, they optimize this sales valuation to create a profit-maximizing shelf layout. Finally, Bianchi-Aguiar, Silva, Guimarães, Carravilla, and Oliveira created a tree diagram that mirrors shopper decision trees by explicitly defining product hierarchies that show how products should be categorized and grouped on shelves [4]. Their optimization function uses this tree to constrain the physical locations of the products within the planogram, in a way that aligns with merchandising best practices.

These hybrid approaches, and other similar methodologies employed by retailers, help to bring the benefits of both customer behavior rules and optimization to planogram creation. Under a joint model, retailers can create planograms that appeal to shoppers and maximize profit drivers such as predicted revenue and demand.

2.5 Project Approach

This project leverages hybrid models to create planograms for AITA™ 's Kiosks, Smart Shelves, and Mobile Hubs. The kiosk model uses a knapsack optimization technique to select the set of products for the device, then uses product usage rules to place items in individual slots within the kiosk. The shelf model uses bin-packing optimization to choose which product groups belong on each rack, and uses sorting rules to order products within each group. Finally, the hub model uses rules dictated by the hub's location within the hospital to choose the products each device will carry. It then uses sorting rules to determine the location of each product within the hub.

By using a combination of optimization and placement rules, the Kiosk, Smart Shelf, and Mobile Hub models create planograms that better address the needs of AITA™ 's customers. Hospital staff will have a more positive experience managing inventory because their planograms will be designed according to rules about their behavior. For example, like retail customers, healthcare providers expect to see products of the same type in the same area of a shelf. So, automatically sorting inventory by product group will make it easier for staff to find and restock items. In addition, adding an optimization component to the models will help to improve the efficiency of AITA™ 's planograms. For example, maximizing the number of procedures that can be automatically picked by the kiosk would save providers' time by reducing the number of products that need to be collected manually.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Data Sources

The AITA™ Kiosk, Smart Shelf, and Mobile Hub models are created using data collected from hospital clients and from libraries maintained by the AITA™ team. This section describes all the relevant and necessary datasets in more detail.

3.1 Preference Cards

Hospitals keep and regularly maintain preference cards for every surgeon / procedure pair. The cards detail which products the surgeon needs to complete that procedure. 3-1 shows an illustrative example of a preference card file, which has 3 primary columns: Preference Card ID, Product SKU #, and Manufacturer. The Preference Card ID is a unique identifier for each preference card. Product SKU # is a unique identifier for every product SKU. Manufacturer shows what company created each SKU.

In other words, in 3-1, 1010 could be the ID of the preference card that Dr. X uses when performing an appendectomy. The preference card file tells us that Dr. X needs 3 products for an appendectomy, P999, P555, and P333. These are the products that would need to be gathered from the stock room and placed in the operating room before the procedure. It should also be noted that the same products can be used across different surgeon / procedure pairs. For example, say that Product SKU # P999 represents an undyed Dafilon suture with a suture size of 4-0 and a length of 45 cm. That same product is used in preference card 1010, and also in card 3030, which would belong to a different surgeon / procedure pair.

Preference Card ID	Product SKU #	Manufacturer
1010	P999	ETHICON
1010	P555	CARDINAL
1010	P333	ETHICON
2020	P444	ARTHREX
2020	P111	ETHICON
3030	P555	CARDINAL
3030	P999	ETHICON
3030	P222	ETHICON
3030	P666	ARTHREX

Figure 3-1: Illustrative example of a preference card file

3.2 Procedure Schedule

Every hospital has a procedure schedule that lists all its past and planned procedures. An illustrative example of a procedure schedule is shown in 3-2. Schedules have 3 main pieces of information: OR Case IDs, OR Rooms, and Preference Cards. Every time a patient is brought into an operating room (OR), their case is assigned a unique ID (the OR Case ID). The OR Room column shows where that case's procedures are being conducted. The preference card ID column shows which preference cards are assigned to each case.

OR Case ID	OR Room	Preference Card ID
1	OR A	3030
2	OR E	4040
2	OR E	2020
3	OR C	1010
4	OR B	3030

Figure 3-2: Illustrative example of a procedure schedule file

For example, OR Case ID 2 could represent the situation when patient Y has two procedures scheduled in OR E at the same time: a joint replacement and an arthroscopy. Preference card 4040 could represent the products Dr. Z needs for a joint replacement, and preference card 2020 could represent the products that Dr. Z needs for an arthroscopy. Products for both procedures would be collected and brought to OR E before the surgeries are scheduled to start.

3.3 Suture Database

AITA™ maintains a suture database with information about every product they have encountered while implementing the Smart System at different hospitals. As shown in 3-3, the database has 7 primary columns that the team uses to sort and manage inventory. The Product SKU # is the unique identifier for every product, and it ties to the SKU # listed in the preference card file. The Product Group column specifies the type of every product. Structured Description lists the name of each product, in a form that would be easily understood by hospital staff. Suture Size represents the diameter of the suture needle. The Closures/Ties column shows if products have a loop, tie, or reel, or if they are a standard suture (in which case the column is left blank). The Needle Point column describes the type / point shape of the needle. And the Needle Length column shows the length of each suture's needle in millimeters.

3-3 shows the attributes for 10 illustrative products. Each product belongs to a different group / type, and has its own unique product SKU #. P222 has a tie, P555 has a loop, P888 has a reel, and the rest of the sutures are standard. There are a wide variety of needle points represented in the example, and the illustrative needle lengths range from 4mm to 37mm.

Product SKU #	Product Group	Structured Description	Suture Size	Closures / Ties	Needle Point	Needle Length
P000	DAFILON	10-0 DAFILON BLACK 1x6" (15cm) DRCM4 DA	10-0		TAPERCUT	4
P111	CHROMIC GUT	2-0 CHROMIC GUT UNDYED 1x27" (70cm) CT-1	2-0		TAPERPOINT	36
P222	PLAIN GUT	0 PLAIN GUT YELLOWISH TAN 1x54" (135cm) TIE	0	TIE	NULL	
P333	NUROLON	4-0 NUROLON WHITE 1x18" (45cm) PC-1 MULTIPASS	4-0		CONVENTIONAL CUTTING	13
P444	PDS PLUS	6-0 PDS PLUS VIOLET 1x27" (70cm) TF	6-0		TAPERPOINT	36
P555	PDS	1 PDS VIOLET 1x96" (244cm) CT-1 LOOPED	1	LOOPED	TAPERPOINT	36
P666	MONOCRYL	0 MONOCRYL UNDYED 1x36" (90cm) CTB-1	0		BLUNT POINT SAFETY	36
P777	MONOCRYL PLUS	3-0 MONOCRYL PLUS VIOLET 1x27" (70cm) RB-1	3-0		TAPERPOINT	17
P888	DAGROFIL	0 DAGROFIL GREEN 1x98" (250cm) REEL	0	REEL	NULL	
P999	V-LOC 180	0 V-LOC 180 CLEAR 1x18" (45cm) GS-11	0		REVERSE CUTTING	37

Figure 3-3: Illustrative example of the suture database file

3.4 Suture Dimension Library

Alongside the suture database, AITA™ also maintains a library documenting the size of each suture (3-4). The file has two main columns. Product SKU # is the unique identifier for each product, and it ties to the SKU #'s used in the suture database and in hospitals' preference cards. The Size column specifies if the product is one of three possible sizes: small, medium, or large.

Products' sizes are dependent on the physical dimensions of their packaging. They are categorized into the groups small, medium, and large depending on which kiosk and hub slots they can fit into. When AITA™ is building its kiosks and hubs, it installs one vertical column at a time. The columns can either hold small products or can hold large/medium products. In a large/medium column, the AITA™ team can add padding around individual bins to turn a large bin into a medium bin. For example, if a large column usually has 10 bins, and AITA™ adds padding to 4 of the large bins, the column will be converted into storage for 6 large products and 4 medium products. The team could also theoretically convert a large bin into a small bin, but this is rarely done, because there are often enough large/medium products to fill the given large/medium slots.

Product SKU #	Size
P000	Small
P111	Small
P222	Large
P333	Medium
P444	Large

Figure 3-4: Illustrative example of the suture dimension library

Chapter 4

Kiosk Planogram Model and Results

4.1 Current Kiosk Planogram Creation Process

In the current kiosk planogram creation process, products are selected and sorted manually. To start, the AITA™ team looks at hospital purchasing data to find the most frequently purchased products. It then chooses the S most popular small products and the L most popular medium and large products, where S is the number of small products the kiosk can hold, and L is the number of medium and large products the kiosk can hold. The team then sorts these products by purchase frequency, and places them in the kiosk with the most popular products in the bottom left-hand corner, closest to the kiosk's picking arm.

Unfortunately, under this manual approach, the kiosk planogram creation process is not as efficient or as accurate as it could be. From an efficiency perspective, the end-to-end process of creating a planogram takes approximately 1 day (8 hours), because it requires manually cleaning and analyzing the hospital's purchasing data. From there, if the hospital wants to iterate on the design by making updates such as adding or removing certain items, it can take another 1-2 hours to implement the necessary changes and check the updated results.

Furthermore, even after the time-consuming process of creating and updating kiosk planograms, they are not always as accurate as hospitals and the AITA™ team would like. Kiosks are meant to save time for hospital staff members by automatically picking products for procedures and placing the necessary products in bins where they can be quickly and easily collected. So, kiosks save the most time if they can supply products for (i.e. cover) the largest number of procedures. Unfortunately, the manual planogram creation process does not always yield the highest possible coverage rate. For example, consider the case where 100% of a procedure's products must be in the kiosk for the order to be sent to the kiosk. If the kiosk is stocked with only the most popular products, many orders would not be sent to the device, because their corresponding procedures / preference cards often contain one or more specialized, less popular items.

4.2 Kiosk Model Objectives

This project's kiosk planogram creation model was built to address the efficiency and accuracy challenges of the manual process. From an efficiency perspective, the model should work automatically and quickly – taking significantly less time to create and update planograms. From an accuracy perspective, the model should use optimization to maximize the kiosk's coverage rate of procedures. As shown in 4-1, the kiosk model has two main steps that enable it to achieve these objectives: 1) selecting products for the kiosk and 2) placing the products in order within the kiosk.

The selection process (step 1) begins with the user inputting a list of parameters and files that the model needs to design the planogram. The model then uses a modified version of the knapsack problem to pick the products needed to maximize the kiosk's coverage rate of historical procedures. The placement process (step 2) uses this list of kiosk products as an input. The model uses historical procedure information to calculate the past usage level of each kiosk product. If there are any empty bins in the kiosk, the model allocates multiple bins to the most popular SKUs by proportionally assigning one or more extra slots based on product usage levels. The model then puts the products in usage order starting with the most popular items in the bottom left corner of the device. This placement allows the kiosk to fill orders more quickly, because the picking arm starts in the bottom left – closest to the most used products.

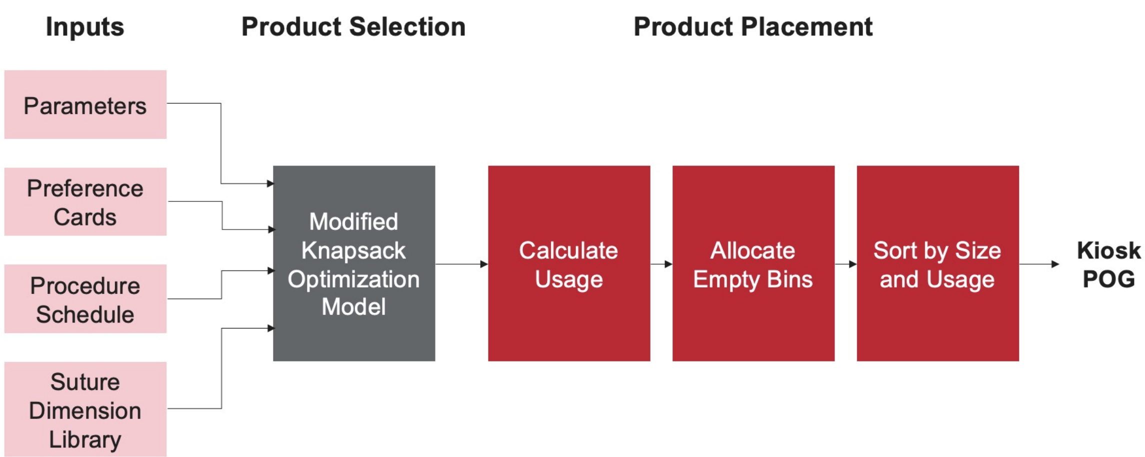

Figure 4-1: Process flow to automatically create kiosk planograms

4.3 Kiosk Model Inputs

The next two subsections detail the parameters that the user must fill in and the files that the user must upload to successfully run the kiosk planogram creation model.

