FLOOR HEALTH PREDICTIVE SUPPORT FOR HIGHLY AUTOMATED DISTRIBUTION CENTERS

by

Emily Stinson

B.S.E., Civil and Environmental Engineering, Princeton University, 2014

Submitted to the MIT Sloan School of Management and the Department of Civil and Environmental Engineering in Partial Fulfillment of the Requirements for the Degrees of

Master of Business Administration
and
Master of Science in Civil and Environmental Engineering

In conjunction with the Leaders for Global Operations Program at the Massachusetts Institute of Technology

June 2019

Signature redacted

Signature of Author...
MIT Sloan School of Management
Department of Civil and Environmental Engineering
May 10, 2019

Signature redacted

Certified by ....
Yanchong (Karen) Zheng, Thesis Supervisor
Associate Professor, MIT Sloan School of Management

Signature redacted

Certified by ....
Saurabh Amin, Thesis Supervisor
Associate Professor, Department of Civil and Environmental Engineering

Signature redacted

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

Signature redacted

Accepted by ....
Heidi Nepf, Chair, Graduate Program Committee
Department of Civil and Environmental Engineering

[THIS PAGE IS INTENTIONALLY LEFT BLANK]

FLOOR HEALTH PREDICTIVE SUPPORT FOR HIGHLY AUTOMATED DISTRIBUTION CENTERS

by
Emily Stinson

Submitted to the MIT Sloan School of Management and
the Department of Civil and Environmental Engineering on May 10, 2019
in Partial Fulfillment of the Requirements for the Degrees of Master of Business Administration
and Master of Science in Civil and Environmental Engineering

ABSTRACT

While automated mobile inventory systems have greatly increased productivity, it has also created a new set of operational challenges. Floor health events, such as fallen product, spills, disabled robots, and floor access can degrade overall floor performance by obstructing access to product, forcing robots to re-route to less efficient paths, exacerbating congestion, increasing idle time, and potentially reducing throughput. Floor health issues are interdependent and have cascading effects, making their impacts difficult to track, visualize, and address. Reactive support and reliance on training and adoption of best practices is not scalable. As the network continues to grow, there is a need to improve real-time visibility and preventative measures into floor conditions.

This project consisted of five main phases: research, hypothesis, testing, evaluation, and implementation. The research phase was dedicated to developing an understanding of the current processes and problem statement. Then a testable hypothesis was constructed based on observations and data exploration. The hypothesis was tested via simulations and statistical analysis. The evaluation phase included analyzing the implications and use-cases of the results. The last phase of the project included developing and implementing selected applications.

The model development phase of the project included simulation experiments where the dependent variable collected was the percentage change in average throughput rate and a multitude of potential explanatory features were tracked. Analysis of this data revealed that some of the best predictors of degradation of throughput rate were the types of floor cells being blocked. There is wide range of impactful applications of these findings, including diagnostic checks to help root cause issues, automated notifications that highlight deteriorating floor conditions, automated user path planning, actionable floor metrics, and prioritization of work. Automated notifications to proactively identify deteriorating floor conditions, real-time prioritization of tasks, and a diagnostic tool were the implementations focused on during this project.

Thesis Supervisor: Yanchong (Karen) Zheng
Title: Associate Professor, MIT Sloan School of Management

Thesis Supervisor: Saurabh Amin
Title: Associate Professor, Department of Civil and Environmental Engineering

[THIS PAGE IS INTENTIONALLY LEFT BLANK]

ACKNOWLEDGEMENTS

The completion of this thesis would not have been possible without the support of many individuals. I would like to express my gratitude to the following people:

My industry executive sponsor at Amazon Robotics and MIT Sloan Alumni, Glenn Ready, for sponsoring this project, welcoming me into the organization, and considering my input and recommendations.

My industry supervisors, Sagar Mohan and Ian Moulton, and their teams for providing valuable data, input, and thoughtful partnership.

My thesis advisors, Karen Zheng and Saurabh Amin, for their support and guidance throughout the project and while writing this document.

The staff from the MIT Leaders for Global Operations program, who make the LGO program possible and enabled this incredible six-month industry experience.

Last but not least, I would like to thank all my classmates, friends, and family, who provided unwavering moral support!

[THIS PAGE IS INTENTIONALLY LEFT BLANK]

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

[THIS PAGE IS INTENTIONALLY LEFT BLANK]

1 INTRODUCTION

1.1 Purpose

Technology has enabled automation of many traditionally manual warehouse operations. In traditional warehouses, humans must walk miles up and down aisles to find and store products. Autonomous mobile robotic technology has enabled highly automated inventory storage systems that eliminate this manual effort and instead allow product to be transported to and from a stationary human. While robotic inventory has greatly increased productivity, it has also created a new set of operational challenges. The primary objective of this project was to identify opportunities to improve visibility to system operators into deteriorating conditions that require manual intervention in highly automated distribution centers.

Data presented in this document has been disguised and details on internally developed products and processes have been abstracted to protect Amazon's proprietary information. Process rates, volumes and costs have been aggregated, normalized or presented as percentages. Formulas and equations describing the relationships between Amazon's processes have been adjusted to communicate concepts notionally.

1.2 Background

1.2.1 Amazon

Amazon.com is a well-known online retailer founded in 1994 by Jeff Bezos and headquarter in Seattle, Washington. Initially started as an online book store, Amazon has since greatly diversified its product offerings. Amazon has also expanded to businesses outside of retail, including producing electronic devices and providing web services. Today, Amazon's e-commerce business is supported by over 175 operating fulfillment centers (FC) and more than 150 million square feet of space across the world, where associates pick, pack and ship millions of customer orders [1]. Amazon is continuously looking for ways to improve throughput, reduce costs, and better serve its customers as it continues to grow its network in the United States and globally.

1.2.2 Amazon Robotics

Amazon Robotics (AR), previously Kiva Systems, was acquired in 2012 to become a wholly owned subsidiary of Amazon.com. Based in North Reading, Massachusetts, AR focuses on automating fulfilment center operations using various methods of robotic technology including autonomous mobile robots, sophisticated control software, language perception, power management, computer vision, depth sensing, machine learning, object recognition, and semantic understanding of commands [2]. The company is particularly known for its manufacture of autonomous mobile robots, shown in Figure 1.1, which enable a highly automated inventory storage system, but the company also provides other hardware and software tools, as well as technical support, to Amazon's network of fulfillment centers.

In 2014, Amazon began integrating the AR robotic inventory system into its fulfillment centers. There are four main types of Amazon fulfillment centers, demonstrated in Figure 1.2 below. Fulfillment centers are either sortable versus non-sortable and either traditional versus robotic. Sortable fulfillment centers contain products less than 18" in size length, allowing all products to fit on standardized inventory shelves. Non-sortable fulfillment centers contain items larger than 18" and irregularly shaped products.

This project focuses on Amazon Robotic sortable fulfillment centers, which are the most highly automated fulfillment centers, the most common building type currently in Amazon's network, and more representative of the fulfillment centers that Amazon plans to build in the future.

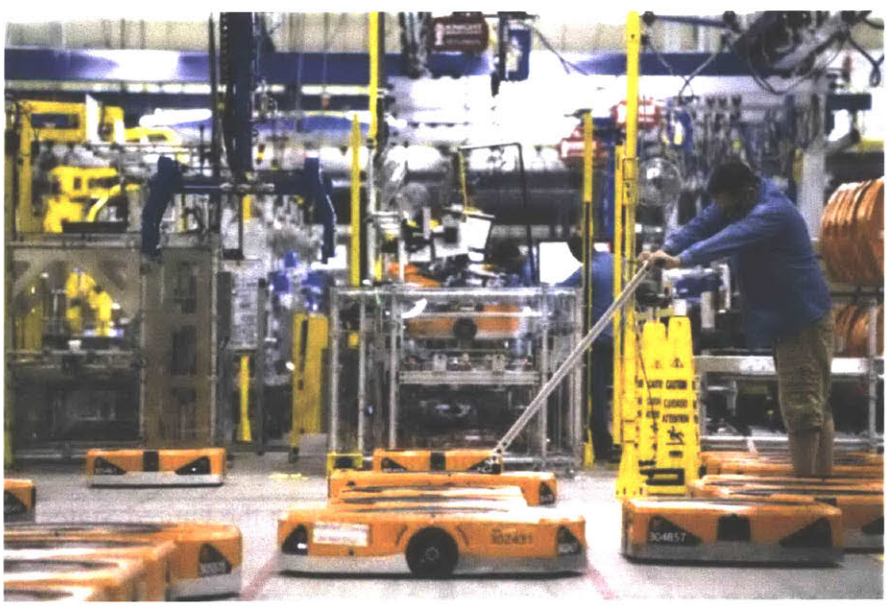

Figure 1.1 - Manufacture of mobile autonomous robots at AR [3]

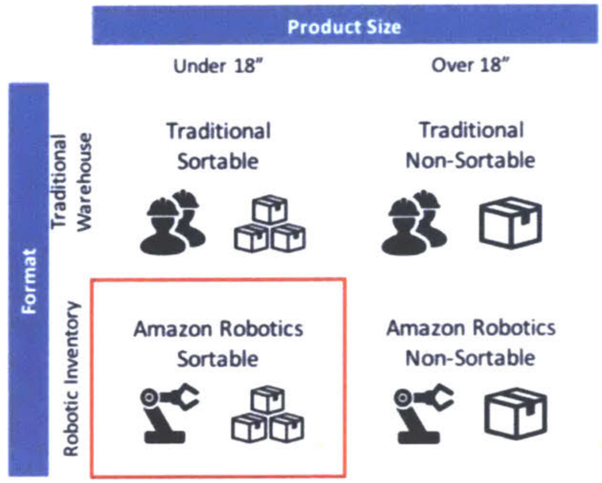

Figure 1.2 – Types of fulfillment centers [4]

1.2.3 AR Mobile Robotic Inventory System

The AR mobile robotic system occupies only a portion of the space in a fulfillment center and can be considered part of the storage operation, as shown in Figure 1.3. A simplified version of an AR system within an Amazon FC is depicted in Figure 1.4. The infrastructure for the inbound processes upstream to receive inventory and the outbound processes downstream to pack and ship customer orders are not depicted here and were outside the scope of this project.

The depicted area is generally referred to as the *AR floor* and confines the major moving parts of the AR system. Fulfillment centers equipped with these systems may have one to ten of these AR floors within the building depending on the building type. The size of an AR floor depends on the building type. Large buildings representative of the major building type in Amazon's fleet would have the scale of several American football fields.

In order to understand the challenges with a mobile robotic inventory system, it is important to have an understanding of the basic system components. The orange mobile autonomous robots, referred to as *drive units*, carry and transport the yellow inventory shelves, called *pods*, which can hold products on all four sides. The pods are moved by the drive units, which position themselves underneath a pod, raise it off of the floor, and drive it to a new location as shown in Figure 1.5. *Stations* are located around the floor perimeter. This is where associates interact with storage pods to store product or pick customer orders as shown in Figure 1.6. The drive units know their location and are able to travel across the floor by tracking barcodes along the floor. A typical AR floor will have tens of thousands of cells, thousands of drive units, and thousands of pods. All system components are connected by Wi-Fi to servers.

As opposed to a traditional fulfillment center, as shown in Figure 1.7, where people must manually search extensive aisles and shelves to retrieve and store inventory, the AR mobile robotic inventory system enables associates to be mostly stationary while the drives fetch and carry the product to and from the associates via the mobile inventory pods.

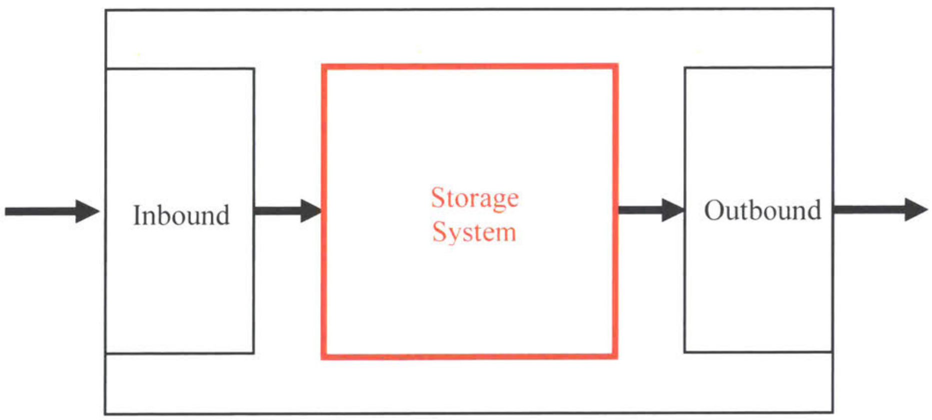

Figure 1.3 – Basic operations of a distribution center [5]

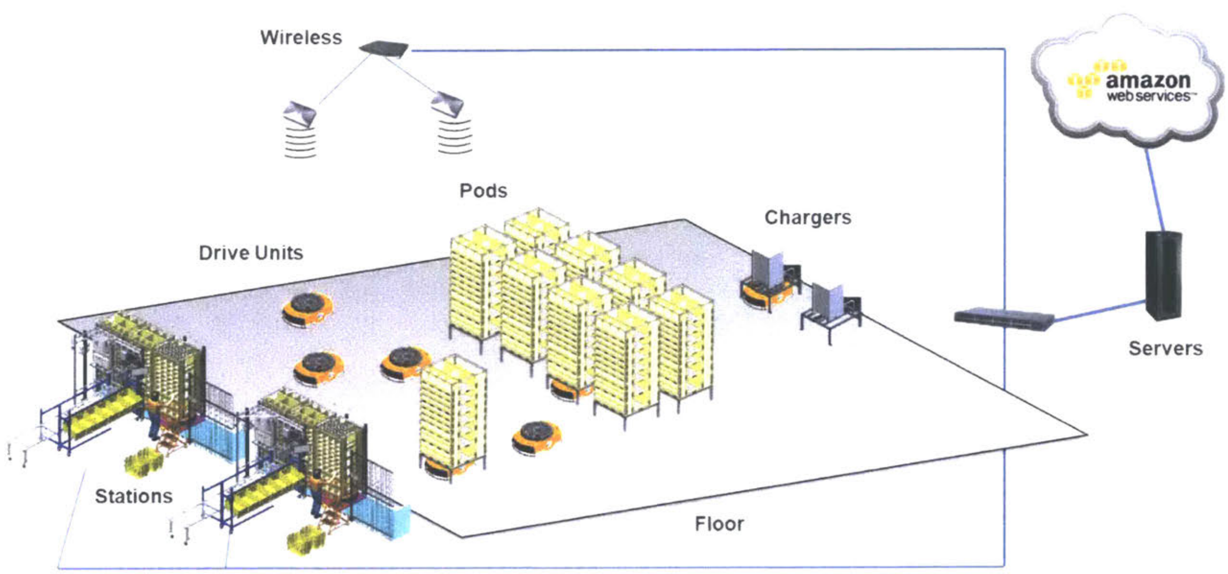

Figure 1.4 - Simplified AR mobile robotic inventory system in a fulfillment center