4.3.1 Input Parameters

Table 4.1: Input Parameters for the Kiosk Model

Parameter	**Description**
Partial Picking Threshold	Decimal between 0 and 1; Specifies the threshold for a preference card to be considered "filled" by the kiosk (e.g. when the threshold is 0.7, 70% of a preference card's products must be in the kiosk for it to be filled by the kiosk); It is not required to always cover 100% of products in a preference card, because hospital staff can manually pick the remaining products from the shelf.
Number of Rows/Columns	Number of rows/columns in the kiosk
	Continued on next page

Table 4.1 – continued from previous page

Parameter	Description
Number of rows in a rack	Number of rows comprising a kiosk's rack; E.g. if a kiosk has 6 rows and each rack holds 2 rows, the kiosk will have 3 racks
Kiosk Capacity	Total number of products the kiosk can hold (integer)
Large Capacity	Number of large products the kiosk can hold (integer)
Medium Capacity	Number of medium products the kiosk can hold (integer)
Location of L/M columns	Which columns can hold Large and Medium products; assumed that the other columns hold small products
Excluded Products	List of products to manually exclude from the kiosk (e.g. if the hospital wants to retire a certain product SKU from use)
Maximum bins	The maximum number of bins assigned to each product; e.g. if the model says that Product 1 should have 5 new bins, but the max number of bins is 3, Product 1 will only be allocated 3 total bins
Manual Bin Count	Manually specify how many bins a product should have; overrides bin count given by the algorithm

4.3.2 Input Files

Table 4.2: Input Files for the Kiosk Model

Input File	Description	Usage
Preference cards	Every preference card and its associated products	Determine which products will optimize coverage rates of preference cards

Continued on next page

Table 4.2 – continued from previous page

Parameter	Description	Usage
Procedure schedule	Lists how many times each preference card appeared in the historical procedure data over a certain time frame	Use in conjunction with preference cards to calculate the usage of products over a certain time period
AITA™ suture dimension library	Has the sizes (small, medium, or large) of each suture	Assess which products can go in the small, medium, and large columns

4.4 Kiosk Product Selection

4.4.1 Illustrative Example

After the user inputs the necessary parameters and uploads the required files, the model starts the first of its two phases: product selection. In the kiosk model, the goal of product selection is to choose products that maximize the coverage rate of historical procedures. The example shown in 4-2 showcases the definition of a coverage rate and illustrates how certain product combinations can lead to more favorable results.

In this example, there is a kiosk with 4 slots. The hospital does 3 types of procedures, and each procedure has its own set of necessary products (as dictated by the preference card). For example, Procedure 2 requires 2 products: green circle and pink drop. From the procedure schedule, we also know how many times each procedure was conducted in the past. Procedure 1 was the most popular procedure with 100 past occurrences, followed by Procedure 2 with 70 occurrences, and Procedure 3 with 50. Furthermore, assume that the partial picking threshold in this scenario is 1. This means that for a procedure to be considered "covered" by the kiosk, 100% of its products must be inside the kiosk. For example, if the kiosk holds only the green circle and yellow diamond products, Procedure 3 is not considered to be "covered" by the kiosk, because one of its products (33% of its total products) is not stocked.

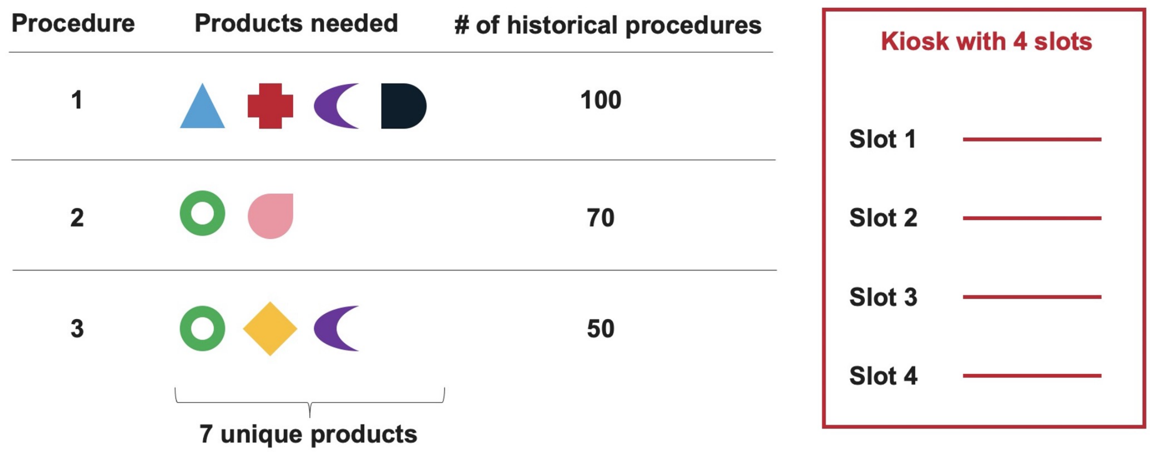

Figure 4-2: Illustrative scenario to select kiosk products

Given this scenario, the kiosk planogram model would have to decide which 4 products to stock in the kiosk. The model would aim to maximize the number of historical procedures covered by those 4 products. As shown in 4-3, the optimal 4 products to stock are the green circle, pink drop, yellow diamond, and purple moon. With these 4 products the kiosk can cover 120 historical procedures (70 from Procedure 2 and 50 from Procedure 3). If, for example, the planogram contained the products blue triangle, red cross, purple moon, and black semi-circle, the kiosk would have only been able to cover 100 historical procedures (from Procedure 1).

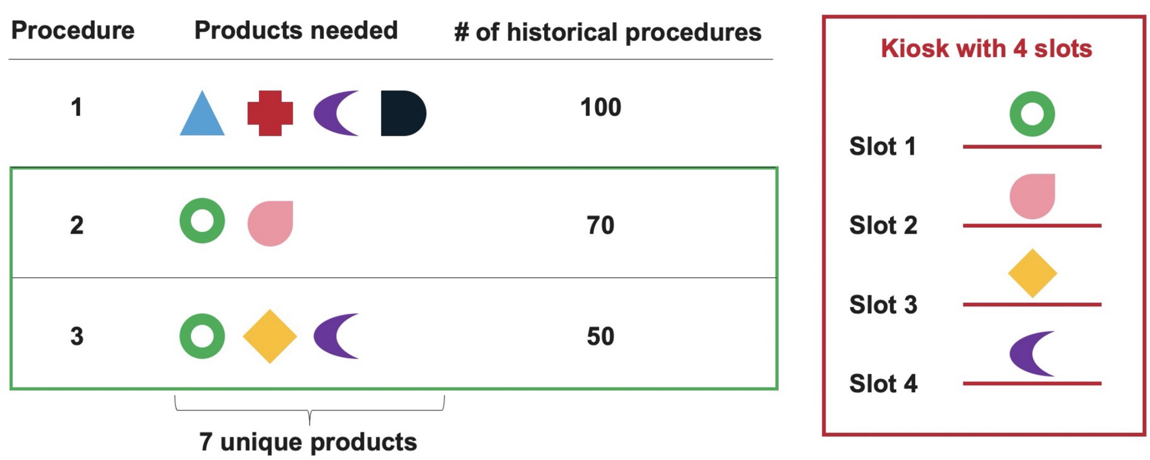

Figure 4-3: Optimized solution to the illustrative scenario

4.4.2 Optimization to Select Products

In real hospital settings, it is much more challenging to select the products needed to maximize coverage of historical procedures. In the illustrative example, there were only 4 kiosk slots, 3 procedures, and 7 products. At real hospitals, there could be hundreds of unique procedure types and products, and dozens of kiosk slots. So, to ensure maximum coverage, the model uses optimization to select products. The following formulation highlights the important features of the optimization model.

Sets

P Set of all products, indexed by p

P_L Set of all large products

P_M Set of all medium products

C Set of all preference cards, indexed by c

T_c Set of all products in each preference card c

Parameters

N Kiosk capacity (maximum number of products in the kiosk)

L Large product capacity (number of large products that fit in the kiosk)

M Medium product capacity (number of medium products that fit in the kiosk)

H Partial picking threshold (percentage over which preference cards are considered to be covered/filled by the kiosk)

h_c Number of historical occurrences of every preference card c (how many surgeries using preference card c took place)

d_c Total number of products listed in each preference card c

Decision Variable

$$x_p = \begin{cases} 1, & \text{if product p is in the kiosk} \\ 0, & \text{if product p is not in the kiosk} \end{cases}$$

Dependent Variables

$$f_c = \begin{cases} 1, & \text{if preference card c is covered (filled) by products in the kiosk} \\ 0, & \text{if preference card c is not covered (not filled) by products in the kiosk} \end{cases}$$

r_c Percentage of preference card c's products that are in the kiosk

Mathematical Formulation

Objective

$$\text{maximize} \sum_{c \in C} f_c h_c r_c \tag{4.1}$$

Subject to

$$r_c \leq \frac{\sum_{p \in T_c} x_p}{d_c} \text{ for all } c \in C \tag{4.2}$$

$$f_c \leq \frac{r_c}{H} \text{ for all } c \in C \tag{4.3}$$

$$\sum_{p \in P} x_p \leq N \tag{4.4}$$

$$\sum_{p \in P_L} x_p \leq L \tag{4.5}$$

$$\sum_{p \in P_M} x_p \leq M \tag{4.6}$$

The objective function 4.1 demonstrates that the model will choose which products x_p to add to the kiosk planogram. It will choose products with the intent of maximizing both the number of historical procedures covered by the kiosk (f_c h_c) and how filled / covered each preference card is (r_c). 4.2 and 4.3 explain how the metrics r_c and f_c are calculated. For example, assume that preference card c has 10 products, and 7 of those products are in the kiosk. r_c shows what percentage of the preference card's products are in the kiosk, so, in this case, r_c would equal 70%.

On the other hand, f_c is a binary indicator that is 1 when preference card c is filled / covered by the kiosk, and 0 when it is not. In other words, 4.3 shows that f_c

will equal 1 when r_c is greater than or equal to the partial picking threshold H, and it will equal 0 when r_c is less than H. In our example, suppose that the partial picking threshold H is 50%. Then preference card c will be filled by the kiosk ($f_c = 1$) because $70\% > 50\%$. However, if the threshold H was increased to 80%, then preference card c would not be "filled" by the kiosk ($f_c = 0$).

It is important to incorporate r_c in the objective function to ensure that procedures are maximally covered by the kiosk. For example, assume that the partial picking threshold H is 50%, and a preference card c has 10 products. If 5 of c's products are in the kiosk ($r_c = 50\%$), c will be considered filled/covered by the kiosk ($f_c = 1$). However, it would be better to put more of c's products in the kiosk ($r_c \geq 50\%$), provided that the total number of historical procedures covered are not reduced. That way, hospital staff will have to manually pick fewer products. When r_c is in the objective function, the model will codify this logic by increasing r_c as much as possible. For example, the model would prefer to have 100% of a preference card's products in the kiosk, over only 50%.

The remaining constraints 4.4, 4.5, and 4.6 ensure that the algorithm respects the capacity constraints of the kiosk. 4.4 states that the total number of products in the kiosk cannot exceed the kiosk's total capacity N. 4.5 stipulates that there can only be L large products in the kiosk, and 4.6 limits the number of medium products in the kiosk to M.

4.5 Kiosk Product Placement

After the first main step of the kiosk model (product selection), the model will have a list of which products should be put into the device. This brings the process to the second step, product placement, where the model must choose the order in which products should be stored in the kiosk. During product placement, there are 3 main decisions the model must make: 1) allocating empty bins, 2) organizing by product size, and 3) sorting products by usage.

4.5.1 Allocating Empty Bins

Sometimes, the capacity of a kiosk is greater than the number of products chosen in the product selection process. This can happen when a hospital has fewer kiosk-eligible products, or more consistent preference cards with high degrees of overlap from surgeon to surgeon. In these situations, there will be extra / empty bins that should be allocated by giving some kiosk products multiple bins.

Empty bins are allocated based on usage, where the most popular products are given multiple bins. This helps to improve operational efficiency, because kiosks with a larger stock of high-usage products will need to be restocked less frequently. Empty bin allocation follows the following rules.

Variables

E Empty bin count

K Kiosk capacity

C Count of products placed in the kiosk

P_E Set of the top E products with the highest usage, indexed by p

U_p Usage of product p

B_p Number of extra bins allocated to product p

Allocation Rule

$$E = K - C \tag{4.7}$$

$$B_p = \text{round}(\frac{U_p}{\sum_{p \in P_E} U_p} * E) \tag{4.8}$$

If a kiosk can hold 100 products, but the model only selects 90 products to add to the device, 4.7 shows that this will leave 10 empty slots in the kiosk. As explained by 4.8, these 10 empty slots are allocated based on product popularity. Assume that the top 10 most popular products have a total usage value of 500 ($\sum_{P_{10}} U_p = 500$). If the most popular product (product A) has an individual usage of 100, it will be allocated 2 of the 10 empty bins ($B_A = \text{round}(\frac{100}{500} * 10) = 2$). This process will continue for the remaining 9 most popular products.

4.5.2 Organizing by Product Size

Kiosk-eligible products come in three sizes: small, medium, and large. This attribute partially dictates product placement because different columns within the kiosk can hold different sizes. Hospitals and the AITA™ team work together to select how many columns should be allocated to small, medium, and large products. Since different hospitals carry a different mix of products, the final size configuration will be dependent on how many large and medium products appear in the preference cards.

4-4 shows an illustrative example of how a kiosk might be laid out by size. In this scenario, the kiosk has 5 columns and 6 rows, and the first 2 columns can hold large and medium products. The diagrams show that the columns and products can be organized in one of two ways – 1) with large and medium items in separate columns or 2) large and medium items mixed together in the same column.

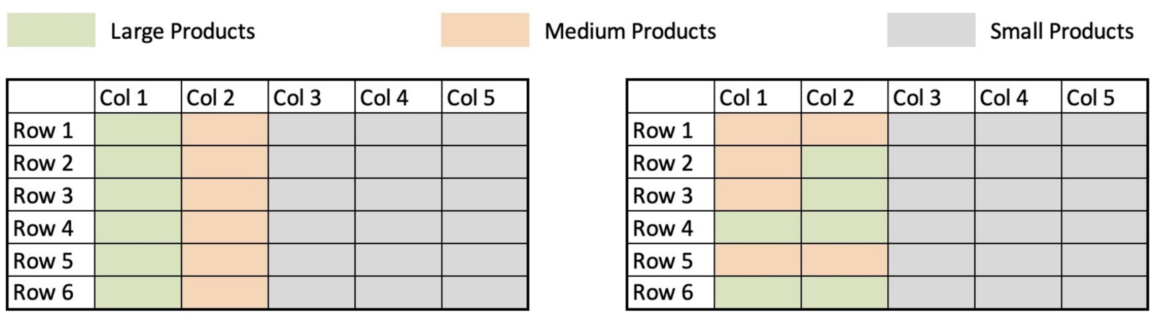

Figure 4-4: (Left) Large and medium products kept in separate columns; (Right) Large and medium products placed in the same column

4.5.3 Sorting Products by Usage

The final step of the product placement process is sorting products by their usage / popularity. The model starts by considering one size group (i.e. the small products, the medium products, the large products, or the medium and large products together). As shown in the formulation below, the model then uses a combination of the hospital's procedure schedule and preference cards to calculate each product's total usage.

Variables

C Set of preference cards, indexed by c

P Set of products, indexed by p

h_c Number of occurrences of preference card c in the historical procedure schedule

U_p Usage of each product p

$$x_{p,c} = \begin{cases} 1, & \text{if product p is in preference card c} \\ 0, & \text{if product p is not in preference card c} \end{cases}$$

Mathematical Formulation

$$U_p = \sum_{c \in C} x_{p,c} * N_c \tag{4.9}$$

Product usage is defined as the number of times a product would have been used in historical procedures. 4.9 demonstrates that this value can be calculated for any given product p by summing all the procedure occurrences for preference cards that contain p. For example, suppose that product A is present in 2 preference cards (card Y and card Z). Also assume that card Y was used 20 times in the procedure schedule and that card Z was used 10 times. Product A would thus have a total usage value of 30.

Once all the products in a given size group have usage values, they should next be placed in the kiosk according to usage order. The kiosk's hardware is structured such that the picking arm (the mechanism that grabs items from bins and puts them into collection lockers) is in the bottom left corner of the device. So, the most popular products should be in the bottom left, closest to the picking arm, to minimize the amount of time needed for the arm to collect all the products. The products are thus placed within their appropriately-sized columns in order of decreasing usage from left to right and from bottom to top. 4-5 shows illustrative usage values of products placed in the kiosk under two different size configurations.

Large Products | Medium Products | Small Products

	Col 1	Col 2	Col 3	Col 4	Col 5
Row 1	11	9	25	20	15
Row 2	13	10	40	35	30
Row 3	15	12	55	50	45
Row 4	16	14	70	65	60
Row 5	19	17	85	80	75
Row 6	20	18	100	95	90

	Col 1	Col 2	Col 3	Col 4	Col 5
Row 1	10	9	25	20	15
Row 2	12	11	40	35	30
Row 3	14	13	55	50	45
Row 4	16	15	70	65	60
Row 5	18	17	85	80	75
Row 6	20	19	100	95	90

Figure 4-5: Products sorted by usage when (Left) large and medium products have separate and when (Right) large and medium products have combined columns

4.6 Kiosk Model Outputs

After the model completes the processes of product selection and product placement, it outputs a file with a visual of the final planogram. 4-6 shows an illustrative view of an output. In this is example, the kiosk has 10 columns and 6 rows that are split into two racks of 3. The large products are in columns 1-2, the medium products are in 3-4, and the small products are in the remaining spaces. Furthermore, after the model's process to allocate empty bins, Product 5 has 3 bins in total and Product 20 has two total bins.

Large Products | Medium Products | Small Products

Rack 1	Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	Col 7	Col 8	Col 9	Col 10
Row 1	Prod 43	Prod 52	Prod 31	Prod 8	Prod 26	Prod 2	Prod 23	Prod 51	Prod 30	Prod 9
Row 2	Prod 4	Prod 22	Prod 3	Prod 39	Prod 15	Prod 53	Prod 37	Prod 11	Prod 6	Prod 41
Row 3	Prod 29	Prod 17	Prod 50	Prod 38	Prod 46	Prod 57	Prod 18	Prod 47	Prod 38	Prod 39
Rack 2	Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	Col 7	Col 8	Col 9	Col 10
Row 4	Prod 36	Prod 40	Prod 45	Prod 19	Prod 34	Prod 31	Prod 56	Prod 27	Prod 54	Prod 24
Row 5	**Prod 5**	Prod 14	Prod 25	Prod 33	Prod 1	Prod 42	Prod 10	Prod 44	Prod 7	Prod 55
Row 6	**Prod 5**	**Prod 5**	Prod 16	Prod 28	**Prod 20**	**Prod 20**	Prod 13	Prod 35	Prod 21	Prod 12

Figure 4-6: Illustrative kiosk model output (view of products within the kiosk)

4-6 shows the main view of the kiosk planogram output. However, the output also includes other views to help the AITA™ team review and confirm the result. The three views show different product attributes, all displayed like 4-6, in kiosk order. 1) The SKU layout (4-6) shows the product SKU number for each item. 2) The description layout shows the full product name of every item (e.g. the SKU number U202H corresponds to the full product name “5-0 Chromic Gut Undyed 1x27" (70cm) RB-1”. 3) The usage layout shows the individual usage values of every product. Taken together, these views can help hospitals and the AITA™ team understand why certain products are placed in particular locations within the kiosk planogram.

4.7 Kiosk Model Results

The kiosk model improves both the accuracy and the efficiency of the kiosk planogram creation process. On the accuracy side, the kiosk model improves the coverage rate of the planogram. In other words, the model-generated kiosk can fill more historical procedures than a manually-created kiosk. To prove this point, the kiosk model was run on a representative hospital's data. The hospital conducts over 10,000 procedures a year, has almost 1,000 unique preference cards on record, and carries an inventory of over 300 types of products. The hospital was already using a planogram built manually by the AITA™ team. 4-7 shows this manually-created planogram's coverage rate as it compares to a planogram built automatically by the kiosk model.

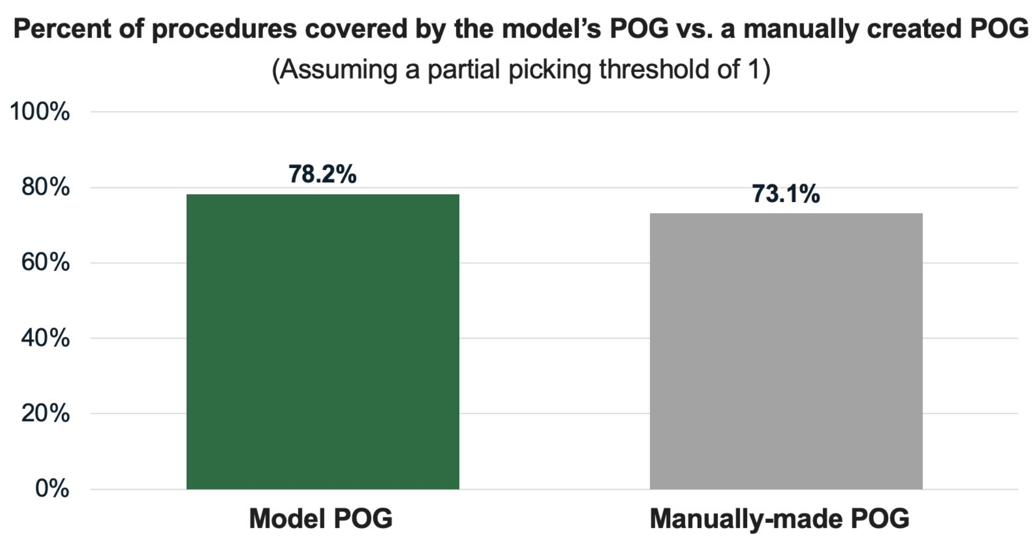

Figure 4-7: Coverage rate of a planogram (number of procedures "filled" by the kiosk divided by total procedures) built by the model vs. manually

4-7 shows that, when the partial picking threshold is set to 1 (indicating that 100%

of a preference card's products must be in the kiosk for the card to be considered "filled" by the kiosk), the model's planogram achieved a higher coverage rate than the planogram that was built manually for the hospital (78.2% vs. 73.1%). These planograms were built with the same constraints and criteria, but the optimization model was able to find additional opportunities to fill more historical procedures. In the case of this hospital, which conducts over 10,000 procedures a year, the model's coverage improvement would result in over 510 additional procedures being filled using a kiosk instead of being filled with manual picking. Assuming that it takes 15 minutes to pick out products by hand for one procedure, automating picking for 510 procedures annually would save 7,650 minutes, or 127 hours, a year for the hospital's staff members.

It is worth nothing though, that in both cases, coverage is being measured using historical procedures. So, if future procedures do not match the same trends as historical events, the coverage rate of the planograms will decrease. However, in the absence of predictive data, historical procedures must be used as a proxy for future trends. Hospitals and the AITA™ team can help mitigate potential discrepancies between future procedures and historical procedures by frequently updating the kiosk planogram.

The kiosk model also provides improved efficiency alongside improved accuracy. 4-8 shows the time saved when an automatic model is used to create and update a planogram, rather than a manual process. Overall, the AITA™ team can save 20.5 hours of work annually (15.7 + 4.8 hours) for every kiosk in every hospital they serve. 4-8 shows that it takes, on average, one day (8 hours) to manually create a planogram. This is because the AITA™ team must calculate individual usage rates for every product and must manually sort products within the kiosk by size and usage. Conversely, the automated kiosk model takes less than half the time to create a new planogram (roughly 3 hours). The time to run the model itself is fast (around 2 minutes), but the team still needs to clean the data, set the parameters, and check the output. Cleaning and preparing the data accounts for a majority of the time (2 hours), because the team must ensure the accuracy and consistency of the input files.

	Frequency	Steps	Model-based process	Manual process	Average Time Saved per Year
Creating a new kiosk planogram	2-4 times per year	Aggregate / clean data	2 hours	1 day (8 hours)	**15.7 hours** = (8 – 2.78) hours * 3 per year
		Adjust parameters	15 minutes		
		Run model	2 minutes		
		Check output	30 minutes		
Iterating on a kiosk planogram	2-8 times per year	Update data / parameters	15 minutes	1-2 hours	**4.8 hours** = (1.5 – 0.53) hours * 5 per year
		Run model	2 minutes		
		Check output	15 minutes		

Figure 4-8: Time taken to create and update kiosk planograms when using the model vs. a manual process

These time savings are further compounded, because the AITA™ team usually creates new planograms for hospitals 2-4 times per year. Since hospitals' procedures and preference cards are frequently being updated, new planograms must be made where every product's usage value would be re-calculated. This helps to ensure that the kiosk continues to maintain a high coverage rate of hospital procedures.

Similarly, 4-8 also shows that it is faster to automatically iterate on a kiosk using the model, rather than updating it manually. Iterations happen approximately 1-2 times when a new planogram is created (or 2-8 times per year). In an iteration, hospital staff members will provide feedback to the AITA™ team on changes they want to see in the planogram. Depending on the types of adjustments that need to be made, it could take anywhere from 1-2 hours to update the design under the manual process. Faster updates would involve adding, removing, or substituting particular products. A more substantial update would involve changing a more fundamental parameter such as the number of large or medium products in the kiosk.

Chapter 5

Shelf Planogram Model and Results

5.1 Current Shelf Planogram Creation Process

Shelves hold a hospital's entire surgical suture inventory and organize the products to make it easy for hospital staff to find items. In the current process to create shelf planograms, the AITA™ team begins by using preference card and purchasing data to identify all the unique product SKUs at the hospital. They then look up the key attributes of each product: the product group, the suture size, the needle length, the needle point, and information about the closure (e.g. ties). The team then manually sorts the products based on these attributes.

The AITA™ team starts by putting products of the same group together. This is complicated by the fact that a shelf is actually made up of a handful of racks (usually 3-5 racks). These racks are placed next to each other in a storage room, and are collectively referred to as a shelf. To reduce confusion, the AITA™ team needs to ensure (to the extent possible) that products from the same product groups are placed together on the same rack. If, for example, a member of hospital staff needed to find a Prolene suture, it would be confusing and time consuming if they had to check multiple racks. To further reduce confusion, once all products in a product group are placed together, the individual products should then be sorted by their remaining attributes: suture size, needle length, needle point, and closure.

The current process requires hundreds of manually-made decisions and is thus inefficient and error-prone. From an efficiency perspective, it takes 1-2 days to create a new planogram because each item must be looked up and placed in the shelf by hand. Furthermore, the AITA™ team must frequently change which product groups are on each rack to minimize the occurrences that a product group is split between racks. Similarly, the process to iterate on planogram designs with hospitals is also time consuming and can take approximately 2-4 hours for each iteration. This is because new items must be researched to determine their position among the sorted products, and because any change in a product group's size might necessitate shuffling groups between racks.