Figure 1.5 - Drive transporting inventory pod [6]

Figure 1.6 - Picking product at a station [7]

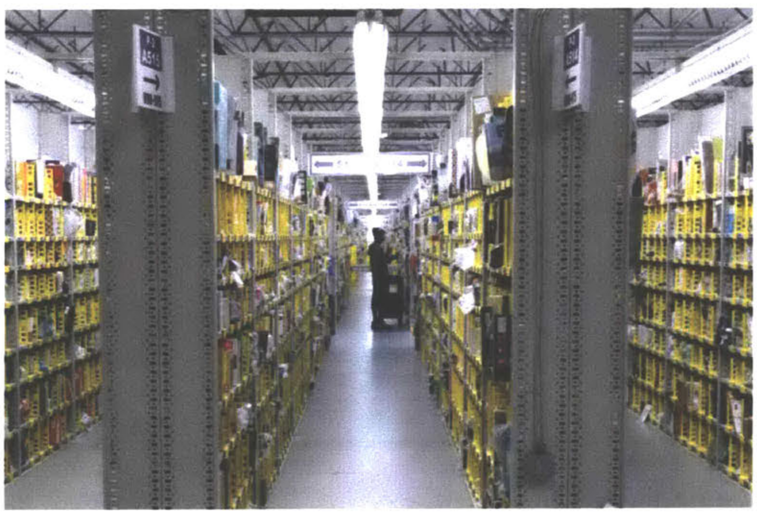

Figure 1.7 – Fulfillment center with manual storage system [8]

1.3 Problem Statement

Although the AR mobile robotic inventory system has greatly increased productivity, the system is susceptible to a new set of challenges that can degrade the efficiency of the system and require human intervention. One of these primary challenges is referred to as *poor floor health*. Poor floor health is depicted in Figure 1.9 by the additions in red. Events that would contribute to poor floor health include when products fall or spill out of the storage pods onto the AR floor. Since inventory is held in place by thin, flexible bands, it is a common occurrence for products to slip from the constantly moving inventory pods and products can easily break open and spill when falling from a high bin on the inventory pod. In fact, a fallen product can be seen in circled in red in Figure 1.8. Spills are particularly detrimental and can remain undetected for quite some time. The affected area will slowly grow as the spill is tracked around the floor by the moving drive units, eventually leading to multiple drive failures. Floor health issues also include when drive units disable due to hardware or software issues, or more often because they run over a fallen or spilled product, or dirty conditions prevent proper reading of the cell barcodes.

Figure 1.8 - Product fallen from moving pod [9]

Furthermore, when any of these issues occur, associates must restrict or take out of service large areas of the floor so that they can safely access the floor to rectify these issues. Thus, floor health issues are often interdependent and have cascading effects, making their impacts difficult to track, visualize, and address. These events can degrade operational performance by obstructing access to products that customers have ordered, forcing robots to re-route to less efficient paths, exacerbating congestion, increasing idle time, and potentially reducing throughput.

Figure 1.10 shows a photograph of an actual AR floor in an Amazon fulfillment center. It is a vast area that is difficult to monitor. Furthermore, this view from above is not typically possible in most fulfillment centers. An associate's view is from same floor level and their visibility into the floor beyond a few feet is severely limited. Poor floor health is the root cause of many high severity issues reported by Operations, particularly during times of high volume turnover, such as during Peak (the November/December holiday retail season) and as new fulfillment centers ramp-up and become more full. Today, AR support teams spend a significant amount of time trying to retrospectively root cause poor floor health, after the performance impact to operations has already been realized. Reactive support and reliance on training and adoption of best practices is not scalable. As Amazon continues to grow its FC network, there is a need to improve real-time and preventative visibility into floor conditions.

Thus, this project was part of a broader mission to provide more predictive and proactive services to identify issues before they impact on-site operations teams at fulfillment centers, in addition to reactively providing support and root causing issues after the fact as is typically the case today. The main objectives of this research project were to 1) develop a floor model that helps quantify the impact of these events and 2) use the model to develop proactive services that could provide Operations better visibility into these conditions.

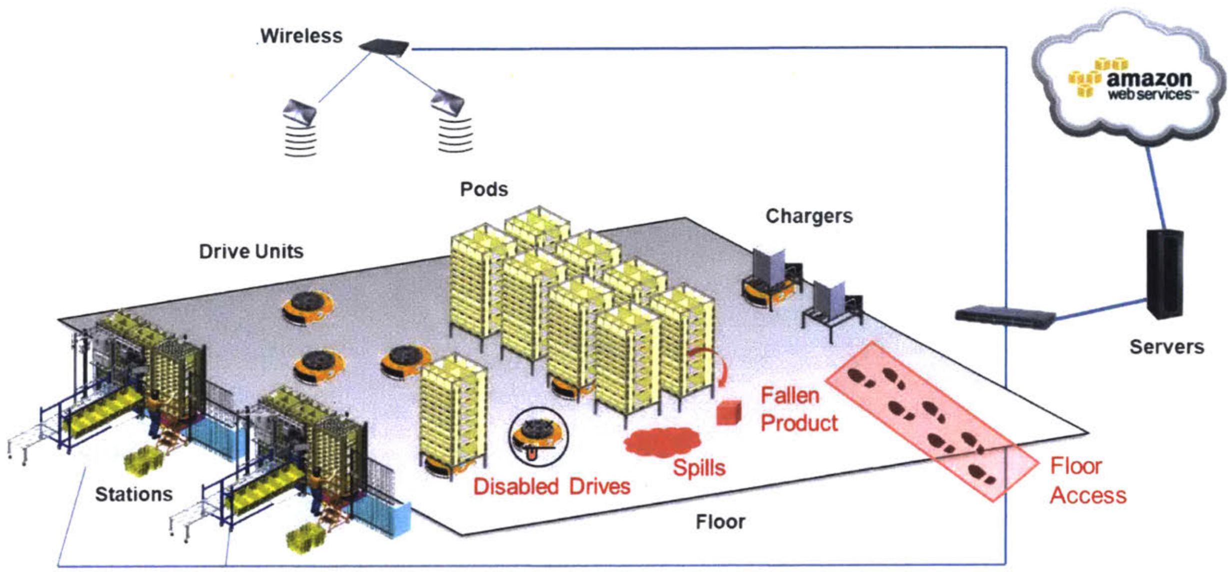

Figure 1.9 – Poor floor health refers to fallen/spilled product, disabled drives, and floor access etiquette

Figure 1.10 – An AR floor in a fulfillment center viewed from above [10]

1.4 Methodology

The approach to this project included the following main phases, which are described in greater detail below:

1. Research: developed understanding of the current processes and problem statement
2. Hypothesis: constructed hypothesis based on observations and data exploration
3. Testing: tested the hypothesis via simulations and statistical analysis
4. Assessment: analyzed the implications and use-cases of the results
5. Implementation: pursued the development and implementation of several applications

Significant time in the beginning of the project was dedicated to meeting stakeholders, understanding the current processes, investigating the problem statement, and exploring available data. Activities in this phase included several interviews with subject matter experts at AR, tutorials on internal tools and products, observation of AR technical support teams responding to FC issues, visits to FCs where the operation of the AR floor could be observed in-person, and interviews and shadowing of on-site operations personnel. These interviews were critical not only to mapping the current processes but also identifying the pain points and friction in current practices that would motivate the direction of the project, as described in Chapter 2. Identifying which teams owned and could provide access to various data sources was also an important activity in this phase.

The information gathered from the research phase was critical to formulating a testable hypothesis and narrowing the scope of the project. Testing the hypothesis consisted of running a controlled experiment via simulations. Considerable time in the testing phase was spent pre-processing data, setting up the simulation inputs, and running the simulations. Statistical analysis was then performed on the resulting simulation data and several models were compared. Chapter 3 will delve into details on the simulations, analysis, and modeling.

The assessment and implementation phases consisted primarily of product design and development activities, as described in Chapter 4. First, the opportunity space to apply the statistical findings was evaluated. Several rounds of additional site visits and interviews were performed with key stakeholders and potential end users to determine the most desired use-cases, implementation requirements and difficulty, and product design features. The final phase of the project was implementation of the selected application.

2 CURRENT STATE ANALYSIS

This chapter delves into information gathered during the research phase of this project. It describes the organizational structure, the available data, and maps the current state of processes pertinent to the problem of poor floor health.

2.1 Organizational Structure

While the AR headquarters for research and development, engineering, and the manufacturing site for the mobile robotic drive units are based out of North Reading, Massachusetts, Amazon's headquarters are located in Seattle, Washington and its network of fulfillment centers span the globe. The geographical dispersion of the company creates unique challenges and complexities in operations and support. The following paragraphs provide a high-level overview of the roles and responsibilities of relevant teams and organizations which are foundational to many of the recommendations and product design choices made later in the project.

Fulfillment center operations are overseen by a hierarchy of operations managers. Leads are assigned to oversee the thousands of associates working on various operations in the inbound, outbound, and AR floor processes. Two of the primary metrics tracked by operations to gauge the performance of the AR floor include the throughput rate (the number of product units extracted from the robotic inventory field at a station per unit of time) and idle time (the amount of time an associate could have been picking product units but didn't have work available to them).

As is common in this industry, a remote technical support team is dedicated to receiving, routing, responding, root causing, and finally closing issues reported by on-site operations teams. Issues are assigned a priority ranking based on their perceived impact to operations. Then the technical support team must route the issue by identifying the right team (e.g. product, software, hardware, firmware, etc.) who has the expertise to help resolve the problem. The technical support team coordinates and facilitates conference calls with the relevant teams and the on-site operations leads. If an issue increases in severity without being resolved, procedures are in place to escalate the issue.

Naturally, the issues can be very specific and easy to route and solve. But more often, the issues are vague and their root causes are unknown. For example, operations may notice irregularities in their performance metrics, such as drop in throughout rate or an increase in idle time, but are unable to immediately pinpoint the root cause. In these cases, the technical support team may spend considerable time investigating the issue by monitoring a variety of internally developed metrics and dashboards and analyzing historical data to get a sense of the conditions leading up to the report of the issue. In many cases, the issue resolves itself overtime (e.g. the performance metrics return to normal) without any special intervention and the technical support team assigns the issue a root cause based on their investigation and provides recommendations for best practices to avoid the issue in the future.

2.2 Floor Health

2.2.1 Floor Health Monitoring

An internally developed kindle application is the tool primarily used by on-site personnel to monitor the health of the AR floor. This application provides a visualization of the floor area and signals where problems have occurred. The drive units, equipped with cameras, detect when an object is in their path. The location of this obstruction is logged and reported in the application. Similarly, if a drive unit disables, its location is logged and highlighted in the application. The application does not show fallen product or spills that have not been visually detected by the drive units.

The current practice is for associates to recover any and all issues regardless of where they are located on the floor as quickly as possible. The associates are provided rules of thumb when deciding which issue to tackle next. It is suggested that associates prioritize issues that block access to a station and those that have dwelled the longest. Many sites have developed their own intuition and best practices. Moreover, shadowing of floor monitors revealed that the most common consideration is actually the proximity of the issue to their current location given the expanse of the floor and the time it would take to walk all the way to the opposite side.

There are many metrics and dashboards used to track the management of floor health. Dwell-time based metrics are some of the most commonly used in the industry and encourage clearing issues that have been on the floor the longest. While the frontline associates monitoring the floor primarily utilize the kindle application to gauge the status of the floor, the overseeing supervisors with access to laptops monitor other dashboards designed to provide a higher-level overview of floor health. For example, Figure 2.1 shows one of the primary dashboards used. The dashboard would give the supervisor insight into the metrics for all five floors under their purview. The "Floor Rating" is calculated from only 11 of the 20 individual metrics listed in the dashboard.

For every factor that the floor violates, 100% / 11 = 9.09% is deducted from the total rating of 100%. While these individual factors highlight undesirable conditions, it is not clear how only these 11 factors were chosen, why they are weighted equally, and their relationship to overall floor performance, thus sometimes making the overall rating difficult to interpret and take action on.

Figure 2.1 – Floor health dashboard (data has been blurred to protect confidential information)

2.2.2 Common Floor Health Problems

Investigation of historical issue tickets assigned a root cause related to poor floor health and shadowing of the AR floor at several fulfillment centers revealed some of the common problems and pain points encountered when dealing with floor health. Many problems are related to the inherent layout of the floor.

The AR floor is a grid system, made up of thousands of square cells. The inventory pods are packed tightly within this grid, with aisles left clear for ease of movement, as illustrated in Figure 2.2. Each cell on the floor has a classification. For example, cells where inventory pods can be placed are categorized as "storage" cells. Cells that remain clear for drive units to move are categorized as "travel" cells. The travel cells along the perimeter ring of the floor that lead into stations are anecdotally referred to as "the highway," as this is the main thoroughfare for the drive units. As part of the strategy to reduce idle time of the associates, several pods line up in front of stations, creating a buffer to ensure there is enough work waiting for the associates. These cells leading into a station have a categorization of "queue" cells. The cell at which the pod is in position for an associate to interact with the pod at the station is the "pick" cell.

As previously mentioned, when an issue occurs on the floor, such as a fallen product or a disabled drive, an associate monitoring the floor must restrict or take out of service large areas of the floor so that they can safely enter the robotic field to recover the issues. Obstructions or disabled drives carrying pods in travel cells, particularly in the highway, are considered particularly problematic and a top priority to rectify as they often prevent product for customer orders from getting to stations and they block valuable travel space that may have been on the optimal routes for other pods carrying product, creating congestion and forcing pods to reroute and take less efficient paths to stations. This can directly increase idle time and degrade pick rates. Whereas, obstructions or disabled drives in the storage area are considered less pressing since they block less heavily traveled areas and create less congestion. Pods can more easily be rerouted around issues in the middle of the storage area, such that they have less potential to increase to idle time and degrade pick rates compared to issues in the travel lanes.

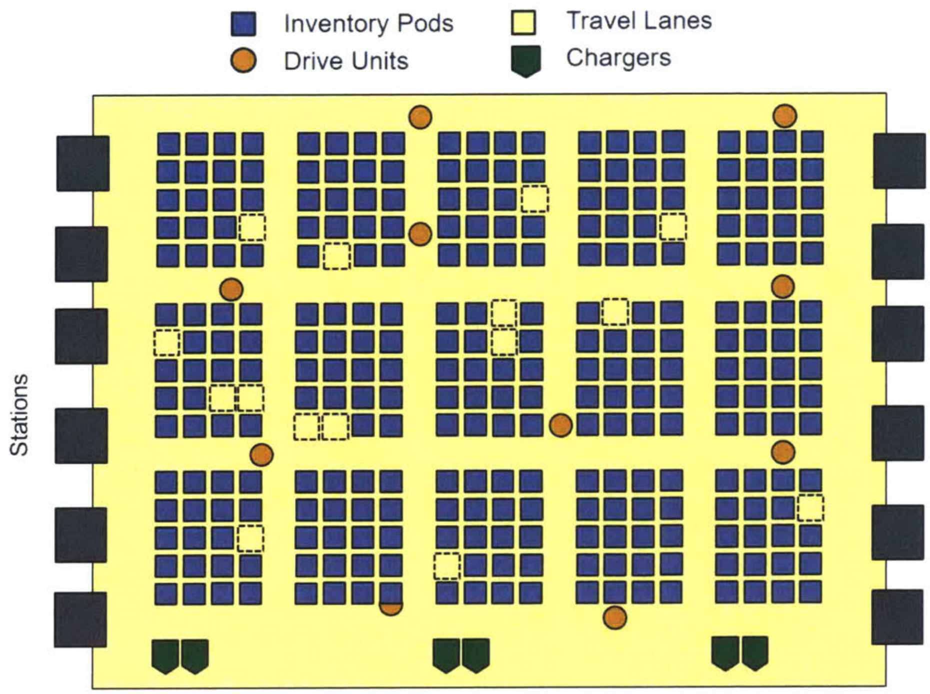

Figure 2.2 – Simplified grid layout of the AR floor

Commonly recurring problems included the following and are described in more detail below:

- Accessing the floor to fix an issue is more impactful than the initial issue
- Poor etiquette when accessing the floor and inefficient routes
- Lack of awareness for what others are doing and compounding floor access effects
- Restricting too much of the floor at one time

Observations revealed that the act of accessing the floor to fix an issue can sometimes be more impactful than the initial issue itself. For example, Figure 2.3 depicts a scenario with a disabled drive blocking a storage cell. Accessing and rebooting that drive unit required blocking many more cells than the disabled drive itself blocked. For safetey reasons, egress to exit the floor must always be maintained. Thus, by the nature of the floor layout, regardless of where an issue is, accessing the floor is inevitably likely to block not just storage cells, but also station queue and travel cells which are critical to keep clear to allow product to be delivered to associates. Somewhat ironically, issues in the middle of the storage area, which are often less impactful than issues in the travel lanes, require blocking the most resources to be resolved.

Thus, it is unclear that the current practice of recovering any and all issues as quickly as possible is the best approach in all cases. Perhaps the impact of some problems are negligible compared to the impact that recovering them would have and therefore, they should be addressed later at a less impactful time, such as during a break.

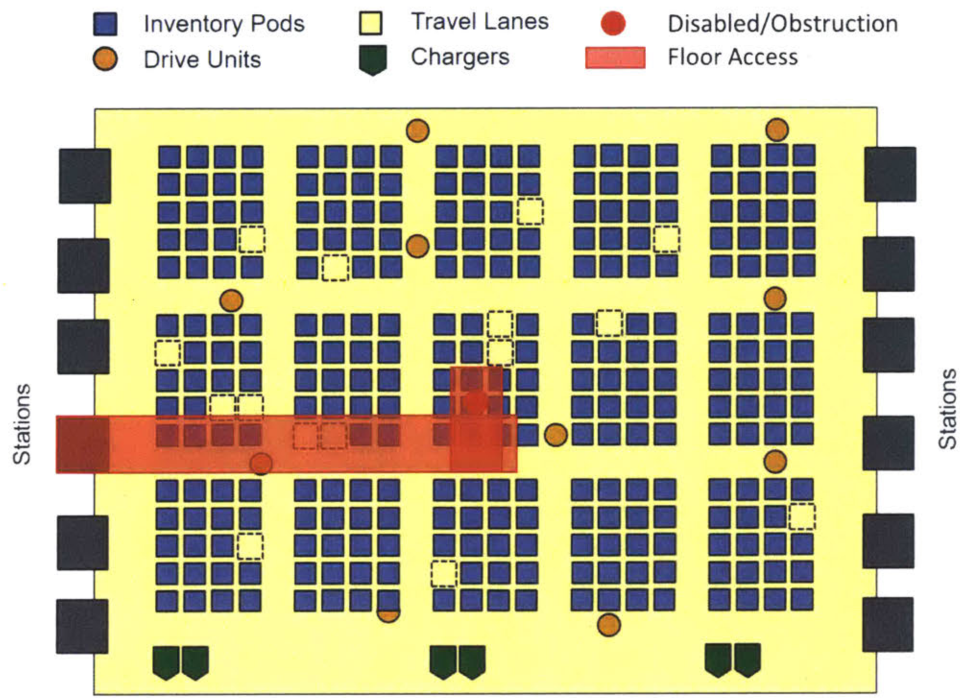

Figure 2.3 – Floor access to recover a disabled drive

Poor etiquette when accessing the floor was another common issue. Associates were sometimes observed to take very circuitous, inefficient paths when accessing the floor. It is difficult to navigate such a large area and the most efficient route is not often clear. Furthermore, other issues can pop-up on the floor in the same vicinity and it is easy to be tempted to add a quick detour to a current path to recover the additional issue as well. While good intentioned, this can really wreak havoc on the dynamics of the floor – creating congestion and rerouting.

Another common issue observed was a lack of awareness for what others were doing on the floor. For example, Figure 2.4 depicts a scenario where two associates monitoring the same floor accidentally cut the floor in half by access the floor on opposite sides and meeting in the middle. This creates a perfect storm, trapping drive units on either side of the barrier, creating massive congestion and rerouting, preventing them from getting to their destined stations, increasing idle time and reducing throughput.

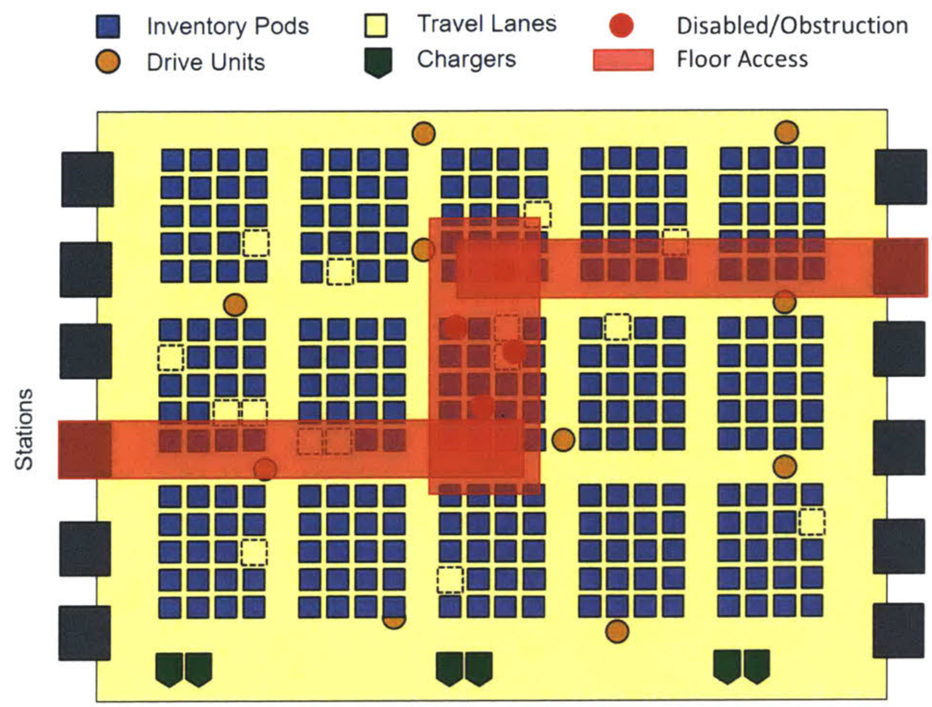

Figure 2.4 – Floor access cuts floor in half

Similarly, restricting too much of the floor at one time and forcing drives to reroute around these long blocked protrusions is also one leading cause of degraded floor performance. Figure 2.5 depicts a scenario where three associates access the floor at once.

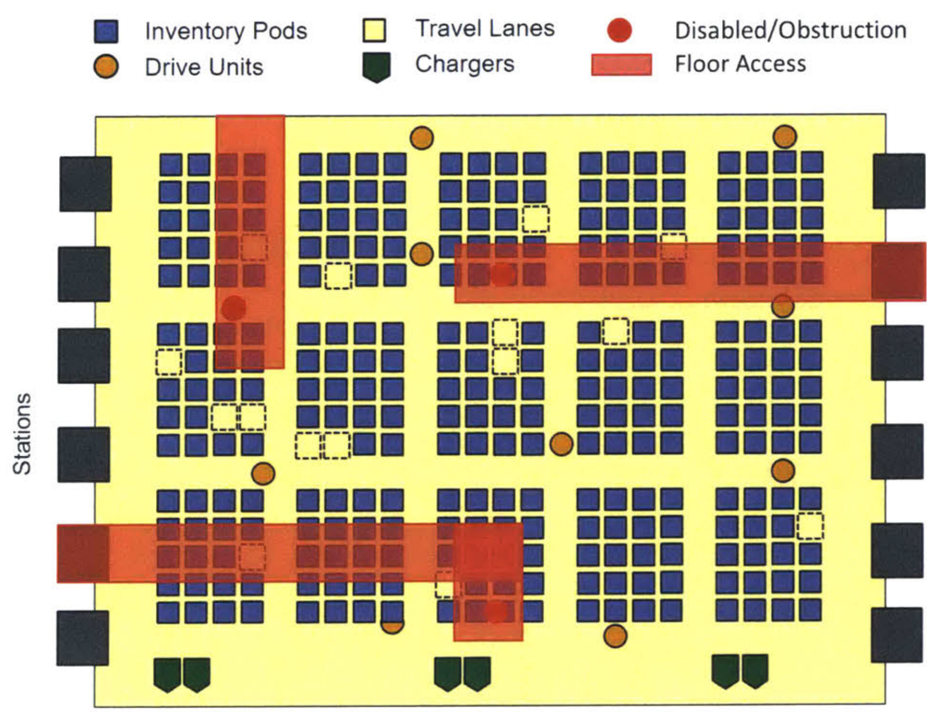

Figure 2.5 – Three associates access the floor at the same time

With such a vast area, these conditions are often not readily apparent to operations. It is easier to uncover these conditions after the performance impact to operations has been realized. Hence, there is a desire within the company to improve real-time visibility and build-in preventative measures into new or existing tools for monitoring floor conditions. These common issues highlight the impact of human variability in decision making and the importance of building the decision making into the tools as much as possible to help decision makers better optimize their actions.

3 STATISTICAL ANALYSIS & MODELING

3.1 Hypothesis

The research phase of the project revealed that the position of an issue on the floor likely affects its impact. For example, an obstruction in a travel path critical to getting to a station likely has a larger impact on floor performance than an obstruction in the middle of the remote storage area. Furthermore, accessing the floor to recover the obstruction in the remote storage area can be more impactful than quickly recovering the closer obstruction in the travel path. Yet, there is no data-driven information provided to Operations to quantify this spectrum of impact and inform decision making. The current policy is to treat these issues equally regardless of their location. The goal of this analysis was to show that the performance impact of floor health events could be correlated to the location characteristics of the events.

More specifically, one potential model was based on the idea that the impact to floor performance is proportional to the "value" of the floor area made inaccessible by poor floor health. And the "value" of that blocked floor area is related to its traffic demand, i.e. how often a drive unit needs to traverse that area. Under this theory, an obstruction in a station queue has a larger impact on floor performance than an obstruction in the middle of the storage area because the obstruction in the station queue is blocking a more "valuable" or high traffic position. The number of times a cell on the floor was traversed over a certain period of time can be extracted from production data. This data was assessed to calculate relative importance scores for each cell on a floor as a potential independent variable, as described in greater detail in Section 3.3.1. Another model was simply based on the types of cells (e.g. travel, storage, queue, etc.) that are blocked by poor floor health events and is described in greater detail in Section 3.3.2.

3.2 Simulations

Preliminary data analysis revealed the difficulty in isolating and quantifying the effects of individual floor health issues using historic production data given the complexity of dynamic factors on the floor. Thus, it was proposed that simulation, where certain factors can be kept constant and specific floor issues can be manipulated, would be the best way to isolate and control for the effects of floor health issues.

Simulations were conducted using an internally developed tool for running real-time simulations based on production data. By allowing full control over system parameters such as the floor map, station schedules and rates, and inventory, this simulation tool enabled A/B testing and is currently used by many internal teams to perform functional testing of new products, features, and alterations to the AR floor before roll-out to production. This simulation tool proved useful for this project in order to isolate and control for the effects of floor health issues.

3.2.1 Simulation Parameters

Simulations were run for four representative medium and large sortable type fulfillment centers. For each site, a control simulation was run with no health events to simulate an idealized, perfectly healthy floor. Then a battery of different health configurations were simulated – real floor health scenarios were randomly selected from the last six months of production data for these four fulfillment centers (production data is only readily accessible for the last six months). Historical data was assessed to determine the typical range of floor health issues to ensure the simulations captured the full spectrum of possible scenarios.