Planograms created under the current process are also not guaranteed to be completely accurate. Given that each product is categorized and sorted by hand, it is possible that mistakes could happen in every planogram implementation. These mistakes usually involve placing a product outside of its correctly sorted order, making it harder to find, and weakening hospitals' trust in the accuracy of the system.

5.2 Shelf Model Objectives

The shelf model improves the efficiency and accuracy of shelf planograms by automatically and quickly creating layouts where the products are always correctly sorted. As shown in 5-1, the model does this in two steps 1) select which product groups should be placed on each rack and 2) sort individual items within each product group.

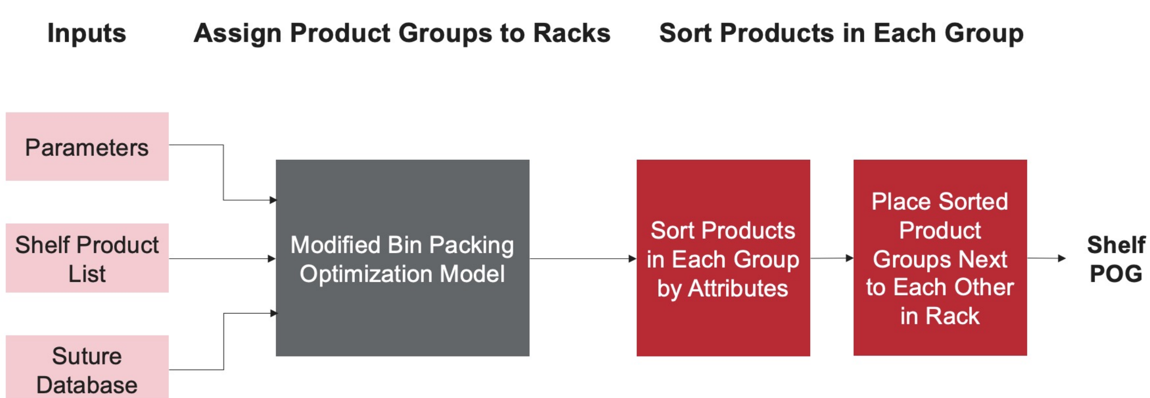

Figure 5-1: Process flow to automatically create shelf planograms

Step 1 (selecting product groups for each rack) is intended to ensure that product groups are not split between multiple racks. The model achieves this with a bin-packing-inspired optimization formulation that minimizes the number of product groups placed on multiple racks. Next, step 2 involves sorting the items within each product group. The shelf model does this by pulling attribute information for each product (suture size, needle length, needle point, and closure) and sorting the products based on these attributes. The final output is a visual of the shelf's planogram where each item is placed in a logical way so that they are easy for hospital staff members to find.

5.3 Shelf Model Inputs

The user must input a set of parameters and files to run the shelf planogram creation model. These inputs will determine which products are placed in each rack, how the products should be physically arranged, and how the final output should be visualized.

5.3.1 Input Parameters

Table 5.1: Input Parameters for the Shelf Model

Parameter	Description
Number of racks	Number of racks that make up the shelf; shelves usually have 3-5 racks that are placed next to each other
Rack rows/columns	Number of rows / columns in each rack
Empty bin placeholders	Specify which product groups should have empty bins at the end of the group, and how many empty bins there should be; E.g. "Prolene – 2" will reserve 2 bins where Prolene products can be placed later
Multiple bins	Specify which products should have multiple bins, and how many bins the products should have

Continued on next page

Table 5.1 – continued from previous page

Parameter	Description
Adjacent Product Groups	Which product groups should be placed next to each other on a rack; e.g. Plain Gut sutures are usually followed by Fast Absorbing Gut sutures and Chromic Gut sutures
Merged Product Groups	Which product groups should be considered as one group; E.g. if the PDS Plus group and the PDS group are merged, the products from each group will be mixed together on the shelf (i.e. the order could be plus, plus, non-plus, plus)
Sorting order	Order in which product attributes should be sorted; Options are: Thickness, Closures, Needle Length, and Needle Point; E.g. If needle point comes before needle length, products will be sorted by point and then by length
Suture Size Order	Order in which suture size should be sorted (e.g. 5-0 sutures come before 6-0 sutures, etc.)
Closure Order	Order in which closures should be sorted (e.g. sutures with ties, then sutures with reels, then looped sutures, then standard sutures)
Needle Length Order	Sort needle length in increasing or decreasing order
Needle Point Order	Order in which needle points should be sorted (e.g. taper point sutures then taper cut sutures)
Product Group Colors	Color that each product group should have in the final visual (e.g. Prolene can be assigned a light green color)

5.3.2 Input Files

Table 5.2: Input Files for the Shelf Model

Input File	Description	Usage
Shelf product list	SKU numbers of all the products that should be added to the shelf; can be based on current inventory, preference cards, or purchasing	See which products need to be placed and organized on the shelf
AITA™ suture database	Lists key attributes for each product SKU; namely product group, suture size, needle length, needle point, and closures	Use attribute information to sort products in the shelf; group determines which racks products are on, and other attributes determine the order of SKUs within groups

5.4 Select Product Groups for Each Shelf Rack

5.4.1 Illustrative Examples

After specifying input parameters and files, the first major step in the shelf planogram creation process is selecting which product groups should go on each rack in the shelf. Shelves usually comprise around 3-5 separate racks that are placed next to each other in a hospital's stock room. To reduce confusion, the AITA™ team aims to minimize the number of product groups that are split across multiple racks. In the ideal case, when no product groups are split, hospital staff would save time because they would only need to search one rack to find a product from a particular group. To bring about this ideal case as much as possible, an optimization model is used to minimize the number of product groups split between racks.

The following example shows why the split product minimization problem is non-trivial. 5-2 shows a simplified shelf with just two, small racks. Each rack has 25 slots arranged into 5 rows and 5 columns. There are 6 product groups (Groups A - F), each with a different number of items. The goal is to place the product groups on the shelf while trying to minimize the number of products that are split across both racks.

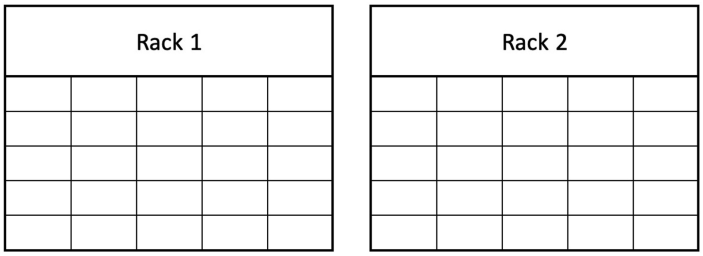

Product Group	Number of SKUs
Group A	11
Group B	4
Group C	12
Group D	6
Group E	9
Group F	8

Figure 5-2: Illustrative product group placement example; 6 product groups with 50 total products that need to be placed in two shelf racks with 25 slots each

Faced with the scenario in 5-2, it is not immediately clear which combination of groups should be placed on each rack. For example, as shown in 5-3, if product groups are placed in alphabetical order, Group C will be split between the two racks, with 10 items on rack 1 and 2 items on rack 2.

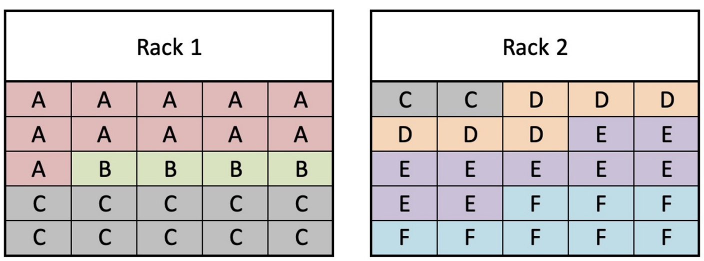

Product Group	Number of SKUs
Group A	11
Group B	4
Group C	12
Group D	6
Group E	9
Group F	8

Figure 5-3: Un-optimized solution to illustrative example; product groups are placed in alphabetical order, resulting in Group C being split across Racks 1 and 2

One possible solution to the problem is shown in 5-4. In this planogram, groups A, D, and F are placed entirely on rack 1, and groups B, C, and E are placed entirely on rack 2. While it is possible to find this solution manually in this example, real shelf planograms are rarely this simple. In practice, there are dozens of product groups spanning hundreds of possible bins, making it highly time consuming to sort product groups into racks by hand. Instead, the shelf model uses optimization to quickly and correctly identify which product groups should be placed on each rack.

Rack 1				
A	A	A	A	A
A	A	A	A	A
A	D	D	D	D
D	D	F	F	F
F	F	F	F	F

Rack 2				
B	B	B	B	C
C	C	C	C	C
C	C	C	C	C
C	E	E	E	E
E	E	E	E	E

Product Group	Number of SKUs
Group A	11
Group B	4
Group C	12
Group D	6
Group E	9
Group F	8

Figure 5-4: Optimized solution to illustrative example, each product group is on just one rack

Even when using an optimization model, there are sometimes situations where product groups need to be split between racks. For instance, 5-5 shows a new example with 2, 25-bin racks and 3 product groups (Group X, Group Y, and Group Z). Given the sizes of the 3 groups, there is no feasible way to place them on the shelf without splitting up one of them. In this case, the model should split the smallest group (Group Y) between the two racks, because this would hopefully cause less disruption and confusion than splitting up a larger group such as X or Z.

Rack 1				
X	X	X	X	X
X	X	X	X	X
X	X	X	X	X
X	X	X	Y	Y
Y	Y	Y	Y	Y

Rack 2				
Y	Y	Y	Y	Y
Y	Z	Z	Z	Z
Z	Z	Z	Z	Z
Z	Z	Z	Z	Z
Z	Z	Z	Z	Z

Product Group	Number of SKUs
Group X	18
Group Y	13
Group Z	19

Figure 5-5: Illustrative example where a product group needs to be split between racks; the smallest product group (Group Y) is chosen to be split

5.4.2 Optimization to Choose Product Groups for Each Rack

The optimization model used to keep product groups together on single racks is described in detail below. Its formulation is based on the bin-packing problem, because the fundamental idea is to fit as many differently sized groups as possible in a fixed amount of space. The model enables fast and accurate allocation of product groups to racks, even when the number of product groups, racks, and rack sizes grows.

Sets

G Set of all product groups, indexed by g

R Set of all shelf racks, indexed by r

Parameters

N Rack capacity (maximum number of products that fit in each rack)

c_g Count of individual products in each product group g (number of product SKUs in each group)

Decision Variable

$$x_{g,r} = \begin{cases} 1, & \text{if product group g is in rack r} \\ 0, & \text{if product group g is not in rack r} \end{cases}$$

Dependent Variable

$$e_r = \begin{cases} 1, & \text{if rack r is empty (if rack r has no products stocked in it)} \\ 0, & \text{if rack r is not empty (if rack r has 1 or more products stocked in it)} \end{cases}$$

Mathematical Formulation

Objective

$$\text{maximize} \sum_r \sum_g (x_{g,r} * c_g) + \sum_r e_r \tag{5.1}$$

Subject to

$$\sum_r x_{g,r} \leq 1 \quad \forall \text{ product group g} \tag{5.2}$$

$$e_r = \begin{cases} 0 & \text{if } \sum_g x_{g,r} > 0 \\ 1 & \text{otherwise} \end{cases} \tag{5.3}$$

$$\sum_g (x_{g,r} * c_g) \leq N \quad \forall \text{ rack r} \tag{5.4}$$

The objective function 5.1 is split into two terms. The first term ($\sum_r \sum_g (x_{g,r} * c_g)$), combined with constraint 5.2 aims to maximize the number of products that are placed in the shelf, while ensuring that no product group is split between racks. In particular, 5.2 guarantees that if $x_{g,r} = 1$, then all products in product group g will be placed solely on rack r. The objective function logic was structured to maximize the number of products shelved (conditional on no product group being split) rather than the number of product groups shelved. This helps the model handle situations like the example shown in 5-5, where product groups must be split between racks. With this objective function, the split product groups would be the smallest categories with the fewest number of SKUs.

The second term of the objective function ($\sum_r e_r$) aims to maximize the number of empty racks in the shelf. An empty rack is defined in constraint 5.3 as a rack that does not contain any product groups. Mathematically, e_r only equals 1 (signaling that rack r is empty) if $x_{g,r}$ equals 0 for every possible product group g. This empty rack logic is important from a space-planning perspective. If the hospital has 4 racks but the model can fit all the products into 3, it will choose to place the products on just 3 racks, because one rack would be left empty (increasing the objective function). This would help to keep the shelves as contained as possible, promoting a more efficient use of space.

The final constraint 5.4 simply ensures that the number of products placed on each rack does not exceed the rack's capacity. Taken together, the objective function and constraints comprise an optimization model that will ensure that as many products from the same group are placed on the same shelf as possible. In a situation where it is not feasible to keep all product groups together (e.g. 5-5), the model will identify which product groups should not be placed on any rack (i.e. will identify all the product groups g where $\sum_r x_{g,r} = 0$). The model will then split these product groups by placing products from the groups into empty slots in each rack.

5.5 Sort Shelf Products within Each Product Group

After product groups are assigned to racks, they must be sorted into a logical order. This will ensure that members of hospital staff can easily find products when they need to pick items for procedures or restock the shelves. Different hospitals might have different ways that they prefer to sort items in their stock rooms. So, the model was created to be adaptable to different sorting styles and rules.

The model enables sorting by 4 different attributes: suture size, closure, needle length, and needle point. A hospital can choose to use any combination of these attributes to sort the individual SKUs within each product group. For example, they could choose to sort a product group by just decreasing needle length alone. Or, they could choose to sort the products by suture size and then to further sort sutures of the same size by needle point.

Furthermore, the model allows for flexibility in the sorting order of each of the 4 attributes. For needle length, the user can indicate if the products should be placed in increasing or decreasing length order. For suture size, closures, and needle point, the user can specify the desired order of the attribute (e.g. by indicating that a tie should be placed before a reel, a reel should be placed before a loop, and a loop should be placed before a standard suture).

For example, suppose a hospital has the following 17 products (5-6) that belong to a certain product group (e.g. Dafilon, Monoplus, Prolene, etc.). The hospital wants to sort its products using all 4 possible attributes in the following order: suture size, closures, needle length, and needle point. The hospital wants suture sizes to be ordered from thickest to thinnest (e.g. 3, 2, 1, 0, 1-0, 2-0, 3-0); closures to be ordered as: ties, loops, reels, and standard; needle length to be decreasing; and needle type to be ordered as: taperpoint, tapercut, straight cutting, and reverse cutting.

SKU #	Suture Size	Closure	Needle Length	Needle Type
Prod 1	0	Tie	48	Tapercut
Prod 2	2-0	Loop	30	Straight Cutting
Prod 3	2	0	14	Taperpoint
Prod 4	2-0	0	42	Reverse Cutting
Prod 5	2	0	28	Taperpoint
Prod 6	1-0	0	26	Tapercut
Prod 7	1	0	36	Reverse Cutting
Prod 8	2-0	0	17	Taperpoint
Prod 9	2	0	20	Tapercut
Prod 10	2-0	0	42	Taperpoint
Prod 11	0	Reel	-	Taperpoint
Prod 12	1	Loop	18	Straight Cutting
Prod 13	1-0	0	26	Straight Cutting
Prod 14	2-0	Loop	24	Tapercut
Prod 15	1	0	24	Taperpoint
Prod 16	1-0	0	26	Taperpoint
Prod 17	1-0	0	26	Reverse Cutting

Figure 5-6: Unsorted products in an illustrative product group

5-7 shows how the model would sort the 17 SKUs in the product group according to the hospital's preferred rules. Suture size is the first attribute used when ordering, so the sorted products decrease from a thickness of 2 to a thickness of 2-0.

Closures are the next attribute used in the sorting process. This means that any suture with a tie, loop, or reel will come before any standard suture of the same size. For example, out of the three products with a suture size of 1, Product 12 comes first, because it has a loop, while the other two products are standard. Meanwhile, for the two sutures with a size of 0, Product 1 comes before Product 11 because the hospital wants ties to come before reels.

Sutures are next sorted by decreasing needle length. For example, there are 3 products (Products 5, 9, and 3) with both a suture size of 2 and a standard closure (i.e. no ties, loops, or reels). So, they are sorted by their needle length and are placed in decreasing order from 28mm to 14mm. It should be noted that, under this example's rules, the suture size and closure attributes will always take precedence over the needle length. For example, Product 12 is the first suture out of the three SKUs with a size of 1, even though it has the shortest needle length of the three (18mm).

Needle type is the final attribute used to sort the SKUs in this product group. In this case, the specified order to follow was taperpoint, tapercut, straight cutting, and reverse cutting. This can be seen clearly in the sorted order of Products 16, 6 ,13, and 17. All 4 products have the same suture size, closure style (i.e. no ties) and needle length. So, they are sorted on the remaining attribute, needle type, and are placed in the same order specified by the hospital in the original set of rules.

SKU #	Suture Size	Closure	Needle Length	Needle Type
Prod 5	2	0	28	Taperpoint
Prod 9	2	0	20	Tapercut
Prod 3	2	0	14	Taperpoint
Prod 12	1	Loop	18	Straight Cutting
Prod 7	1	0	36	Reverse Cutting
Prod 15	1	0	24	Taperpoint
Prod 1	0	Tie	48	Tapercut
Prod 11	0	Reel	-	Taperpoint
Prod 16	1-0	0	26	Taperpoint
Prod 6	1-0	0	26	Tapercut
Prod 13	1-0	0	26	Straight Cutting
Prod 17	1-0	0	26	Reverse Cutting
Prod 2	2-0	Loop	30	Straight Cutting
Prod 14	2-0	Loop	24	Tapercut
Prod 10	2-0	0	42	Taperpoint
Prod 4	2-0	0	42	Reverse Cutting
Prod 8	2-0	0	17	Taperpoint

Figure 5-7: Products sorted by suture size, closures, needle length, then needle type

The model orders products within each group using a technique called hierarchical, or multiple-key sorting. Under this system, data can be sorted by multiple criteria in a particular order. The data are first sorted by one criteria and ties are broken by the second criteria. If there are any ties on both the first and second criteria, they are broken using the third criteria (and so on). This method allows for a fast, automatic way to sort products by any combination and order of attributes that a hospital sees fit. The result is an organized and intuitive shelf, where hospital staff can easily retrieve and replace individual products.

5.6 Shelf Model Outputs

After the model selects which product groups should be placed on each rack and sorts the individual products within each group, it outputs a visual of the shelf planogram. 5-8 shows an illustrative view of a model-generated output. In this example, the shelf had two racks, each with 48 slots split into 8 columns and 6 rows. In reality, shelves are usually larger than 96 total slots, because they must fit all the items in a hospital's surgical inventory.

Rack 1	Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	Col 7	Col 8
Row 1	Monocryl 5	Monocryl 9	Monocryl 10	Monocryl 2	Monocryl 4	Monocryl 7	Monocryl 3	Monocryl 8
Row 2	Monocryl 11	Monocryl 1	Monocryl 6	Fast Absorb 3	Fast Absorb 1	Fast Absorb 2	Plain Gut 1	Plain Gut 3
Row 3	Plain Gut 2	Prolene 14	Prolene 17	Prolene 10	Prolene 20	Prolene 8	Prolene 16	Prolene 4
Row 4	Prolene 11	Prolene 18	Prolene 7	Prolene 1	Prolene 13	Prolene 3	Prolene 9	Prolene 19
Row 5	Prolene 5	Prolene 2	Prolene 21	Prolene 12	Prolene 6	Prolene 15	PDS 4	PDS 7
Row 6	PDS 3	PDS 6	PDS 5	PDS 1	PDS 8	PDS 2	Empty Slot 1	Empty Slot 2

Rack 2	Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	Col 7	Col 8
Row 1	Vicryl 5	Vicryl 1	Vicryl 7	Vicryl 2	Vicryl 6	Vicryl 8	Vicryl 9	Vicryl 3
Row 2	Vicryl 4	Stratafix 2	Stratafix 8	Stratafix 1	Stratafix 7	Stratafix 9	Stratafix 3	Stratafix 4
Row 3	Stratafix 6	Stratafix 5	Ethilon 13	Ethilon 8	Ethilon 12	Ethilon 17	Ethilon 10	Ethilon 3
Row 4	Ethilon 5	Ethilon 14	Ethilon 1	Ethilon 15	Ethilon 9	Ethilon 2	Ethilon 18	Ethilon 11
Row 5	Ethilon 16	Ethilon 4	Ethilon 6	Ethilon 7	Ethibond 4	Ethibond 2	Ethibond 6	Ethibond 3
Row 6	Ethibond 7	Ethibond 1	Ethibond 5	Empty Slot 3	Empty Slot 4	Empty Slot 5	Empty Slot 6	Empty Slot 7

Legend
Monocryl
Stratafix
Ethilon
Coated Vicryl
Fast Absorbing Gut
Plain Gut
Ethibond
Prolene
PDS

Figure 5-8: Illustrative output of the shelf planogram model

The final shelf output has a few notable features. First, each product group is color-coded, and the legend is included in the shelf output. Color-coding helps the AITA™ team ensure that product groups are split up and organized as expected. In this example, no product group needed to be divided between the two racks. The second notable feature is the order of the products. The output comes with products sorted within their groups according to the hospital's rules. For instance, Monocryl 5 coming before Monocryl 9 in 5-8 could be because Monocryl 5 is a thicker suture. The final notable feature is the 7 total empty slots in the shelf. Since shelves are meant to hold all the surgical products in a hospital, shelves often have extra capacity in case new products need to be added.

Aside from the primary product view shown in 5-8, the shelf model's output also includes additional views that are intended to help the AITA™ team and hospital staff understand and iterate on the result. The 3 other main views are the description, group, and printing layouts. Much like in the kiosk model, the description view shows the full product name of every item (e.g. 3-0 Coated Vicryl Plus Undyed 1x27" (70cm) PS-2 Multipass). The group layout shows the product group of every item so users can confirm that the color-coding and sorting matches their expectations. Finally, the printing layout includes both the SKU number and the full product name. It is formatted with specific dimensions and borders so it can be turned into a large-scale print-out that is placed in the hospital and used in the physical shelf implementation process.

5.7 Shelf Model Results

Much like the kiosk model, the shelf model improves both the accuracy and efficiency of the planogram creation process. From an accuracy perspective, the model's planogram is free of manual errors. When placing product groups in the shelf, if there is a feasible way to fit each product group on a single rack, the model's optimization algorithm will find it. From a sorting perspective, the model will correctly organize products according to the hospital's pre-defined rules, without making mistakes.

	Frequency	Steps	Model-based process	Manual process	Average Time Saved per Year
Creating a new shelf planogram	1 time per year	Aggregate / clean data	1 hour	1.5 days (12 hours)	**9.7 hours** = (12 – 2.28) hours * 1 per year
		Adjust parameters	30 minutes		
		Run model	2 minutes		
		Check output	45 minutes		
Iterating on a shelf planogram	4-8 times per year	Update data / parameters	5 minutes	2-4 hours	**14.3 hours** = (3 – 0.62) hours * 6 per year
		Run model	2 minutes		
		Check output	30 minutes		

Figure 5-9: Time taken to create and update shelf planograms when using the model vs. a manual process

Meanwhile, the model's efficiency improvements are summarized in 5-9. In total, using the model as opposed to the manual process results in time savings of 24 hours per hospital per year. 9.7 of these hours saved come from the reduced amount of time to create a new shelf planogram, and 14.3 of these hours come from the time saved iterating on a shelf planogram.

Under the current, manual process, it takes 1.5 days (12) hours to create a shelf planogram. This process takes longer than kiosk planogram creation, because there can be hundreds more products in a shelf than in a kiosk. These products must then be sorted one-by-one by hand and the output must be carefully checked to ensure it is error-free. In contrast, the process to create a new shelf planogram using the model takes only around 2 hours. This is an hour shorter than the time required to build a kiosk, because it takes less time to aggregate and clean the input data. The only two necessary input files are a list of all the unique surgical products in the hospital (which can be obtained from preference cards or purchasing data), and the suture database that is continuously maintained and updated by the AITA™ team.

It is also faster to iterate on a planogram's design using the model. Under the current, manual process, it can take 2-4 hours to update a planogram. A small change could constitute adding or removing products, but even these actions can be time consuming, because they necessitate shifting the order of products within the group. Updating a shelf can take even more time if product groups grow or shrink, because it could require moving the groups to minimize splitting between racks. However, the most time-consuming manual update would likely be changing the underlying sorting order (e.g. needle length before needle type), because it would necessitate moving every product on the shelf. Conversely, all of these changes can be solved quickly and easily with the model. Changing products or even the sorting order just requires quick tweaks of the input parameters. Overall, this means a shelf iteration with the model only takes around 30 minutes (including the time needed to check the output), shaving hours off the manual process.

Similar to the kiosk model, the time savings from automatically creating and updating planograms using the model are compounded when the AITA™ team refreshes the shelves. The team usually creates a new planogram once a year by re-assessing the hospital's preference cards or purchasing data to get an updated list of all the unique surgical items in use. This process happens less often than the kiosk refresh process because the kiosk holds just a subset of items and is thus more sensitive to changes in the frequency of procedures. However, as items are added and removed from hospitals, the shelf still needs to be updated/iterated on, because it is intended to accurately reflect the complete surgical product inventory. These smaller iterations happen roughly 1 to 2 times a quarter, or 4-8 times a year.

Chapter 6

Hub Planogram Model and Results

The final model enables the automatic creation of hub planograms. While kiosks are designed to cover as many procedures as possible, and shelves are used to store and organize a hospital's full surgical inventory, hubs are leveraged to manage add-on requests. Add-on requests occur when surgeons require additional products during a procedure. This could happen for several reasons, including a change to the surgical plan, an extended surgery duration, or the need for a replacement product. Given the varied nature of add-on requests, hubs must be able to provide convenient and quick access to a wide range of products. So, they are usually stationed close to operating rooms (ORs) and contain products that are commonly used in that group of rooms.