The objective of this statistical analysis effort was to identify the relationships between the characteristics of the health events and the change in floor performance caused by the health events as compared to the control. Anecdotal feedback obtained during the research phase of the project suggested that models should be linked to throughput rate, the primary performance metric, in order to garner traction and implementation support. Thus, the dependent variable collected was the percentage change in average throughput rate and a multitude of potential explanatory features were tracked in addition to the relative importance scores and cell types blocked, including:

- # of disabled drives
- # of obstructions
- # of floor access paths
- total # of concurrent issues
- # of restricted cells due to floor access
- # of travel cells blocked[1]
- # storage cells blocked
- # of station queue cells blocked
- # of pick cells blocked
- # of charger cells blocked
- total # of cells blocked
- total sum of relative importance scores of cells blocked (derived from traffic history as described in Section 3.3.1)

Statistical analysis was performed on the dataset collected from the simulations with the goal of determining the simplest and most effective model: one with high predictive power but only a few essential, interpretable variables since interpretability would be critical to getting buy-in from stakeholders, particularly the Operations leaders at the fulfillment centers.

3.3 Models

The simulations resulted in 300 usable data points for statistical analysis. About 80% of the data was used for training models and the other 20% for testing. Python via Jupyter notebooks was used for data pre-processing and formatting the simulation inputs and the software package R was primarily used for the statistical analysis. After investigating several pairwise and multivariate regression models, the two discussed in the subsequent sections emerged as the most promising.

3.3.1 Traffic History Model

As previously introduced, one potential model was based on the idea that the impact to floor performance is proportional to the "value" of the floor area made inaccessible by poor floor health where the "value" of that blocked floor area is related to its traffic demand. One measure of the relative importance of a position on the floor can be obtained by dividing a cell's traversal count by the count for the most traversed cell. To illustrate this methodology, example relative importance scores are listed in Table 3.1 for a particular AR floor in a fulfillment center. Using this method, the most traveled cell in the floor has a relative importance score of 1. In this example, the most valuable cell on this floor was traversed 9971 times by a drive unit so it has a relative importance score of 1. The relative importance score of the second most traversed cell is derived by 9166/9971 = 0.919266.

When there are several concurrent floor health issues on the floor, the value of the floor space affected by these issues at that moment in time can be estimated by summing the relative importance scores for all the blocked cells. For example, if the top three cells in this scenario were blocked due to poor floor health, the total score would be the sum of these three: 1 + 0.919266 + 0.877846 = 2.797112.

Using the simulation data for one of the fulfillment centers, Figure 3.1 shows this total score on the x-axis versus the drop in throughput rate on the y-axis. This total sum of the relative importance scores derived from the drive traffic history was highly correlated to the degradation in throughput rate with a correlation coefficient of 0.80.

Table 3.1 – Example top 10 most important cells

Ranking	Traversal Count	Relative Importance Score
1	9971	1
2	9166	0.919266
3	8753	0.877846
4	8467	0.849163
5	8248	0.827199
6	8088	0.811152
7	8088	0.811152
8	8080	0.81035
9	8004	0.802728
10	7858	0.788085

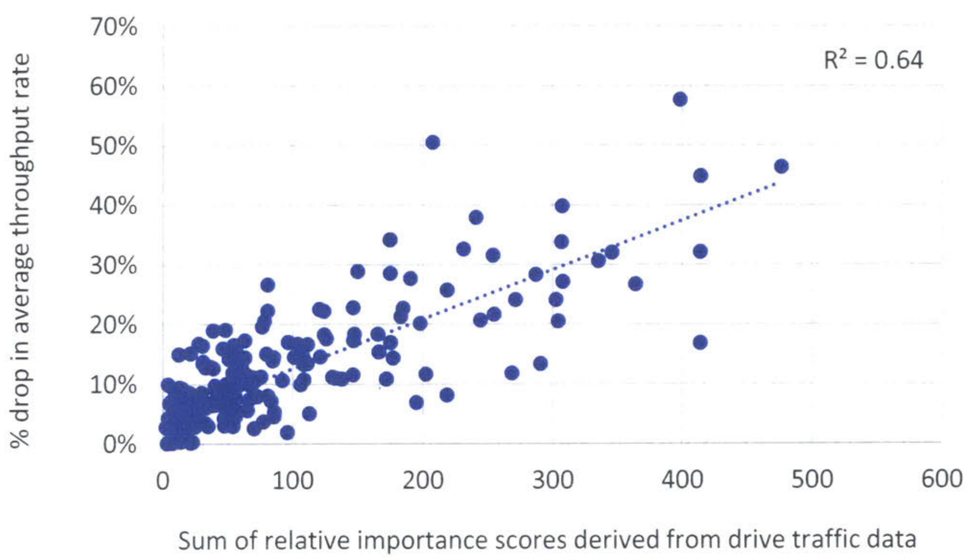

Figure 3.1 – Model based on the value of blocked floor area derived from traffic history data

This analysis also revealed some interesting insights about the most traveled cells. The most traveled cells are largely driven by the floor configurations. In Figure 3.2 below, the areas shaded in purple are the stations along the perimeter of the floor while the yellow and green shaded areas designate travel ("highway") and storage cells, respectively. The arrows indicate the direction of travel permitted through that cell. The teal square highlights the top five most traveled cells in this particular scenario. All of the top five most traveled cells are in the highway and at intersections of major travel lanes. Furthermore, the third, fourth and fifth most traveled cells are all in travel lanes that lead directly into or out of station queues. The first and second most travelled cells do not conform this pattern but closer examination reveals that the travel directions permitted by the surrounding cells force traffic to converge at these points. The least travelled cells, not depicted here, are often at the edges and corners of the map, highway cells between the enter and exit queues of stations, and stations queues not being used.

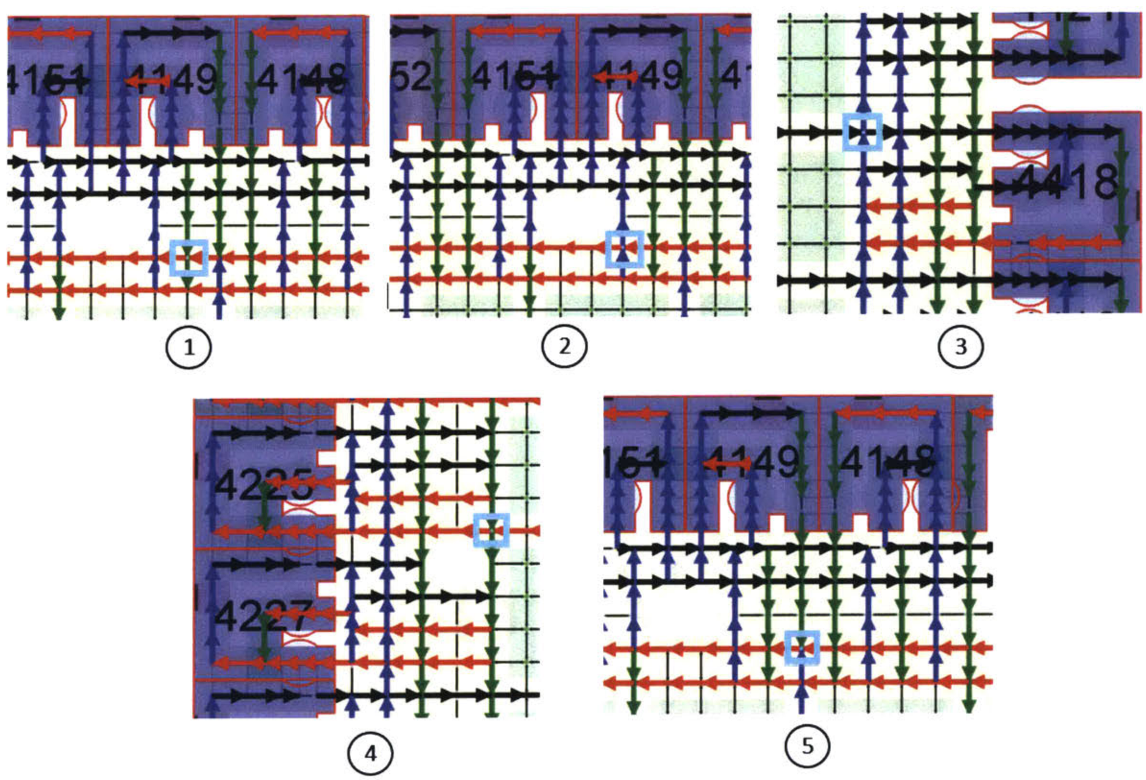

Figure 3.2 – Example top 5 most traversed cells for an AR floor in a fulfillment center

3.3.2 Cell Type Model

The second best predictors of the change in throughput rate were the types of cells being blocked. First, the pairwise correlations were investigated, followed by multiple regression.

In particular, the percentage of travel cells blocked (correlation of 0.77), percentage of queue cells blocked (correlation of 0.80), and percentage of pick cells blocked (correlation of 0.79) are the most important cell types, as shown in Figures 3.3 – 3.5. The percentage of cells blocked was used to normalize to the size of the floor. The different colors represent the different fulfillment centers.

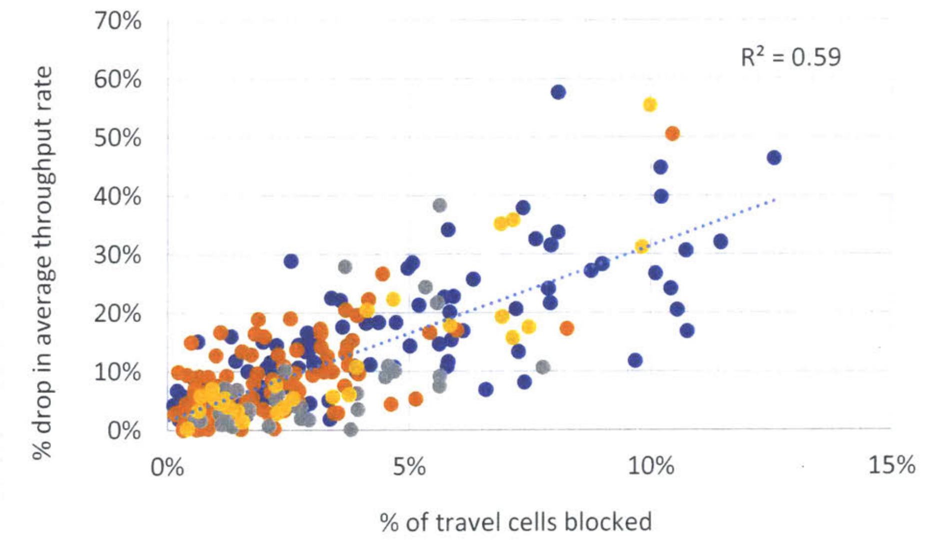

Figure 3.3 – Percentage of travel fiducials blocked versus change in throughput rate

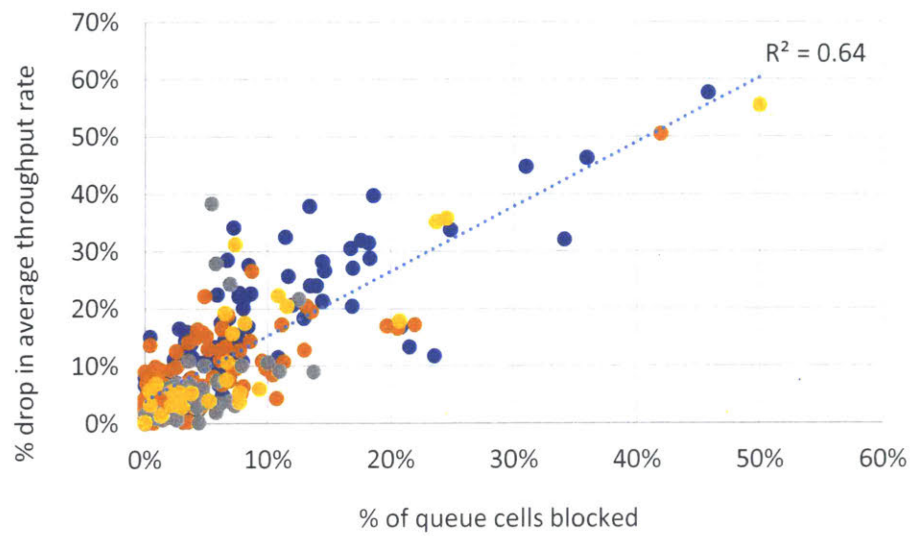

Figure 3.4 – Percentage of station queue cells blocked versus change in throughput rate

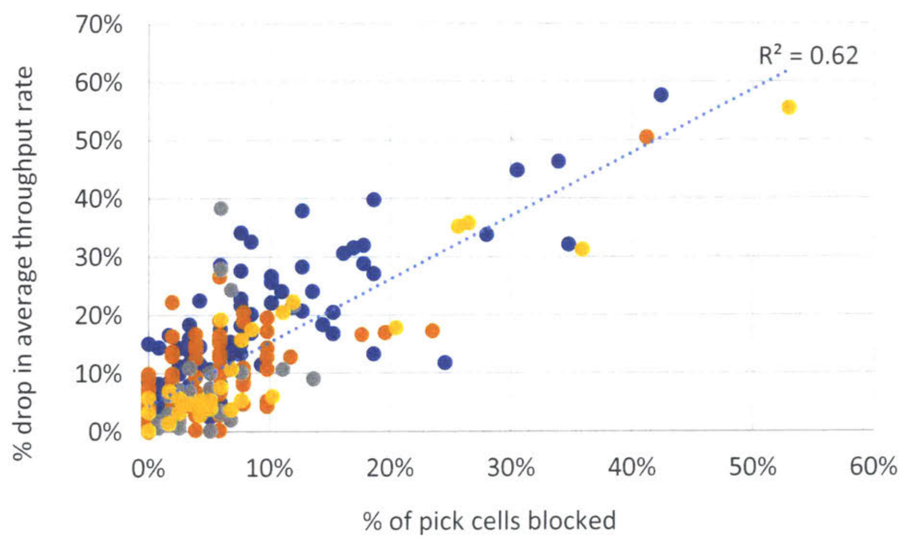

Figure 3.5 – Percentage of pick cells blocked versus change in throughput rate

A linear regression model including the four cell types most correlated to the change in pick rate has high goodness-of-fit. However, some variables do not have significant p-values. Variables with the largest p-value were removed so that only the statistically significant variables remained, as shown in Figures 3.6 – 3.7 below.

```
Call:
lm(formula = ratechange ~ travelperc + queueperc + storageperc +
    pickperc)

Residuals:
      Min        1Q    Median        3Q       Max
-0.206208 -0.031620 -0.005341  0.031466  0.242658

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)  0.02023    0.00556   3.638 0.000330 ***
travelperc   1.00118    0.43436   2.305 0.021938 *
queueperc    0.58374    0.17510   3.334 0.000978 ***
storageperc  0.66342    0.64401   1.030 0.303874
pickperc     0.20643    0.16517   1.250 0.212467
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.05797 on 267 degrees of freedom
Multiple R-squared:  0.693,     Adjusted R-squared:  0.6884
F-statistic: 150.7 on 4 and 267 DF,  p-value: < 2.2e-16
```