6.1 Current Hub Planogram Creation Process

In the current hub planogram creation process, the AITA™ team manually selects and sorts the products to be placed in the device. They start by mapping out how many hubs should be placed in a hospital, and which ORs each hub will service. For example, in the illustrative hospital set-up shown in 6-3, there are eight total operating rooms, with three concentrated on one end of the hall, two in the middle of the hall, and three concentrated on the other end of the hall. So, it might make sense to assign three hubs to the hospital: one to provide products to ORs A, B, and C; one to provide products to ORs D and H, and one to provide products to ORs E, F, and G.

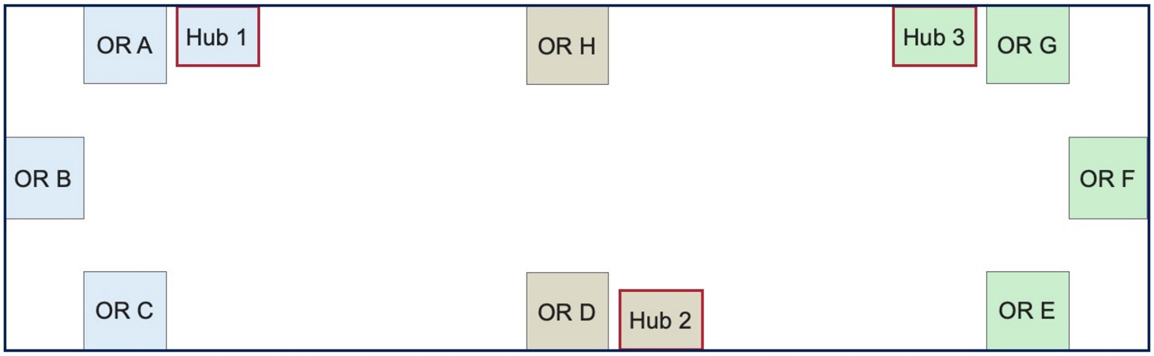

Figure 6-1: Illustrative layout of a hospital's operating rooms (ORs) and where hubs might be placed

Once it is determined which operating rooms will be served by each hub, the AITA™ team next chooses which products to place in each hub. The selection process is usually based on specialty because many hospitals coordinate their layout to correspond to different types of surgical procedures. In this case, ORs A, B, and C might correspond to Orthopedic and General surgeries; ORs D and H might correspond to Gynecology and Urology, and ORs E, F, and G might correspond to Plastic and Vascular surgery. The team chooses products for hubs based on these specialties. For example, they would place the most popular Orthopedic and General products into Hub 1, where popularity is either based on purchase / usage data or based on interviews with members of the hospital staff who work in those ORs.

Once products are selected, they must be sorted and placed in the hub. Much like the shelf, hubs are sorted based on their product attributes: type, suture size, closures, needle length, and needle point. They are organized as such to ensure that members of the hospital staff can easily pick out and replenish products. The only notable difference between the shelf sorting style and the hub sorting style is that hubs, like kiosks, have size restrictions. Only a certain number of columns can hold large or medium products, while the rest of the columns only hold small products. So, the AITA™ team must choose products that meet these size requirements, and must organize them such that products are placed into the correctly sized columns, while also following product attribute sorting guidelines.

Much like the kiosk and shelf processes, the hub planogram design process is not as efficient or accurate as it could be. On the efficiency side, hub planograms take the most time to design and iterate on, because they combine picking requirements from the kiosk and sorting requirements from the shelf. In particular, it takes approximately 2 days to design a new planogram, because the most popular products for each specialty must be determined, and because they must be sorted by hand. Then, it could take 1-3 hours to iterate on the planogram's design, where a shorter iteration would involve swapping out products, and a longer iteration would involve changing the device's underlying sorting order.

Aside from the time-consuming process required to design hub planograms, the final results are also not usually as accurate as the AITA™ team and hospitals would like. To start, similar to the shelf, there is a potential for missorted products. If a product is placed in the wrong bin, it would be harder for hospital staff members to find it when they are rushing to complete an add-on picking during a procedure. There are also scenarios where hubs do not have the right assortment of products in the first place. This situation occurs when surgeries for one specialty are assigned to ORs that are usually designated for another specialty. In fact, such cross-specialty OR assignment happens frequently, as ad-hoc and unpredictable procedures often require scheduling flexibility. If, for example, general surgeons are assigned to ORs for plastic and vascular procedures, the hubs closest to them might not have the right selection of add-on products.

6.2 Hub Model Objectives

The automated hub planogram creation model solves the manual approach's accuracy and efficiency deficits. Rather than choosing products based on specialty, it selects the products that were used the most often historically in its corresponding ORs. Then, it automatically sorts these products according to a set of rules that the hospital can set and change with ease.

As shown in 6-2, the model can be broken down into 2 main components: 1) product selection, and 2) product sorting. During the product selection process, the

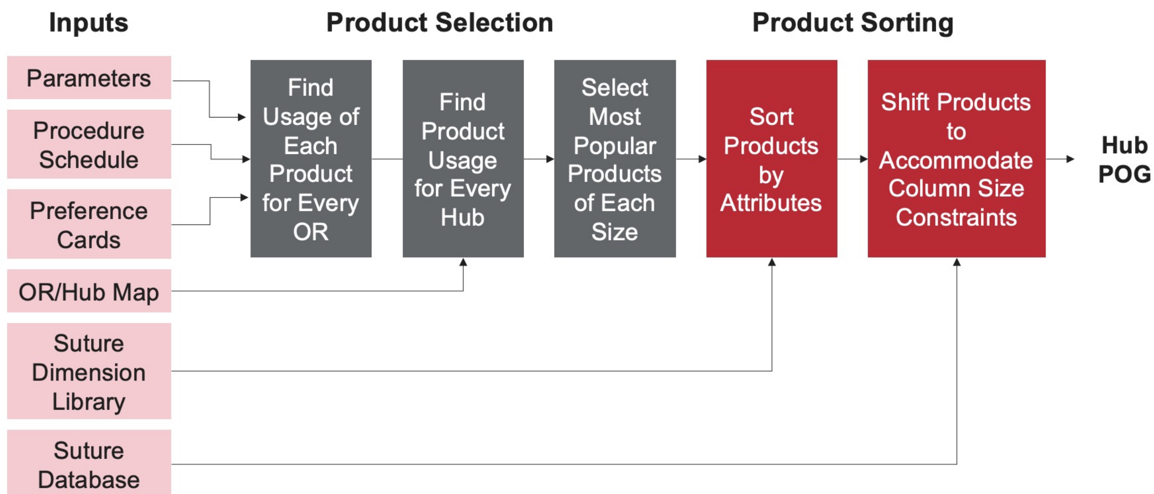

Figure 6-2: Process flow to automatically create hub planograms

model automatically parses through procedure schedule and preference card data to determine the usage of each product in every OR. It then aggregates this information across all the ORs corresponding to a particular hub, and selects the L most popular large and medium products and the S most popular small products, depending on the size constraints of the device.

The model uses preference card data to select products because preference cards are often the only reasonable proxy available for add-on pickings. In an ideal world, the model could stock the hub using information about the exact products that were historically picked as add-ons for procedures, but this information is often not tracked by hospitals. So, the model uses preference cards to inform product selection under the assumption that add-ons will often come from the pool of products that were already commonly being used in the hub's ORs. For example, in a hub serves plastic surgery ORs, an add-on for a plastic surgery procedure would likely be covered by that procedure's preference card or by other plastic surgery preference cards that are served by the hub.

After product selection, the model next sorts products with the same attribute-based process as in the shelf model. The only difference between the hub's and shelf's sorting process is that the hub's products are often shifted to accommodate the size constraints of each column. After these two steps are completed, the final output is a visual depiction of the hub's planogram. This planogram will contain products that are likely to be needed during an add-on request and will be sorted such that they are easy to find.

6.3 Hub Model Inputs

The hub model begins with the user uploading a set of input parameters and files. These inputs will determine which products are selected for each hub, how the products are sorted within each hub, and how the final output should be configured.

6.3.1 Input Parameters

Table 6.1: Input Parameters for the Hub Model

Parameter	**Description**
Number of Rows / Columns	Number of rows / columns in each hub; used for product selection and the final visualization
Location of Large / Medium Columns	Which columns can fit large and medium products; can either specify that columns should hold just large or just medium products, or that large and medium products can be placed in the same columns together
Manually Add Products or Change Quantities	Add a new product to a hub that is not already present; or change the quantity of a product already in a hub (e.g. if a product is popular, specify that it should have 2 bins)
Manually Remove Products	Specify if a certain product should be removed from a particular hub (e.g. the plastic surgery team is switching to a new product, so remove Product A from Hub 3)

Continued on next page

Table 6.1 – continued from previous page

Parameter	Description
Adjacent Product Groups	Which product groups should be placed next to each other in the hubs; e.g. Plain Gut sutures are usually followed by Fast Absorbing Gut sutures and Chromic Gut sutures
Merged Product Groups	Which product groups should be considered one group in the hubs; E.g. if the PDS Plus group and the PDS group are merged, the products from each group will be mixed together on the shelf (i.e. the order could be plus, plus, non-plus, plus)
Sorting order	Order in which product attributes should be sorted; Options are: Thickness, Closures, Needle Length, and Needle Point; E.g. If needle point comes before needle length, products will be sorted by point and then by length
Suture Size / USP Order	Order in which suture size / USP should be sorted (e.g. 5-0 sutures come before 6-0 sutures, etc.)
Closure Order	Order in which closures should be sorted (e.g. sutures with ties, then sutures with reals, then looped sutures, then standard sutures)
Needle Length Order	Specify if needle length should be sorted in increasing or decreasing order
Needle Point Order	Order in which needle points should be sorted (e.g. taper point sutures then taper cut sutures)
Product Group Colors	Color that each product group should have in the final visual (e.g. Prolene can be assigned a light green color)

6.3.2 Input Files

Table 6.2: Input Files for the Hub Model

Input File	Description	Usage
OR / Hub map	Indicates which ORs should be serviced by each Hub (e.g. Hub 1 contains products for ORs A, B, and C)	Shows which ORs need to be considered when choosing the most popular products for a hub
Procedure schedule	Lists every procedure in a hospital over a certain time frame; includes the OR where the procedure took place and the preference card corresponding to the procedure	See how often each preference card was used in each OR; used with the OR/Hub Map and the Preference Cards file to select the most popular products for each hub
Preference cards	Every preference card and its associated products	Determine which products would be used the most often in the ORs that correspond to a particular hub
AITA™ suture dimension library	Has the sizes (small, medium, or large) of each suture	Assess which products can go in the small, medium, and large hub columns
AITA™ suture database	Lists key attributes for each product SKU; product group, suture size, needle length, needle point, and closures	Use attribute information to sort products in the hub

6.4 Hub Product Selection

After the user specifies the necessary parameters and uploads the required input files, the hub planogram creation model begins the process of product selection. First, the model uses the OR/Hub map, the procedure schedule, and the preference card files to calculate the historical usage of each product in a hub's corresponding operating rooms. Next, the model selects the L most popular large and medium products and the S most popular small products to place in the hub.

6.4.1 Illustrative Example: Calculating Product Usage

The process used to calculate each product's usage can be seen in the illustrative example laid out by 6-3 and 6-4. As shown in 6-3, the model begins by using the OR/Hub map and the hospital's procedure schedule to create a preference card summary for every hub. The OR/Hub map shows the ORs serviced by each hub (e.g. Hub 1 provides add-on products for ORs A, B, C, and D). The hospital's procedure schedule shows which preference cards were used at each OR over a given time period (e.g. OR A used 3 preference cards: A, D, and B). Combining this information leads to a preference card summary for each hub. For example, as shown in the procedure schedule, the 4 ORs serviced by Hub 1 had 10 total procedures, each with a corresponding preference card. The preference card summary shows that Card A was used 4 times, Card B was used 2 times, Card C was used 3 times, and Card D was used once.

1) OR / Hub Map

Hub	OR
1	A
1	B
1	C
1	D
2	E
2	F
2	G
2	H
3	I
3	J

&

1) Procedure Schedule

OR	Preference Card ID
A	Pref Card A
A	Pref Card D
A	Pref Card B
B	Pref Card A
B	Pref Card B
C	Pref Card C
C	Pref Card A
C	Pref Card C
C	Pref Card C
D	Pref Card A

➡

2) Hub 1 Pref. Card Summary

Hub 1 Pref. Cards	Number of Occurences
Pref Card A	4
Pref Card B	2
Pref Card C	3
Pref Card D	1

Figure 6-3: Illustrative creation of a preference card summary for a given hub (Hub 1)

6-4 shows that, after the preference card summary for Hub 1 is created, it can be combined with the preference card details file to calculate the usage for every product. The preference card details file shows which products appear in every preference card. For example, Preference Card A contains 3 products: Products W, Y, and Z. So, to calculate the historical usage of each product, the hub model checks which preference cards contain the product, and sums the number of times those preference cards were used over the time period. For instance, Product Z was used in 3 cards: A, C, and D. Preference Card A was used 4 times, Card C 3 times, and Card D once. So, Product Z has a historical usage of 8 occurrences $(4 + 3 + 1)$.