Figure 3.6 – Statiscal analysis in R Studio for model based on four cell types

```
Call:
lm(formula = ratechange ~ travelperc + queueperc)

Residuals:
      Min        1Q    Median        3Q       Max
-0.209666 -0.033495 -0.004596  0.030605  0.243456

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept) 0.021010   0.005439   3.863 0.000141 ***
travelperc  1.423136   0.214544   6.633  1.8e-10 ***
queueperc   0.715850   0.078107   9.165  < 2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.05801 on 269 degrees of freedom
Multiple R-squared:  0.6903,    Adjusted R-squared:  0.688
F-statistic: 299.8 on 2 and 269 DF,  p-value: < 2.2e-16
```

```
> vif(lm(ratechange~travelperc+queueperc))
travelperc  queueperc
  2.673957   2.673957
```

Figure 3.7 – Statiscal analysis in R Studio for model based on two cell types

The insignificant p-values for storage and pick is partly explained by high correlation between cell types. The correlation matrix below shows that the percentage of pick cells blocked is highly correlated to the percentage of queue cells blocked and similarly for the travel and storage types. This makes sense given the inherent layout of the floor. Since pick cells are surrounded by queue cells near the perimeter of the floor, a floor access path that requires blocking a pick cell is inevitably likely to also block queue cells. Similarly, travel and storage cells are located in close proxmity more to each other in the center part of the floor. Accessing a storage cell necessitates blocking a travel cell.

Multicollinearity can assessed by computing a score called the variance inflation factor (VIF), which measures how much the variance of a regression coefficient is inflated due to multicollinearity in the model. The smallest possible value of VIF is one (absence of multicollinearity). A general rule of thumb is that a VIF value exceeding 5 warrants further investigation and a VIF value exceeding 10 indicates a problematic amount of collinearity.

Thus, a model including just the percentage of travel and queue fiducials blocked was chosen to mitigate the effects of multicollinearity while maintaining high goodness of fit, shown in Figure 3.8. A model simply based on the total percentage of all cells on a floor that are blocked, regardless of type, is still a relatively strong indicator of potential degradation in throughput rate and is shown in Figure 3.9 for comparison below. Additional models less sensitive to multicollinearity, such as Ridge and Lasso regressions, were also investigated. However, the linear regression model performed better and was more intuitively explainable.

Table 3.2 – Correlation matrix for cell types

	Rate Change	**Total Blocked**	**Travel**	**Storage**	**Queue**	**Pick**
Rate Change	1.00					
Total Blocked	0.78	1.00				
Travel	0.77	0.99	1.00			
Storage	0.56	0.87	0.84	1.00		
Queue	0.80	0.81	0.69	0.47	1.00	
Pick	0.79	0.80	0.78	0.46	0.96	1.00

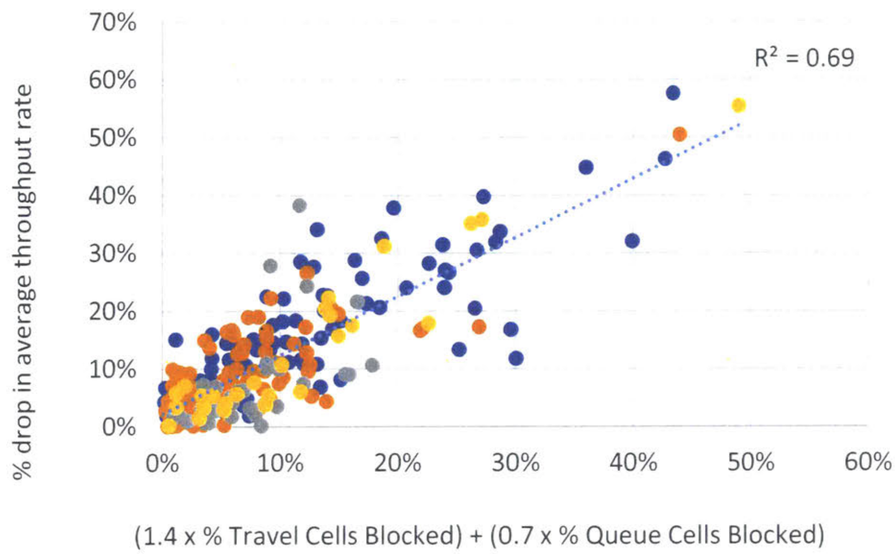

Figure 3.8 – Model based on percentage of travel and queue cells blocked

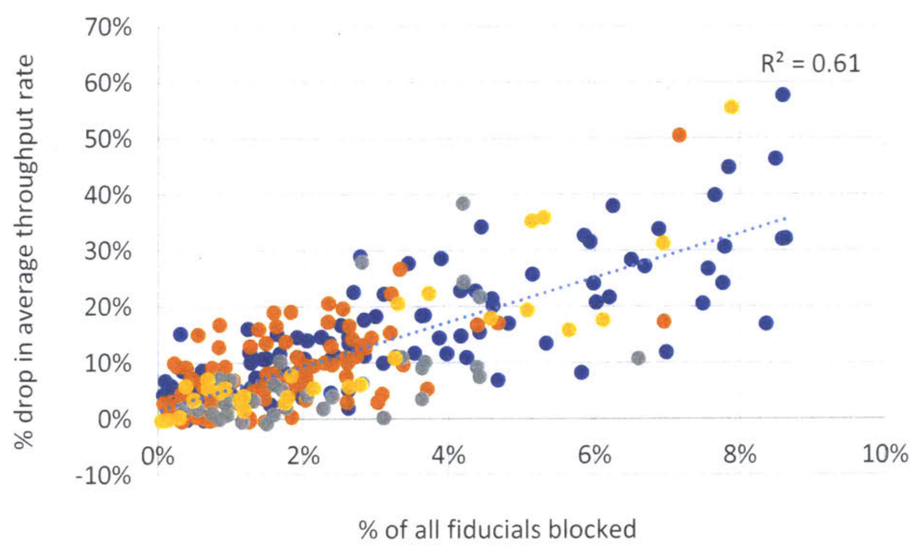

Figure 3.9 – Model based on percentage of all cells blocked

3.4 Model Limitations

While simulations allowed for a programmatic way to isolate and control for the effects of floor health issues, there are a few main limitations of this approach worth noting: limited availability of training data, generalizability to non-sortable fulfillment center types, and applicability to new technologies and adjustments to system algorithms.

Simulations were conducted in real-time, limiting the amount of training data that could be collected during the span of the project. Furthermore, all simulations start from the beginning of a shift such that all drives originate from the storage area. Consequently, there is some ramp-up time before the floor simulator achieves steady-state. Thus, it was recommended that a 3-hour simulation would be the minimum run-time needed to obtain meaningful data. Moreover, part of the power of the tool is that it is connected to live services. While this lends to the accuracy of the tool, it also means that any time a live service experiences an outage, the tool also goes out of service. Consequently, out of approximately 600 simulation runs, only 300 usable data points were collected. There were not enough data points to enable usage of more sophisticated machine learning techniques.

Given limited resources and time, scope was narrowed to collect simulation data for the sortable building type only, since these are the most highly automated building type, the most common building type currently in Amazon's network, more representative of the fulfillment centers that Amazon plans to build in the future, and often the most susceptible to poor floor health. Initially, simulation data was collected primarily for the most representative large fulfillment center. To avoid over fitting site-specific conditions such as staffing configuration and floor map nuances, simulations were extended to three other representative sortable buildings. However, there is always a risk that the model is overfit to these specific sites. Future work could include collecting more training data for additional sortable buildings, as well as applying the same methodology established during this project to develop a model for non-sortable building types.

AR is constantly developing and introducing new technologies and updates to the software programs and algorithms which often affect the system behavior. Thus, it is important to note that the coefficients of the model determined in this project are not static. Future work could include implementing machine learning techniques to dynamically update the model using real-time data.

3.5 Model Selection

Considering the implementation difficulty of each model was important during the model selection process. The larger goal of the project was to start developing additional features or entirely new tools to help operations have better visibility into poor floor health. While the drive traffic data proved valuable in developing a model based on the relative importance of cells blocked, using the drive traffic data as an input would be challenging when developing a usable product. The drive traffic data is currently only stored for the purposes of visualizing a heat map. Significant manual effort was required to process the data for the purposes of this analysis. Additional infrastructure would be required to start collecting and storing this data in real-time.

Moreover, the model that includes drive traffic data as an input does not perform substantially better than the model simply based on the types of cells blocked. Furthermore, drive traffic patterns are highly dynamic and subject to changes in staffing configurations, product placement, floor layout, etc. It is expected that a model based on drive traffic data is less generalizable and would need to be updated and tailored to specific site conditions more frequently than a model based on cell types.

Thus, for the goals of this project, it was determined that the additional complexity of implementing a model based on the drive traffic data was not worth the marginal gain in model performance.

4 USE-CASES

4.1 Opportunity Space

There are a wide range of impactful applications of the earlier findings, including but not limited to the following:

- **Diagnostic Check**

  The project team was already working on a diagnostic tool for the technical support team responsible for responding to and root-causing issues reported by the fulfillment centers. The tool runs a series of broad-spectrum diagnostics to help quickly narrow down possible root causes but there is no check for floor health yet. A check could be added based on the model of cell types that are blocked to help identify which areas of the floor to focus on to mitigate the issue. This additional check could improve escalation response time for floor health related events.

- **Automated Notifications of Poor Floor Conditions**

  A new floor health score (e.g. score = 1.4 Travel + 0.7 Queue) could be tracked and floor health events whose score exceeds a certain threshold could prompt automated alerts to notify the AR technical support team, on-site operations managers, or the front-line associates of deteriorating floor conditions. Automated alerts could be implemented through a variety of existing alert systems. Automated notifications for other operational problems are already implemented, but not yet for floor health.

  Analysis of historic high severity events indicate that in most cases deteriorating conditions could have been detected more than an hour in advance of on-site teams informing AR of the issue. Auto-notifications would allow the AR technical support team to proactively engage with on-site operations teams to identify impending issues before they become high severity events, thereby potentially reducing impact to performance and the number of escalations.

- **Track and Incentivize Better Floor Access Practices**

  As previously described, poor etiquette when accessing the floor to recover an issue can be more impactful and disruptive to operations than the initial issue itself. The average score for user areas could be tracked at the individual associate level, or for each shift, or on a site-level basis. This could provide insight into which individuals, shifts, and sites, might benefit from additional coaching in floor access etiquette and best practices.

No metrics currently incentivize careful floor access planning and adherence to best practices. Dwell time based metrics currently encourage getting on and off the floor as fast as possible but do not recognize the performance impact that different routes can have.

- **Automated and Optimized Floor Access Path Planning**

  Even better than just tracking the efficiency of a user's floor access would be to automatically provide associates the optimal, least impactful path to recover a floor issue, removing the guess work and variability in human decision making. This ideal application is obviously very complex. More feasible, shorter-term implementations could include automated notifications when a user tries to access the floor with a path that exceeds a certain threshold score.

- **Intuitive and Actionable Floor Health Metric in Roboscout**

  These findings can also be used to improve the interpretability of existing floor health metrics. Displaying the total percentage of a floor's cells that are inaccessible and the specific types of cells that are blocked, which have now been shown to be highly correlated to potential degradation in performance, can help on-site teams better focus their efforts.

- **Prioritization of Floor Health Issues**

  The current practice is to address any and all issues on the AR floor as quickly as possible and dwell time metrics reinforce this behavior. Associates address issues based on general rules of thumb. They often weigh whether or not it is a blocking or not blocking access to a station, the issue's dwell time, and the proximity of the issue to their current location, as they decide which floor issue to clear next. The developed model based on cell types blocked could be used to provide more specific guidance on the priority of issues. The score could be calculated for each issue on the floor in a moment in time, providing an apples-to-apples comparison between all the issues on the floor where the highest scoring issues should have the highest priority as they have the most potential to negatively impact performance. Furthermore, a new motivating metric could be the number of high priority issues or the cumulative score of the issues cleared in a shift.

The next sections describe in more detail the use-cases that were deeply explored in the implementation phase of this project.

4.2 Automated Notifications of Poor Floor Conditions

The development of automated notifications was identified as one of the target implementations for the project given that it could be built-off the existing services within the project team. While automated notifications could be directed to the AR technical support team, the longer-term goal was to identify the most efficient and effective means of directly engaging the appropriate on-site operations teams. The proposed initial auto-cutting scheme consists of tracking the score in real-time, ideally on at least a 5-minute basis. Based on the model findings and analysis of historic high severity issues, an initial hypothesis for an effective alerting threshold structure would consist of two trigger points:

1) if the score exceeds a threshold of 0.02 for three readings over a 20-minute period
2) if the score exceeds a threshold of 0.05 for two readings over a 10-minute period

This structure aims to detect sustained, unintended problems rather than one-time issues or controlled conditions like maintenance windows or station cleanings. Figure 4.1 illustrates how this alerting scheme would have worked on an issue reported from earlier this year and Appendix A illustrates examples for nine other issues reported this year.

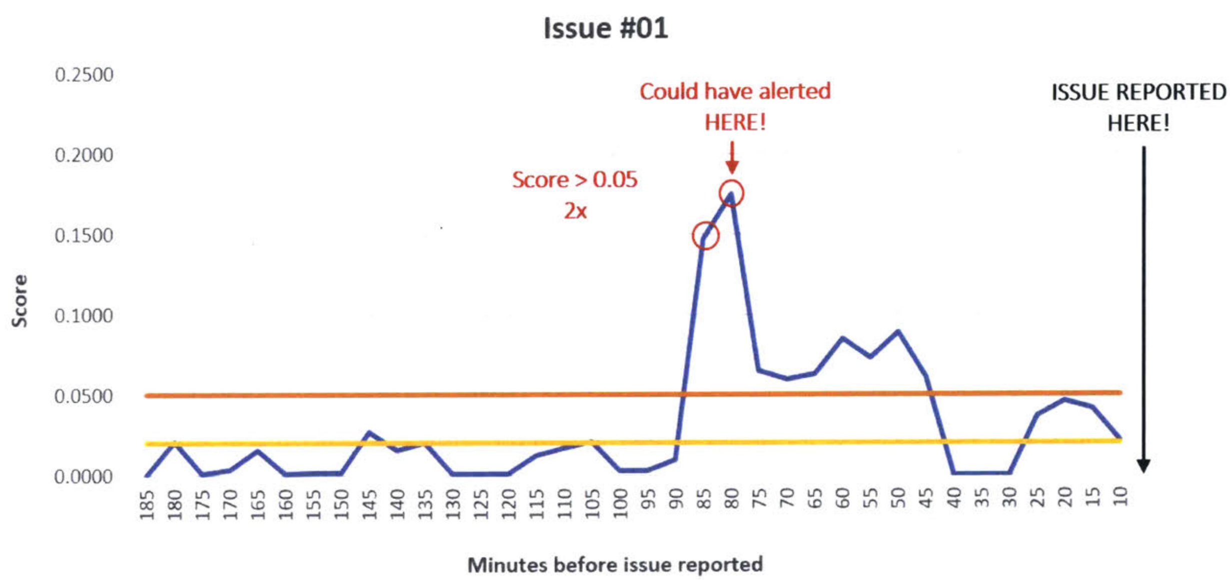

Figure 4.1 - Score three hours leading up to the report of an issue

4.3 Task Prioritization

Another implementation that was pursued in the final phase of the project was using the developed floor health model to inform the priority ranking of floor health issues in the kindle-compatible visualization tool used network-wide to manage and monitor AR floors. The floor health model, simply based on the percentage of travel and station queue cells blocked provides an apples-to-apples score by which to prioritize tasks. By highlighting the most impactful issues on the floor, prioritizing floor issues is expected to more efficiently mitigate poor floor conditions, thereby reducing high severity events. Prioritizing tasks in the kindle-compatible monitoring application is a compelling opportunity as it directly engages in real-time the front-line associates who interact with the floor. Moreover, it does so with minimal disruption to their established work flows by integrating with their existing primary tools. The existing infrastructure for automated notifications necessitated use of a laptop or access to email, neither of which are convenient for the target user.

Development work for the task prioritization use-case included creating additional APIs to allow various live services to communicate and transfer data. A flow diagram illustrating the complexity of calls between these services to enable the prioritization is shown in Appendix B. Actual product and service names are anonymized to protect proprietary information.

A pilot at a fulfillment center in partnership with the processing engineering teams to test the new prioritization scheme is recommended to evaluate 1) the performance of the new prioritization logic and 2) the user adoption of the feature and adherence to completing the tasks in priority order, before launching to the broader network. The pilot fulfillment center should be conducive to conducting an A/B test. For example, a site with floors identical in layout would facilitate an A/B test such that the new prioritization logic could be implemented on one floor while maintaining the existing system on other floor.

This application was also compelling to implement because there will be no perceptible difference to the end user since the prioritization logic is abstracted from the user on the backend. So additional training and onboarding would not be needed to enable an A/B test on different floors but allows comparison of the prioritization methods.

4.4 Diagnostic Tool

During code freeze when the prioritization application was put on hold, a python script consisting of SQL queries was developed to analyze historic floor health data (almost in real-time: 5 to 20 minutes delayed). This script collects the x,y location features of obstructions, disabled drives, and user areas from the Amazon database and determines the number and type of cells affected by these events by merging with facility data. This script was used to investigate if the proposed model based on cell types blocked could help root cause some the unusual spikes in idle time metrics observed during Peak 2018. Appendix C provides additional details on this analysis.

In summary, analysis revealed that there is indeed a relationship between the idle time metric and the proposed score based on cells blocked. Spikes in the score are often accompanied by spikes in the idle time metric. However, there are some instances where they do not align. Investigation into cases of high model scores but low idle time often revealed a cleaning blitz or station cleanings where large swaths of the floor are blocked. This is a reminder that controls will be needed to mitigate for false alerts in the automated notification use-case.

In the reverse case, when idle times were high but the score was low, the interpretation was that some of the idle time was not driven by poor floor health. Investigation of idle time root causes data revealed that root causes not related to floor health such as ramp-up, takt overestimation, and low work target, were particularly high in these cases.

To further investigate the hypothesis that the score should be more correlated to idle time driven by floor health, the idle time metric was broken down by root causes that are often associated with floor health and those considered not related to floor health. As shown in Figure 4.2 below, there is a stronger correlation between the score and floor health driven idle time than the score and the non-floor health driven idle time as expected.

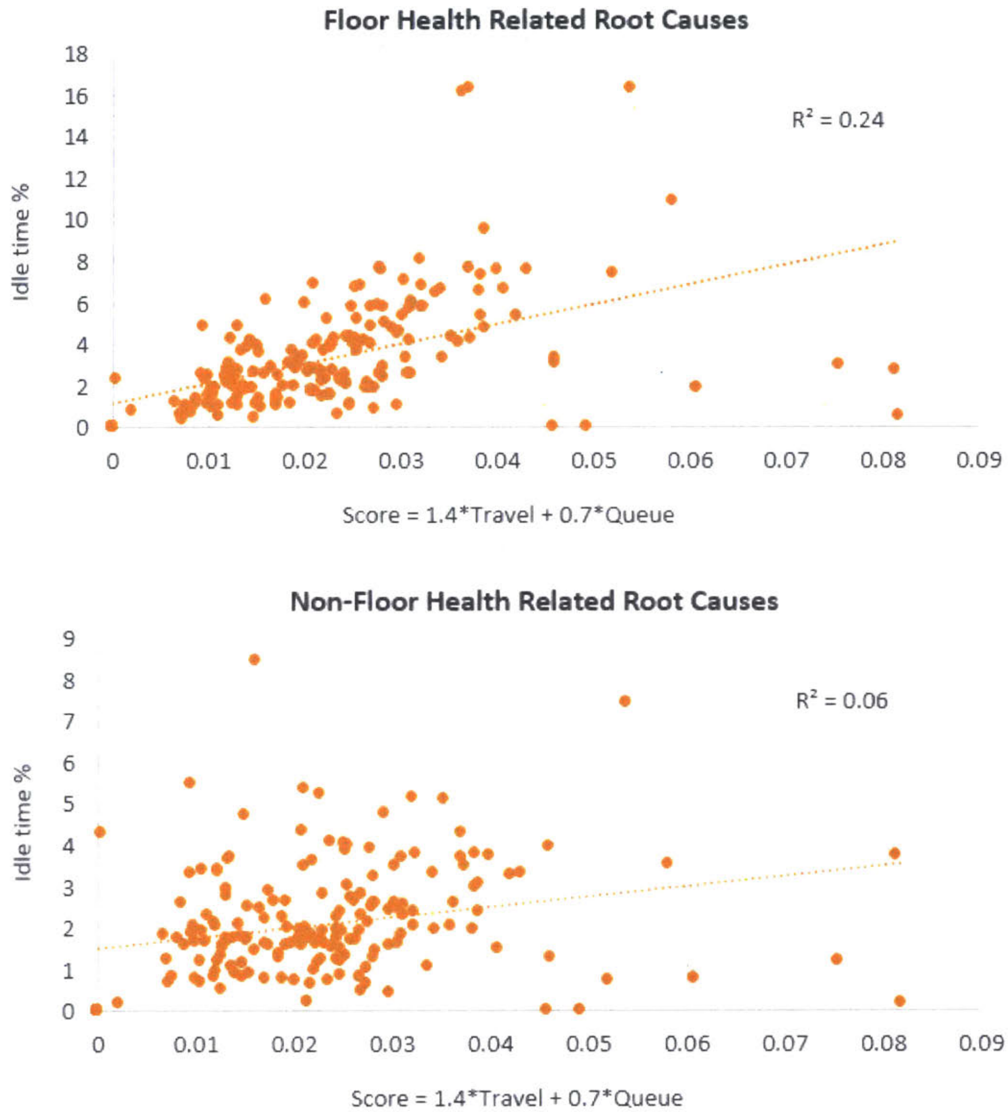

Figure 4.2 – Score is more correlated to floor health related idle time root causes

5 CONCLUSION

In summary, this project resulted in a variety of new processes to assess floor health. The discovery that the types of cells blocked is highly correlated to degradation in performance has powerful implications. This model can be used to create floor health troubleshooting tools, automated notifications of deteriorating conditions, better informed floor access management, a more intuitive and actionable floor health metric for operations, and the prioritization of issues.

In addition to pursuing these uses-cases, additional strategic recommendations were provided based on observations and information gathered throughout the project. These recommendations aim to improve user adoption of new AR tools and technologies, consolidate the multitude of disparate tools and dashboards, and align metrics and incentive structures to encourage desired behavior.

- **User Adoption**

  It is recommended that AR establish processes and procedures for rolling-out new technologies as well as processes on deprecating old tools. It is also recommended that AR invest in collecting simple user engagement metrics, such as number of page views, average session time, average session frequency, and screen flow, to better assess user adoption rather than relying on self-reported surveys where there are incentives to provide desired responses.

- **Analysis Paralysis**

  Consolidating all metrics and root causing tools into a centralized location, such as FC Console, Roboscout, or Tableau, could improve the efficiency in which on-site teams monitor and react to system conditions. It is recommended that new ideas be integrated into commonly used existing dashboards or services to the extent possible before creating new ones.

- **Misaligned Metrics and Incentive Structures**

  In the long-term, it is recommended that AR reevaluate some of its metrics and assess the intended and unintended incentives they create. For example, rather than measuring the number of blocking andons exceeding 10 mins and the number of disabled drives and obstructions exceeding 15 mins, a more motivating metric aligned with the model findings might be the number of high priority andons or the cumulative score of the andons an AFM clears in a shift.

In conclusion, significant progress was made in the development and implementation of automated alerts that proactively identify deteriorating floor conditions and real-time prioritization of floor health issues. A diagnostic tool in the form of a script of SQL queries was delivered to help more quickly narrow down if poor floor health is a root cause of a performance problem. These real-time, preventative measures are expected to be particularly useful for new sites experiencing the growing pains of increased floor health issues during ramp-up, allowing Amazon to scale more effectively.

Appendix A: Deep Dive into 2018 Floor Health Related Issue Tickets

Deep dives were performed on several high severity issue tickets from 2018 to investigate the potential efficacy of the initial proposed score/threshold alerting scheme. The following pages show a time series of the conditions leading up to the report of the issue by operations and illustrate where alarms could have been triggered in advance. In all 8 cases evaluated, deteriorating conditions could have been detected on average 61.8 minutes in advance of on-site teams informing AR of the issue.

Ticket ID	**Alert triggered?**	**Minutes in advance**
01	Yes	80
02	Yes	92
03	Yes	96
04	Yes	83
05	Yes	122
06	Yes	61
07	Yes	58
08	Yes	26
	Average	**61.8**

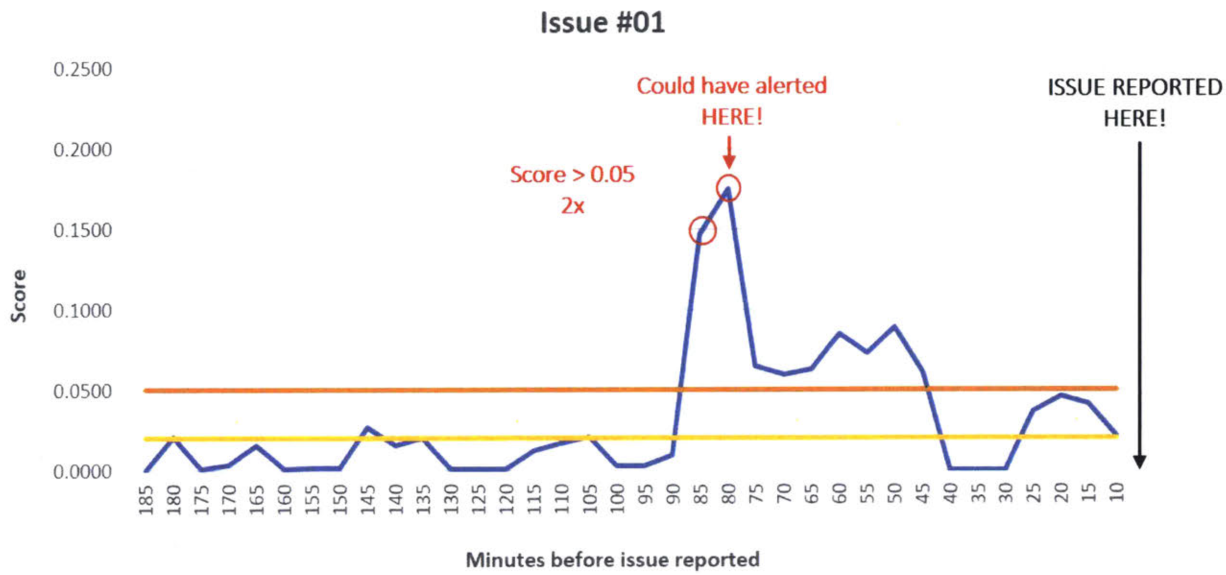
Issue #01
Score
0.2500
0.2000
0.1500
0.1000
0.0500
0.0000
Could have alerted
HERE!
Score > 0.05
2x
ISSUE REPORTED
HERE!
185 180 175 170 165 160 155 150 145 140 135 130 125 120 115 110 105 100 95 90 85 80 75 70 65 60 55 50 45 40 35 30 25 20 15 10
Minutes before issue reported