2) Hub 1 Pref. Card Summary

Hub 1 Pref. Cards	Number of Occurences
Pref Card A	4
Pref Card B	2
Pref Card C	3
Pref Card D	1

&

3) Preference Cards

Pref. Card.	Product
Pref Card A	Prod W
Pref Card A	Prod Y
Pref Card A	Prod Z
Pref Card B	Prod W
Pref Card C	Prod W
Pref Card C	Prod X
Pref Card C	Prod Y
Pref Card C	Prod Z
Pref Card D	Prod W
Pref Card D	Prod Z

➡

4) Usage Calculations

Product	Historical Usage
Prod W	**10**
	= A + B + C + D
	= 4 + 2 + 3 + 1
Prod X	**3**
	= C = 3
Prod Y	**7**
	= A + C = 4 + 3
Prod Z	**8**
	= A + C + D
	= 4 + 3 + 1

Figure 6-4: Illustrative calculation of historical usage for each product using a preference card summary and preference card information

6.4.2 Mathematical Formulation to Calculate Product Usage

Variables

H Set of all hubs in a hospital, indexed by h

O_h Set of operating rooms covered by hub h, indexed by o

P Set of all the products in a hospital, indexed by p

C Set of all the preference cards used over a certain time period, indexed by c

$N_{c,h}$ Total number of times each preference card was used in one of hub h's ORs

$U_{p,h}$ Historical Usage of product p in hub h

$$y_{p,c} = \begin{cases} 1, & \text{if product p is in preference card c} \\ 0, & \text{if product p is not in preference card c} \end{cases}$$

$$x_{c,o} = \begin{cases} 1, & \text{if preference card c was used in OR o} \\ 0, & \text{if preference card c was not used in OR o} \end{cases}$$

Mathematical Formulation

For every hub h:

$$N_{c,h} = \sum_{c \in C} \sum_{o \in O_h} x_{c,o} \tag{6.1}$$

$$U_{p,h} = \sum_{c \in C} y_{p,c} * N_{c,h} \quad \forall \text{ product p} \tag{6.2}$$

Equation 6.1 shows how the model mathematically calculates the number of times each preference card c was used in an OR covered by hub h. The model iterates over all the procedures that occurred in a given time frame and counts how many times each preference card c was assigned to one of the ORs o covered by hub h. Next, equation 6.2 uses this value ($N_{c,h}$) to calculate the historical usage of each product. The model iterates over all the preference cards in the hospital. Any time a preference card c contains product p, the model adds $N_{c,h}$ (the number of times that preference card c was used in an OR covered by hub h in the procedure schedule) to product p's usage.

6.4.3 Choosing Products by Size

After the model calculates each product's historical usage, it uses these values to choose which products should be placed in the hub. To try to maximize the number of add-on requests covered by the hub, the model places items with the highest historical usage levels into the device. However, if hubs can hold N total products, the model cannot simply choose the most popular N products, because the devices have size constraints. Hospitals can choose how many small, medium, and large products each hub can store, and (similar to the kiosk) they can choose whether to stock medium and large products in the same column, or in separate columns.

4-4 shows the two illustrative options for stocking medium and large products in the hub. Suppose the hub has one column dedicated to large products and one column dedicated to medium products (as shown on the left figure of 4-4), and assume that each column can hold C products. In this case, the model would aggregate all the large products, rank them by historical usage, and choose the C most popular large items. It would repeat the same steps for medium products. On the other hand, if the hub was organized more like the right figure of 4-4, with large and medium products shelved together, the model would aggregate all the large and medium products together, and would pick the 2C medium and large items with the highest historical usage values. In both cases, the remaining spots would be filled with the most popular small products.

6.5 Sort Hub Products by Size and Attributes

After products are selected, they must be placed into the hub in sorted order. The hub sorting process exactly matches the shelf model's design. Products are first sorted by product group, and hospitals can then choose to sort them by any or all of the following attributes: suture size, closures, needle length, and needle point. Not only can the number of sorting attributes and the order of attributes be adjusted, but the "correct" sorted order for each attribute can also be changed (e.g. needle length can increase or decrease).

The only major difference between how the shelf is sorted and how the hub is sorted, is that the hub is limited by size restrictions, while the shelf is not. The hub's columns are configured such that they can only hold one category of sizes: small, medium, large, or large and medium. This limitation on which products fit into each column often requires shifting the sorting order.

For example, the top part of 6-5 shows how products in a hub would be sorted if there were no size constraints. The natural sorting order (as dictated by the attribute rules established by the hospital) intermixes small, medium, and large products. However, physically implementing this sorted order is impossible, because the columns that hold small products cannot hold large or medium products. So, the products must be arranged according to the diagram shown in the bottom part of 6-5. In this example, column 6 holds medium products and column 7 holds large products. So, if the hub model encounters a medium or large product while placing sorted items into the device, it will move this item to the next available slot in a medium or large column.

Large Medium Small

SKU order with no size constraints

Hub 1	Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	Col 7
Row 1	A	B	C	D	E	F	G
Row 2	H	I	J	K	L	M	N
Row 3	O	P	Q	R	S	T	U

SKU order with a medium and large column

Hub 1	Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	Col 7
Row 1	A	C	D	E	F	J	B
Row 2	G	H	I	L	M	O	K
Row 3	P	Q	R	S	U	T	N

Figure 6-5: (Top) The ideal sorted order of products in a hub with no size constraints; (Bottom) the actual sorted order of products in a hub where columns 1-5 are small, 6 is medium, and 7 is large

6.6 Hub Model Outputs

Once the model completes its two main steps of selecting products for each hub and sorting the products within each hub, it outputs a file with a visualization of the final planogram. 6-6 shows an illustrative output for a hospital with two hubs: Hub 1 and Hub 2. Both hubs have 48 total slots, and both carry small products in columns 1-6, medium products in column 7, and large products in column 8.

Similar to the shelf planogram output, the hub planogram is color-coded according to product group (e.g. Monocryl, Fast Absorbing Gut, etc.). However, unlike the shelf, the hub has size constraints that shift the order of its planogram. For example, in 6-6, Monocryl 8 and Monocryl 3 are separated from the rest of the Monocryl products, because they are large items that can only fit in column 8.

It is also worth noting that the set of products in Hub 1 is different than the set of products in Hub 2. They largely contain different product groups, with Hub 1 stocking Gut, Permahand, V-LOC, and Ethibond sutures, and Hub 2 holding Prolene, Vicryl, PDS, and Stratafix sutures. Even the product groups that they do share (Monocryl and Ethilon) have different sets of products within each hub. This disparity in the products stocked stems from the fact that Hub 1 and Hub 2 cater to different procedures. The output indicates that operating rooms close to Hub 1 require a different set of products than the ORs close to Hub 2.

Hub 1	Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	**Col 7**	*Col 8*
Size	Small	Small	Small	Small	Small	Small	**Medium**	*Large*
Row 1	Monocryl 5	Monocryl 9	Monocryl 10	Monocryl 2	Monocryl 4	Monocryl 7	**Fast Gut 2**	*Monocryl 8*
Row 2	Monocryl 11	Monocryl 1	Monocryl 6	Fast Gut 4	Fast Gut 7	Fast Gut 3	**Plain Gut 3**	*Monocryl 3*
Row 3	Plain Gut 2	Plain Gut 4	Ethilon 13	Ethilon 8	Ethilon 12	Ethilon 17	**Ethilon 7**	*Ethilon 3*
Row 4	Ethilon 10	Perma 4	Perma 6	Perma 8	Perma 7	Perma 2	**Perma 1**	*V-LOC 5*
Row 5	Perma 3	Perma 5	V-LOC 2	V-LOC 7	V-LOC 3	V-LOC 1	**V-LOC 4**	*V-LOC 6*
Row 6	V-LOC 8	Ethibond 4	Ethibond 3	Ethibond 9	Ethibond 11	Ethibond 8	**Ethibond 13**	*V-LOC 9*

Hub 2	Col 1	Col 2	Col 3	Col 4	Col 5	Col 6	**Col 7**	*Col 8*
Size	Small	Small	Small	Small	Small	Small	**Medium**	*Large*
Row 1	Monocryl 9	Monocryl 10	Monocryl 4	Monocryl 6	Prolene 10	Prolene 20	**Prolene 8**	*Monocryl 8*
Row 2	Prolene 18	Prolene 7	Prolene 1	Prolene 13	Prolene 3	Ethilon 13	**Ethilon 8**	*Ethilon 3*
Row 3	Ethilon 12	Ethilon 17	Ethilon 10	Ethilon 5	Ethilon 14	Ethilon 1	**Ethilon 15**	*Ethilon 9*
Row 4	Ethilon 2	Ethilon 18	Ethilon 11	Ethilon 16	Ethilon 4	Ethilon 6	**Ethilon 7**	*Vicryl 5*
Row 5	Vicryl 1	Vicryl 2	Vicryl 6	Vicryl 9	Vicryl 3	Vicryl 4	**Vicryl 8**	*Stratafix 7*
Row 6	Stratafix 2	Stratafix 1	Stratafix 9	Stratafix 4	PDS 6	PDS 8	**PDS 5**	*Stratafix 3*

Legend
Monocryl
Fast Absorbing Gut
Plain Gut
Ethilon
PERMAHAND SILK
V-LOC
Ethibond
Prolene
Coated Vicryl
PDS
Stratafix

Figure 6-6: Illustrative hub model output

Aside from the main product view shown in 6-6, the hub model also provides other visualizations to help the AITA™ team confirm the validity of the model's result. The 4 other views shown are the description, group, size, and printing layouts. The description view shows the full product name of every item, the group layout shows the product group of every item, and the size layout shows every item's size (small, medium, or large). Finally, the printing layout includes both products' SKU numbers, as well as their full name. This view can be directly printed out and shown to hospital staff members during the planning and implementation phases.

6.7 Hub Model Results

Similar to the kiosk and shelf models, the hub planogram creation model improves the accuracy and efficiency of the planogram creation process. The accuracy improvements are two-fold. First, the model's planogram does not have the sorting errors that are sometimes present in manually created models. Second, the model stocks the products that cater to the real needs of operating rooms, rather than hypothetical needs. For example, even if hospitals try to dedicate one set of ORs to plastic surgery, in practice, the ORs are often used for a mix of specialties. The manual process would stock the most popular plastic surgery products, while the model would stock the most popular products that were actually used historically in the ORs.

As shown in 6-7, the model also creates large efficiency gains for hospitals and for the AITA™ team. In total, using the model instead of the manual creation process will save an annual average of 50.1 hours per hub per hospital. The majority of this time disparity comes from the 41.9 hours saved when creating a new hub planogram with the model. The remainder of the time saved (8.2 hours) stems from a more efficient process to iterate on an existing hub planogram.

	Frequency	Steps	Model-based process	Manual process	Average Time Saved per Year
Creating a new hub planogram	2-4 times per year	Aggregate / clean data	1 hour	2 days (16 hours)	**41.9 hours** = (16 – 2.03) hours * 3 per year
		Adjust parameters	30 minutes		
		Run model	2 minutes		
		Check output	30 minutes		
Iterating on a hub planogram	2-8 times per year	Update data / parameters	5 minutes	1-3 hours	**8.2 hours** = (2 – 0.37) hours * 5 per year
		Run model	2 minutes		
		Check output	15 minutes		

Figure 6-7: Time taken to create and update hub planograms when using the model vs. a manual process

Under the current manual process, it takes roughly 2 days to create a planogram. This is the longest creation process because it involves both manual product selection, as well as manual product sorting. Meanwhile, the model takes only around 2 hours from start to finish – including time to aggregate the data, run the model, and check the output. These 14 hours in time savings are further compounded because hubs are updated around 2-4 times per year. Similar to the kiosk, these updates ensure that hubs are accurately reflecting changes in procedure trends and in preference cards.

6-7 also shows that the model saves time when iterating on / adjusting hub planograms. Manual iterations usually take around 1-3 hours, with different changes taking different amounts of time. The fastest changes would involve adding or removing products and shifting the sorted products accordingly. The longest changes would involve changing which operating rooms are covered by each hub or changing the underlying sorting rules that dictate the products' order. However, with the model, these changes all take a flat time of around 20 minutes, because they would all just require a quick tweak in the model's input parameters. Furthermore, these time savings of around 0.5 to 2.5 hours per iteration happen roughly 1-2 times when a new hub is implemented, or 2-8 times a year.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 7

Conclusion

7.1 Overview

This project created automated models that improved the planogram creation process for the AITA™ Smart System. The original creation process was manual, leading to time-consuming steps and imperfect results. Meanwhile, the new, model-based approach enables fast planogram creation, easy iteration, and accurate solutions. These results are achieved through 3 separate models for each of the AITA™ devices: the Kiosks, Smart Shelves, and Mobile Hubs.

1. **Kiosk Model:** Uses knapsack optimization to select which products would have the highest coverage rate of historical procedures. Then allocates empty bins to popular products and places products in the kiosk based on size (small, medium, large) and usage.

2. **Shelf Model**: Uses bin-packing optimization to fit as many product groups as possible into fixed-sized racks without splitting up the groups. Sorts products within each group based on suture size, closures, needle length, and needle type attributes.

3. **Hub Model**: Calculates the usage of every potential hub product and selects the most popular products to add to the device. Sorts products based on size

(small, medium, large) and based on attributes (suture size, closures, needle length, and needle type).