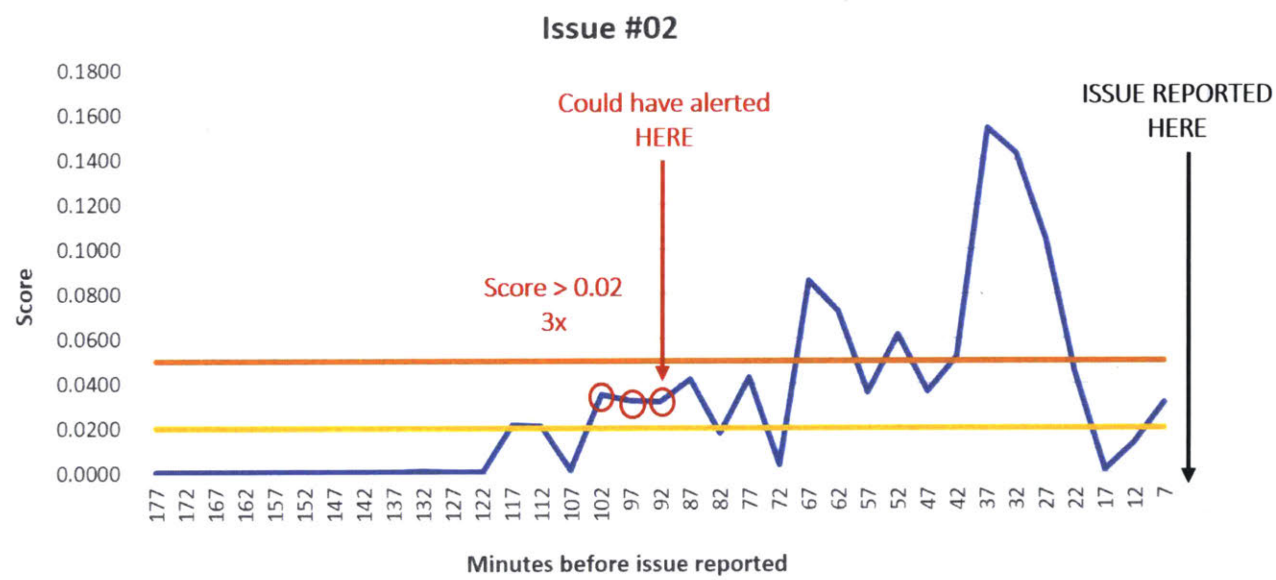
Issue #02
Score
0.1800
0.1600
0.1400
0.1200
0.1000
0.0800
0.0600
0.0400
0.0200
0.0000
Could have alerted
HERE
Score > 0.02
3x
ISSUE REPORTED
HERE
177 172 167 162 157 152 147 142 137 132 127 122 117 112 107 102 97 92 87 82 77 72 67 62 57 52 47 42 37 32 27 22 17 12 7
Minutes before issue reported

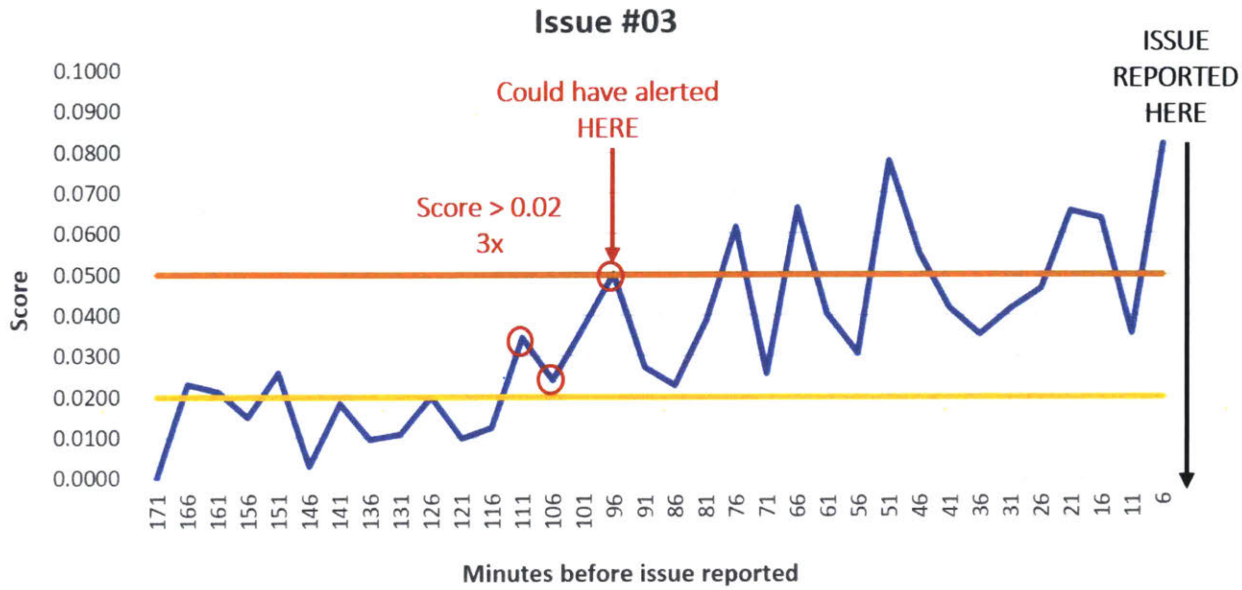
Issue #03
ISSUE REPORTED HERE
Could have alerted HERE
Score > 0.02 3x
Score
0.1000
0.0900
0.0800
0.0700
0.0600
0.0500
0.0400
0.0300
0.0200
0.0100
0.0000
171 166 161 156 151 146 141 136 131 126 121 116 111 106 101 96 91 86 81 76 71 66 61 56 51 46 41 36 31 26 21 16 11 6
Minutes before issue reported

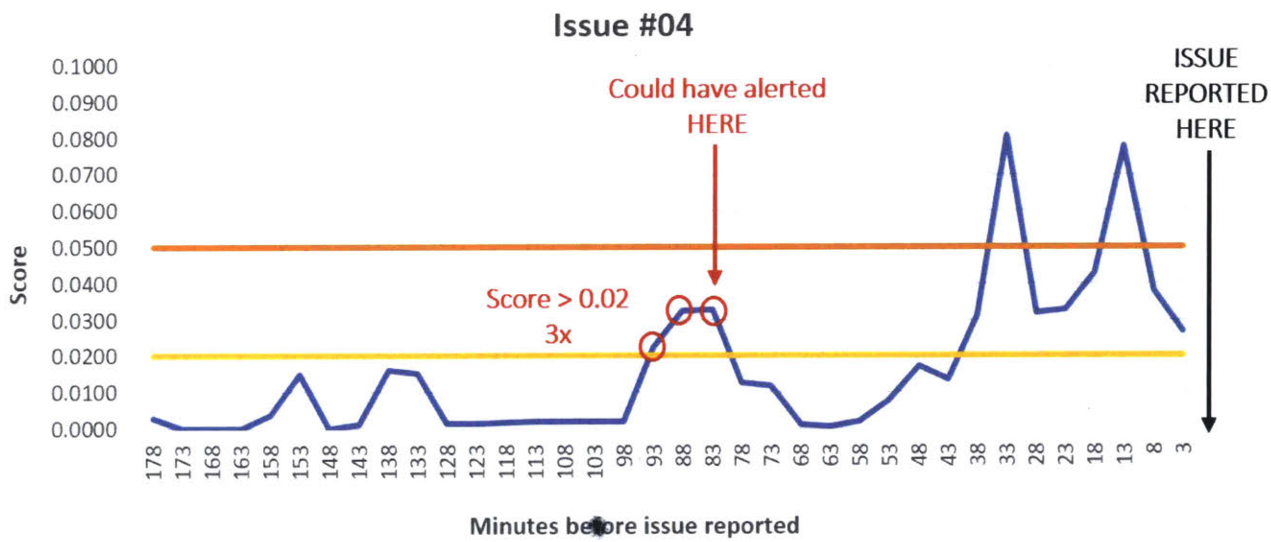
Issue #04
ISSUE REPORTED HERE
Could have alerted HERE
Score > 0.02 3x
Score
0.1000
0.0900
0.0800
0.0700
0.0600
0.0500
0.0400
0.0300
0.0200
0.0100
0.0000
178 173 168 163 158 153 148 143 138 133 128 123 118 113 108 103 98 93 88 83 78 73 68 63 58 53 48 43 38 33 28 23 18 13 8 3
Minutes before issue reported

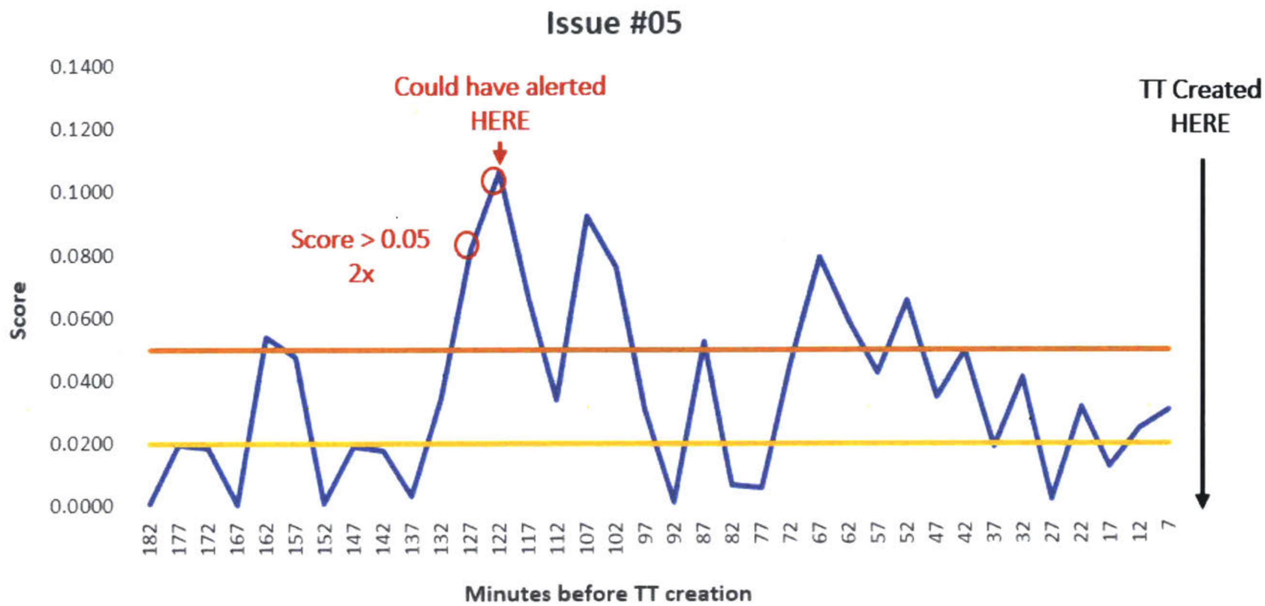
Issue #05
Could have alerted HERE
Score > 0.05 2x
TT Created HERE
Score
Minutes before TT creation
0.1400
0.1200
0.1000
0.0800
0.0600
0.0400
0.0200
0.0000
182 177 172 167 162 157 152 147 142 137 132 127 122 117 112 107 102 97 92 87 82 77 72 67 62 57 52 47 42 37 32 27 22 17 12 7

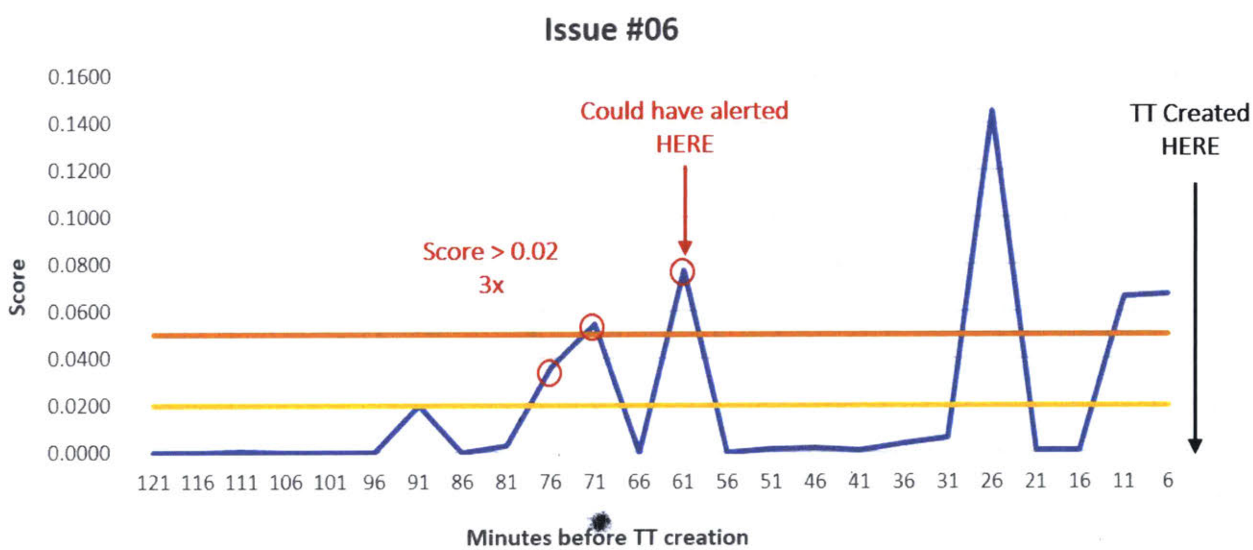
Issue #06
Could have alerted HERE
Score > 0.02 3x
TT Created HERE
Score
Minutes before TT creation
0.1600
0.1400
0.1200
0.1000
0.0800
0.0600
0.0400
0.0200
0.0000
121 116 111 106 101 96 91 86 81 76 71 66 61 56 51 46 41 36 31 26 21 16 11 6

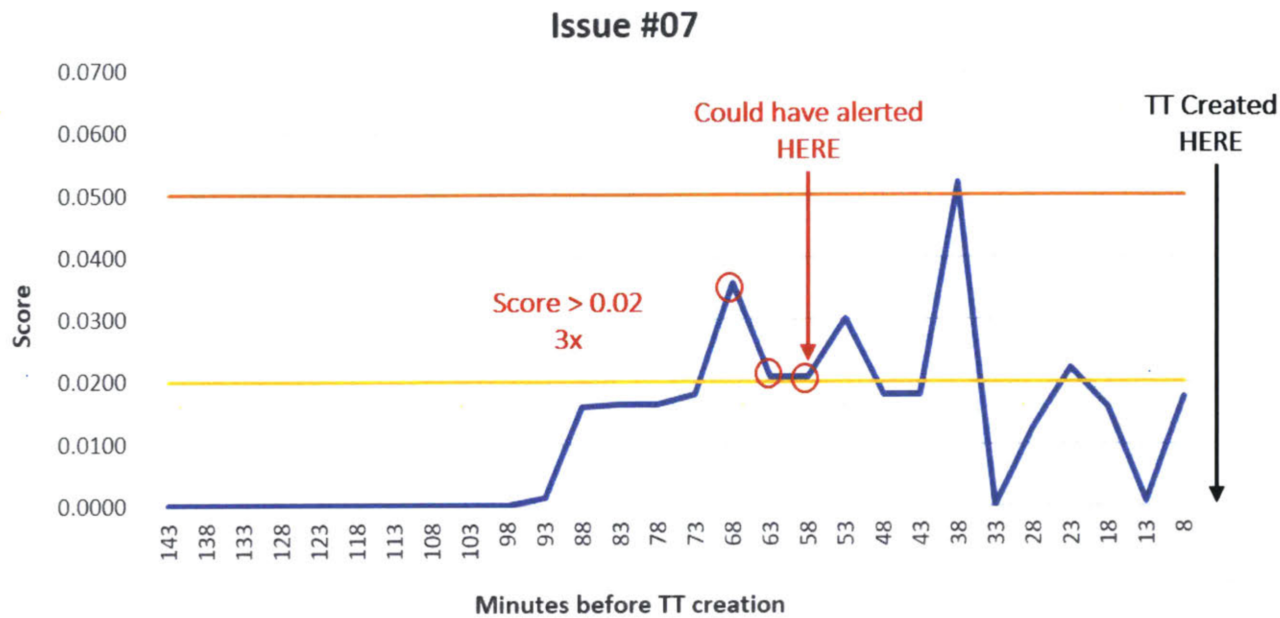
Issue #07
0.0700
0.0600
0.0500
0.0400
0.0300
0.0200
0.0100
0.0000
Score
Could have alerted
HERE
TT Created
HERE
Score > 0.02
3x
143 138 133 128 123 118 113 108 103 98 93 88 83 78 73 68 63 58 53 48 43 38 33 28 23 18 13 8
Minutes before TT creation

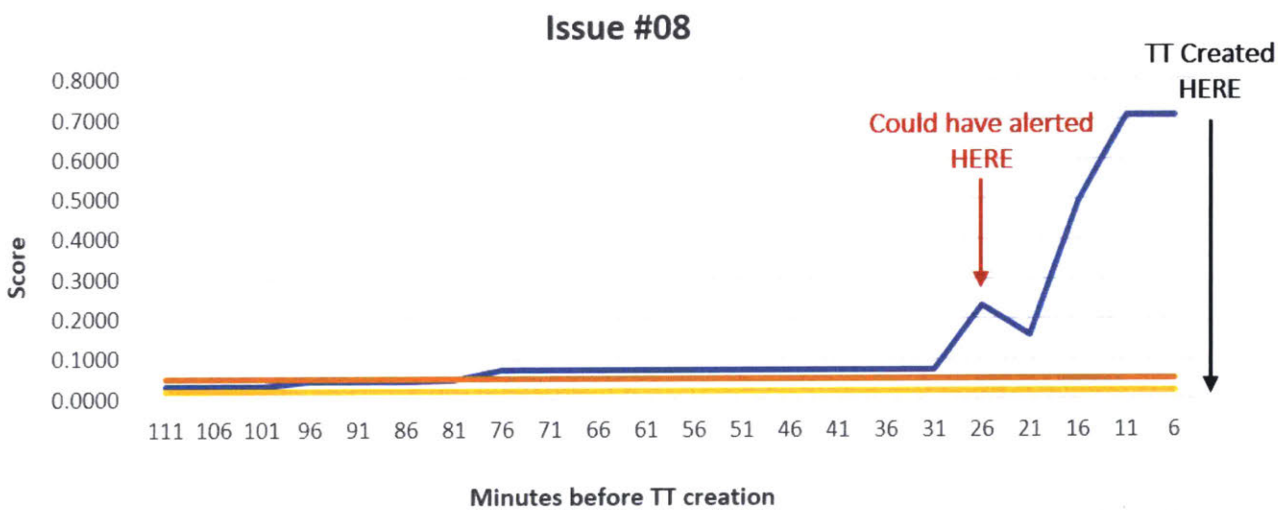
Issue #08
TT Created
HERE
0.8000
0.7000
0.6000
0.5000
0.4000
0.3000
0.2000
0.1000
0.0000
Score
Could have alerted
HERE
111 106 101 96 91 86 81 76 71 66 61 56 51 46 41 36 31 26 21 16 11 6
Minutes before TT creation

Appendix B: Flow Diagram of Service Calls for Task Prioritization Application

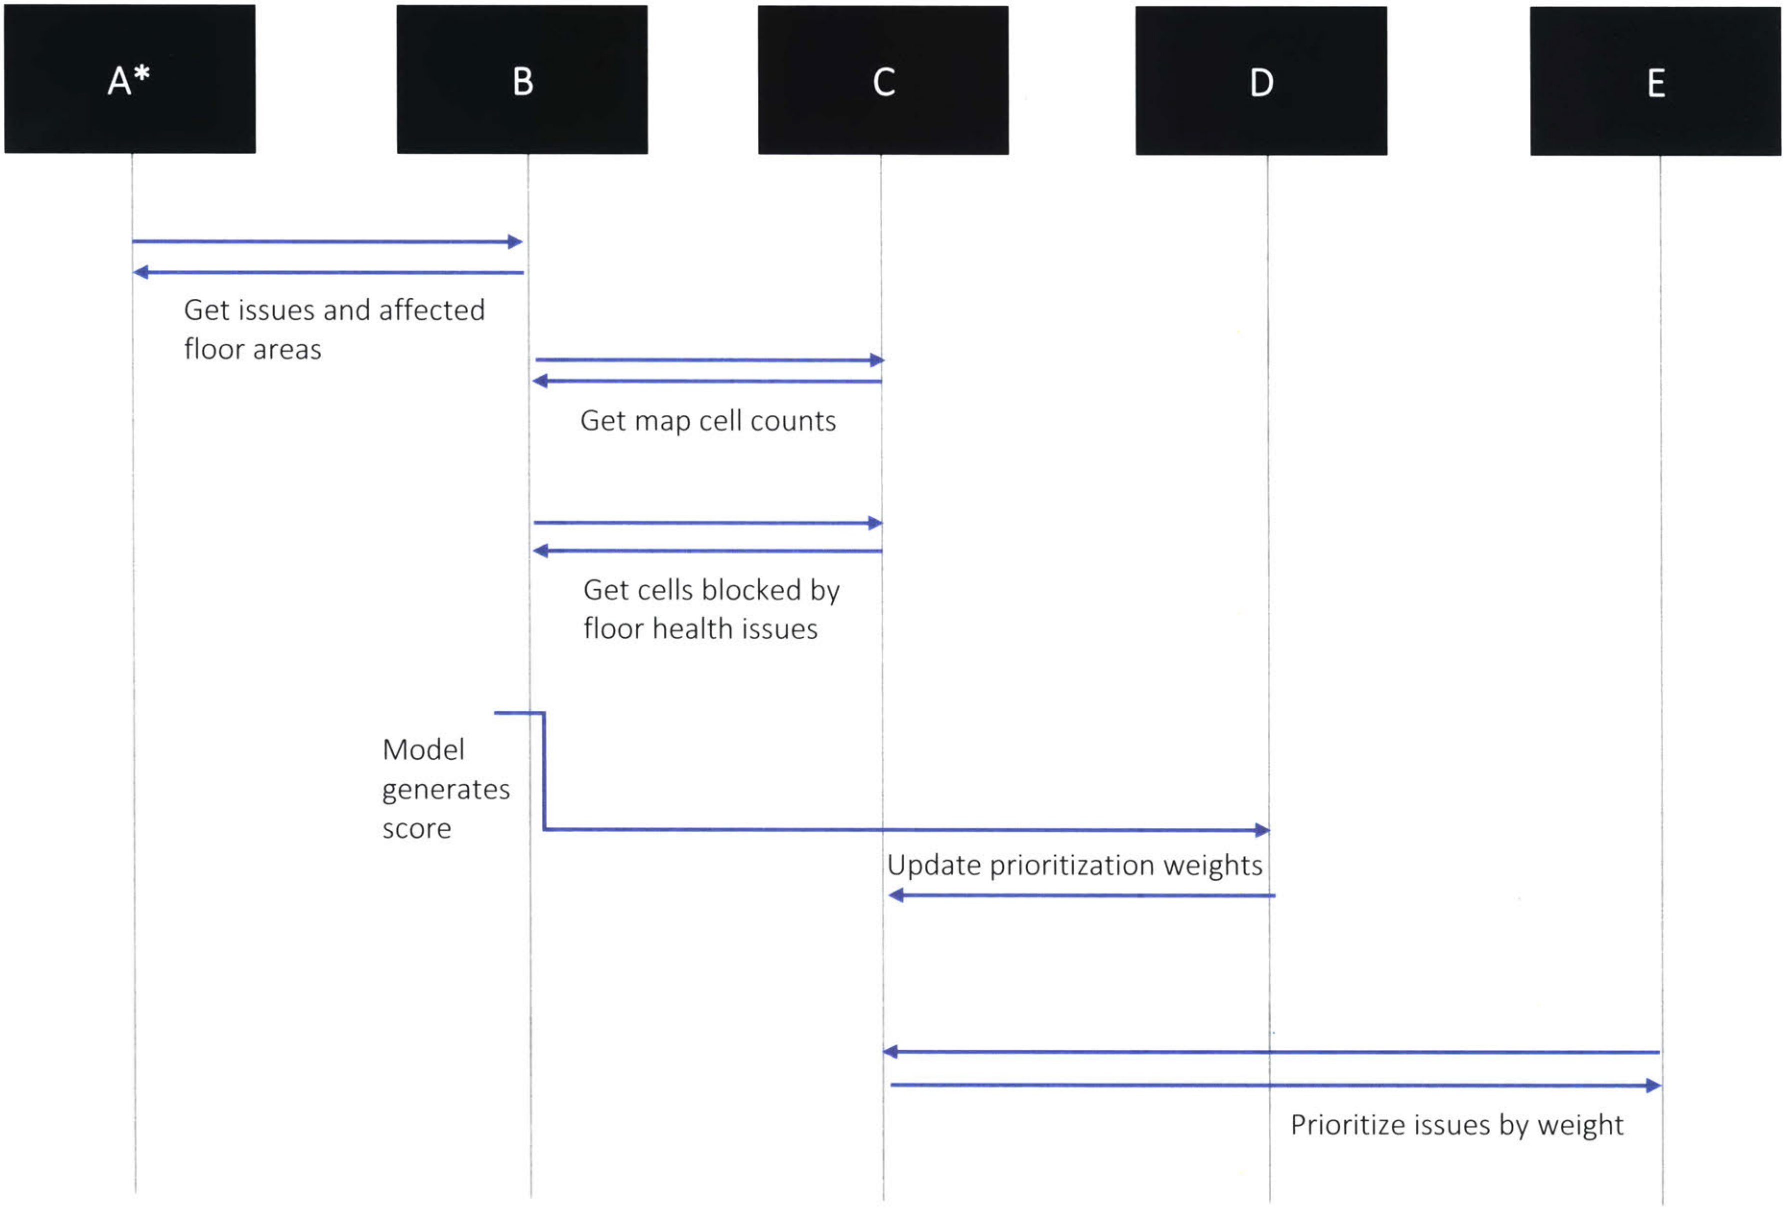

*Actual product/service names are anonymized to protect propriety information

REFERENCES

[1] Our Fulfillment Centers. [Online]. Available: https://www.aboutamazon.com/amazon-fulfillment/our-fulfillment-centers

[2] Our Vision. [Online]. Available: https://www.amazonrobotics.com/#/vision

[3] Gonzalez, Angel. Amazon's robots: job destroyers or dance partners? 2017. [Online]. Available: https://www.seattletimes.com/business/amazon/amazons-army-of-robots-job-destroyers-or-dance-partners/

[4] Small, Aaron Alexander. Gamification as a Means of Improving Performance in Human Operator Processes. 2017. EBSCOhost, search.ebscohost.com/login.aspx?direct=true&db=cat00916a&AN=mit.002565479&site=eds-live&scope=site.

[5] Stowe, James DeWitt. Throughput Optimization of Multi-Agent Robotic Automated Warehouses. 2016. EBSCOhost, search.ebscohost.com/login.aspx?direct=true&db=cat00916a&AN=mit.002433581&site=eds-live&scope=site.

[6] Robots at the doorstep: Amazon's Christmas gift to customers. 2014. [Online]. Available: https://www.sott.net/article/289671-Robots-at-the-doorstep-Amazons-Christmas-gift-to-customers

[7] Wingfield, Nick. As Amazon Pushes Forward with Robots, Workers Find New Roles. [Online]. Available: https://www.nytimes.com/2017/09/10/technology/amazon-robots-workers.html

[8] Kopytoff, Verne. How Amazon Crushed the Union Movement. [Online]. Available: http://time.com/956/how-amazon-crushed-the-union-movement/

[9] Burke, Michael. Robots make work easier at Amazon fulfillment center. [Online]. Available: https://journaltimes.com/business/local/robots-make-work-easier-at-amazon-fulfillment-center/article_b4075ab5-b38b-5a44-b571-89ffaf631c61.html

[10] La Gorce, Tammy. Despite Decision, Amazon Has Huge NJ Presence. [Online]. Available: https://njmonthly.com/articles/jersey-living/despite-decision-amazon-huge-nj-presence/

[11] Herrero-Pérez, D., and H. Martínez-Barberá. "Decentralized Traffic Control for Non-Holonomic Flexible Automated Guided Vehicles in Industrial Environments." Advanced Robotics, vol. 25, no. 6/7, Apr. 2011, pp. 739–763. EBSCOhost, doi:10.1163/016918611X563283.