7.2 Results Summary

The AITA™ team has already successfully implemented the three models and has used them to design and update kiosk, shelf, and hub planograms. Using these models has enabled the team to improve upon both the accuracy and the efficiency of the manual creation process. A summary of the improvements can be found below in 7-1, 7-2, and 7-3.

Model	Accuracy improvement
Kiosk	• Coverage rate of a representative hospital increased from 73.1% to 78.2%
Shelf	• Minimal number of product groups split between racks • No sorting mistakes
Hub	• Higher coverage of add-on requests • No sorting mistakes

Figure 7-1: Summarized Accuracy Results

The summarized results in 7-2 and 7-3 show that, in a typical hospital with 1 kiosk, 1 shelf, and 2 hubs, this project's models are projected to save, on average, 145 hours (18 business days) per hospital per year. The greatest source of time saved is in the automated process to design new hub models (50.1 hours saved annually for every hub), which is further compounded when a hospital leverages multiple hubs.

Create a new planogram		Kiosk model	Shelf model	Hub model
Frequency (occurrences per year)		2 - 4	1	2-4
Manual completion time		1 day (8 hours)	1.5 days (12 hours)	2 days (16 hours)
Model completion time	Aggregate/clean data	2 hours	1 hour	1 hour
	Adjust parameters	15 minutes	30 minutes	30 minutes
	Run model	2 minutes		
	Check output	30 minutes	45 minutes	30 minutes
Average Time Saved per Year		**15.7 hours**	**9.7 hours**	**41.9 hours**

Figure 7-2: Summarized Efficiency Results (Creating a New Planogram)

Iterate on a planogram		Kiosk model	Shelf model	Hub model
Frequency (occurrences per year)		2 - 8	4 - 8	2 - 8
Manual completion time		1 - 2 hours	2 – 4 hours	1 - 3 hours
Model completion time	Update data / Parameters	15 minutes	5 minutes	5 minutes
	Run model	2 minutes		
	Check output	15 minutes	30 minutes	15 minutes
Average Time Saved per Year		**4.8 hours**	**14.3 hours**	**8.2 hours**

Figure 7-3: Summarized Efficiency Results (Update / Iterate on a Planogram)

7.3 Implications

Now that the AITA™ team has implemented this project's kiosk, shelf, and hub models, the models' improved accuracy and efficiency will benefit stakeholders across the chain of care. The AITA™ data science team will save an average of 145 hours per hospital by creating and iterating on planograms using the model. As AITA™ continues to expand to serve additional hospital clients, these time savings will scale, helping to free up significant capacity for the AITA™ data science team to focus on value-add business development projects. The data science team will also be able to more easily test new hypotheses about their system (e.g. checking if coverage rates improve if more large and medium slots are added to the kiosk).

Meanwhile, the AITA™ commercial team will be able to garner higher customer satisfaction rates due to the faster implementation times and more accurate planograms. Taken together, these efficiency and accuracy gains will enable the total team to serve more clients- scaling their business without needing to expand their headcount.

On the provider side, hospitals will be able to more quickly integrate the AITA™ Smart System into their existing inventory management frameworks. Once the system is installed, hospitals will then receive faster, more regular updates than before. Furthermore, thanks to the improved accuracy of their planograms, providers will see improvements in their own operations due to higher coverage rates of their procedures and add-on requests. Hospitals will also be better equipped for waste management, because the regular planogram updates will more quickly flag which products are low-runners and should potentially be eliminated from the inventory.

Finally, patients will benefit from improvements to the AITA™ planogram creation process, because of the efficiency gains realized by their providers. Higher-coverage hubs will reduce idle time in surgeries, decreasing the chance of complications. Higher-coverage kiosks will increase the number of procedures that can be filled automatically, freeing capacity for members of the hospital staff to spend more time on direct patient interaction and care.

7.4 Future Work

Now that the AITA™ team is using the models presented in this project to automatically build and update planograms, the team can further improve their process by continuing to streamline other parts of the Smart System.

- **Recommend device quantities**: Currently the AITA™ team recommends how many kiosks, shelves, and hubs a hospital should have based on a manual assessment of their inventory. The team could automate this process by creating a model that calculates key metrics (e.g. the coverage rate of N kiosks or the total number of unique SKUs) and using the metrics to recommend the ideal number of devices to be placed in the hospital.

- **Product eligibility**: The AITA™ team currently manually evaluates each product in a hospital to assess if it is physically compatible with Smart System devices (e.g. if it is the right size or if it has a barcode for the kiosk to scan). This time consuming process could potentially be replaced with a machine learning model that predicts the likelihood of a product's eligibility based on catalog information (e.g. brand, product group, full description, etc.)

- **Suture consolidation**: At the beginning of each implementation the AITA™ team manually reviews all the eligible products in the hospital and assesses if there are any opportunities for consolidation. For example, it the team identifies two products that are identical in every attribute except for thread color, they will recommend that the hospital consolidates its inventory by buying just one of the SKUs going forward. This process could be automated through an attribute matching program that automatically flags which products could be reasonable substitutes.

- **Demand forecasting**: Demand forecasts could be built using machine learning models on past procedure schedules and purchasing data. These forecasts could improve the Smart System on multiple fronts. First, AITA™ could use demand predictions to help hospitals make ordering decisions, reducing the risks of unnecessarily high inventory costs or of stock-outs. Second, the AITA™ team could use demand forecasts to build planograms that meet the needs of future, rather than the needs of historical procedures.

With these additions and with the automated planogram creation models presented in this thesis, AITA™ can help to generate even more value across the healthcare system. Boosting the efficiency and accuracy of the AITA™ Smart System can help drive profits for Johnson & Johnson, reduce provider burden, and increase time available for direct patient care.

THIS PAGE INTENTIONALLY LEFT BLANK

Bibliography

[1] Neurosurgical Atlas. "Suturing and Closure". In: *Neurosurgical Atlas* (2024). URL: https://www.neurosurgicalatlas.com/volumes/principles-of-cranial-surgery/suturing-and-closure.

[2] Ruibin Bai. "An investigation of novel approaches for optimising retail shelf space allocation". PhD thesis. University of Nottingham, 2005.

[3] Constant Berkhout. "Universal Merchandising Guidelines". In: *Assortment and Merchandising Strategy: Building a Retail Plan to Improve Shopper Experience*. Cham: Springer International Publishing, 2019, pp. 165–182. ISBN: 978-3-030-11163-2. DOI: 10.1007/978-3-030-11163-2_8. URL: https://doi.org/10.1007/978-3-030-11163-2_8.

[4] Teresa Bianchi-Aguiar et al. "Allocating products on shelves under merchandising rules: Multi-level product families with display directions". In: *Omega* 76 (2018), pp. 47–62. ISSN: 0305-0483. DOI: https://doi.org/10.1016/j.omega.2017.04.002. URL: https://www.sciencedirect.com/science/article/pii/S0305048317303699.

[5] A. Hande Erol Binguler, Serol Bulkan, and Mustafa Agaoğlu. "A Heuristic Approach for Shelf Space Allocation Problem". In: *Journal of Management and Information Science* 4.1 (2016), pp. 38–44. DOI: 10.17858/jmisci.89213.

[6] Burcu Caglar Gencosman and Mehmet A. Begen. "Exact optimization and decomposition approaches for shelf space allocation". In: *European Journal of Operational Research* 299.2 (2022), pp. 432–447. ISSN: 0377-2217. DOI: https://doi.org/10.1016/j.ejor.2021.08.047. URL: https://www.sciencedirect.com/science/article/pii/S0377221721007396.

[7] Kateryna Czerniachowska and Marcin Hernes. "Simulated Annealing Hyper-Heuristic for a Shelf Space Allocation on Symmetrical Planograms Problem". In: *Symmetry* 13.7 (2021). ISSN: 2073-8994. DOI: 10.3390/sym13071182. URL: https://www.mdpi.com/2073-8994/13/7/1182.

[8] Tobias Düsterhöft, Alexander Hübner, and Kai Schaal. "A practical approach to the shelf-space allocation and replenishment problem with heterogeneously sized shelves". In: *European Journal of Operational Research* 282.1 (2020), pp. 252–266. ISSN: 0377-2217. DOI: https://doi.org/10.1016/j.ejor.2019.09.012. URL: https://www.sciencedirect.com/science/article/pii/S0377221719307520.

[9] Ethicon. *AITA Smart System*. 2024. URL: https://www.jnjmedtech.com/en-US/product/aita-ethicon-suture-management.

[10] Ethicon. *Wound Closure Listings*. 2024. URL: https://www.ethicon.com/na/epc?lang=en-default.

[11] Hakan Gecili and Pratik J. Parikh. "Joint shelf design and shelf space allocation problem for retailers". In: *Omega* 111 (2022), p. 102634. ISSN: 0305-0483. DOI: https://doi.org/10.1016/j.omega.2022.102634. URL: https://www.sciencedirect.com/science/article/pii/S030504832200041X.

[12] H. Neil Geismar et al. "Maximizing Revenue Through Two-Dimensional Shelf-Space Allocation". In: *Production and Operations Management* 24.7 (2015), pp. 1148–1163. DOI: https://doi.org/10.1111/poms.12316. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/poms.12316.

[13] Alexander Hübner, Fabian Schäfer, and Kai N. Schaal. "Maximizing Profit via Assortment and Shelf-Space Optimization for Two-Dimensional Shelves". In: *Production and Operations Management* 29.3 (2020), pp. 547–570. DOI: https://doi.org/10.1111/poms.13111. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/poms.13111.

[14] Faiz Tuma Jessica Rose. "Sutures And Needles". In: StatPearls Publishing, 2023.

[15] Johnson Johnson. "Ethicon Product Catalog 2023". In: (2023).

[16] Ayyad Zartasht Khan et al. "Clinical Review: Suture Materials". In: *Tidsskr Nor Legeforen* 143 (2023). DOI: 10.4045/tidsskr.22.070.

[17] Miles Kurtz. *Planogram Optimization in Support of Small Format Retail Inventory Management*. 2023.

[18] Neha Mehta, Neha Mehta, and Pawan Kumar Chugan. "Impact of Visual Merchandising on Consumer Behavior: A Study of Furniture Outlets". In: *Universal Journal of Management* 2.6 (2014), pp. 207–217. URL: https://ssrn.com/abstract=2481859.

[19] Christopher M. Miller et al. "Optimizing and Evaluating Retail Assortments for Infrequently Purchased Products". In: *Journal of Retailing* 86.2 (2010). Special Issue: Modeling Retail Phenomena, pp. 159–171. ISSN: 0022-4359. DOI: https://doi.org/10.1016/j.jretai.2010.02.004. URL: https://www.sciencedirect.com/science/article/pii/S0022435910000114.

[20] Medical Device Network. "Surgical sutures market share in the US: recent trends". In: (2023).

[21] Erjen van Nierop, Dennis Fok, and Philip Hans Franses. "Interaction Between Shelf Layout and Marketing Effectiveness and Its Impact on Optimizing Shelf Arrangements". In: *Marketing Science* 27.6 (2008), pp. 1065–1082. DOI: 10.1287/mksc.1080.0365. URL: https://doi.org/10.1287/mksc.1080.0365.

[22] Manuel Ostermeier, Tobias Düsterhöft, and Alexander Hübner. "A model and solution approach for store-wide shelf space allocation". In: *Omega* 102 (2021), p. 102425. ISSN: 0305-0483. DOI: https://doi.org/10.1016/j.omega.2021.102425.

[23] Robert P. Rooderkerk and Donald R. Lehmann. "Incorporating Consumer Product Categorizations into Shelf Layout Design". In: *Journal of Marketing Research* 58.1 (2021), pp. 50–73. DOI: 10.1177/0022243720964127. URL: https://doi.org/10.1177/0022243720964127.

[24] David Rosenfeld. *Suture reel dispenser with end of suture indicator*. U.S. Patent 6467612B1, May. 2000.

[25] Darshan S. Shah, Bryan K. Lawson, and Michael Yaszemski. "2.5.9 - Description and Definition of Adhesives, and Related Terminology". In: *Biomaterials Science (Fourth Edition)*. Fourth Edition. Academic Press, 2020, pp. 1181–1198. ISBN: 978-0-12-816137-1. DOI: https://doi.org/10.1016/B978-0-12-816137-1.00075-1. URL: https://www.sciencedirect.com/science/article/pii/B9780128161371000751.

[26] Ivana Štulec, Kristina Petljak, and Anja Kukor. "The Role of Store Layout and Visual Merchandising in Food Retailing". In: *European Journal of Economics and Business Studies* 2.1 (2016), pp. 138–151. URL: https://doi.org/10.26417/ejes.v4i1.p138-151.

[27] Ming-Hsien Yang. "An efficient algorithm to allocate shelf space". In: *European Journal of Operational Research* 131.1 (2001), pp. 107–118. ISSN: 0377-2217. DOI: https://doi.org/10.1016/S0377-2217(99)00448-8. URL: https://www.sciencedirect.com/science/article/pii/S0377221799004488.

[28] Ju Zhao, Yong-Wu Zhou, and M.I.M. Wahab. "Joint optimization models for shelf display and inventory control considering the impact of spatial relationship on demand". In: *European Journal of Operational Research* 255.3 (2016), pp. 797–808. ISSN: 0377-2217. DOI: https://doi.org/10.1016/j.ejor.2016.05.025. URL: https://www.sciencedirect.com/science/article/pii/S0377221716303459.