

Predicting Hospital Capacity and Efficiency

Authors removed for double blind

Keywords: Predictive Analytics, Data Warehousing, Patient Movement, Discrete Event Simulation

Abstract: Hospitals and healthcare systems are challenged to service the growing healthcare needs of the population with limited resources and tightly restrained finances. The best healthcare organizations constantly seek performance improvement by adjusting both resources and processes. However, there are endless options and possibilities for how to invest and adapt, and it is a formidable challenge to choose the right ones. The challenge is that each potential change can have far reaching effects. This challenge is exacerbated even further because it can be very expensive for a hospital to experience logjams in patient movement. Each and every change has a “ripple” effect across the system and traditional analytics cannot calculate all the ramifications and opportunities associated with such changes. This project uses historical records of patient treatment plans in combination with a virtual discrete event simulation model to evaluate and predict capacity and efficiency when resources are added, reduced or reallocated. The model assigns assets as needed to execute the treatment plan, and calculates resulting volumes, length of stay, wait times, cost. This provides a valuable resource to operations management and allows the hospital to invest and allocate resources in ways that maximize financial benefit and quality of patient care.

1 INTRODUCTION

Hospitals and healthcare systems, especially US-based academic healthcare institutions, are under constant pressure to streamline and achieve more with limited resources and finances. With the growth of electronic medical record systems (EMRs), healthcare data warehouses, and business intelligence technologies, there is more information available to identify problems and opportunities for improvement. For example, we have robust dashboards which show length of stay for patient cohorts and how these values compare to benchmarks. We have information around patient time in “boarder” status, i.e. being boarded in one department when they belong in another. Examples include patients in the emergency room who have been admitted to the hospital, patients in surgical recovery rooms who are no longer under the effects of anesthesia, and patients in intensive care units who no longer require critical care.

The challenge though is use all this information to make good decisions and initiate valuable change. We can create change by adding or reallocating resources and by changing processes. Resources that could be adjusted include staff and staff schedules,

beds and bed accommodations, operating rooms and operating room allocation schedules, imaging resources, clinic and ambulatory surgery locations, staffing and hours and more. Example processes that can be manipulated include: house and bed management; emergency room registration, triage and rooming procedures; preventive care initiatives; isolation process; discharge procedures; care management protocols; urgent care facilities. It is challenging enough to predict the cost and return on investment (ROI) from such changes to make educated decisions. However, what really makes the challenge difficult is that each change can have far-reaching effects. For example, adding emergency department resources could create a logjam of ED patients waiting for an inpatient bed, or cause significant congestion in inpatient departments. Adding an OR room or adjusting operating schedules can create problems placing patients after perioperative recovery. Even a simple change like adjusting discharge order times can create a logjam of patients waiting for the orders to be processed. The challenge is exacerbated because it can be very expensive for a hospital to experience logjams in patient movement. For example, an admitted patient who is still in an emergency department bed is acquiring additional cost without compensation, and

very likely getting reduced care. A patient occupying a regular bed but requiring critical care treatment is not only receiving reduced treatment but is also creating additional staff requirements for the non-critical care department. Each and every change has a “ripple” effect across the system and traditional analytics cannot calculate all the ramifications and opportunities associated with such changes.

In addition to the operational challenges above, there are both opportunities and risks for the actual clinical treatment of the patient. It is important to ascertain that any change in process or staffing does not reduce care. Generally, we measure quality of care by key performance indicators (KPIs) including readmission rate, mortality rate and rate of hospital-acquired infections. Additionally, disposition should be included to indicate that a discharge to home is a better outcome than dispositions such as discharge to skilled nursing facility. Finally, patient experience (customer satisfaction) results can be included.

Additionally, treatment processes can be improved. These include disease-specific clinical protocols such as initiatives for sepsis, exacerbated COPD, congestive heart failure, pneumonia and stroke. There are also initiatives around specific clinical events such as ventilation, blood transfusions, administration of broad spectrum anti-biotics, management of central lines and foley catheters. Each of these can have far-reaching effects. For example, improving ventilation protocols can reduce reintubation and instances of ventilator-acquired pneumonia, and can result in reduced readmissions to ICU units. Improved management of vancomycin can reduce instances of kidney failure and the requirement for dialysis. A significantly successful clinical program can increase the stress on the hospital by increasing the survival rate and putting pressure on lower acuity units to absorb more patients.

There is simply no meaningful way for anyone to manually determine all the possible ramifications, both positive and negative, from a change in resource allocation, process or treatment. In this project we are proposing a digital simulation model which coupled with bio-statistics can provide much more meaningful and actionable insight.

2 BACKGROUND

In our investigation and design of this project we looked at literature and industry use cases surrounding:

- Predictive analytics in hospitals
- Tools for what-if scenarios for healthcare models

■ Predictions models in other industries

Discrete event simulation is the process of creating a digital model of specific events and rules. Every treatment or movement in a hospital can be considered an event. Tools and systems for discrete event simulation are sometimes referred to as “digital twins”. Digital twins create a complete virtual model of a physical asset or process. Digital twins can be networked together. Discrete event simulation and digital twins are not new technology. They have been used for years to simulate events that are too hard to always test manually (Fishman 1978), such as the impact of heat or wind speed on a plane (Tuegel 2011). Also, they are used to simulate entire systems which are connected by well-defined rules. For example, they could be used in automobile manufacturing to determine if there is a cost advantage to a more expensive paint which dries more quickly (Grieves 2014). This would be based on the capacity of the additional manufacturing systems to absorb additional auto capacity more rapidly. There are many commercial digital twin products available on the market, as well as open source products such as Ditto (Glocker 2017) and OMNet++ (Varga 2001) and SimJava (Howell 1998).

Despite the long history of this technology, it has become much more mainstream in recent years. Digital twins were recently identified by Gartner as one of the ten significant trends of 2017 (Gartner 2017). The reason for this sudden growth is IoT (the internet of things). As there are more small devices connected to the network, and so much is known about these devices and sensors, there is more opportunity for digital modelling. For example, consider a smart, connected thermostat. If it knows the goal is that the house is 72 when you get home between 5:30 and 6:00pm, it can also get the weather report and the current temperature. Then a digital model can be used to simulate the weather and your air conditioner. It can then determine the most energy and cost efficient way to achieve the goal of 72 degrees by that time.

Despite the growth of this technology, it has rarely been applied to healthcare. Healthcare is a much less predictable model. A patient treated quickly and politely may still decide to leave the hospital against medical advice. A healthy patient might suddenly have a totally unexpected aneurism. A patient whose life was just saved may still be bitter and unsatisfied. Despite being optimally staffed for a normal Friday night any particular Friday night may have a mass shooting or an earth quake. System-wide predictive

analytics in healthcare is always challenging and discrete even simulation models are no exception.

In specific contexts, there has been significant work on predictive analytics in healthcare. Sutter Health was able to accurately predict 30 day readmission rates and leverage the information to target and reduce readmissions (Ng 2014) (Jamei 2017). Similarly, (Bardhan 2014) was able to predict readmission for congestive heart failure after analysis of data from 67 hospitals. Parkland Hospital has been able to accurately predict daily census and several hospitals have used similar technology to improve coding and reimbursement (Bradley 2010). (Maguire 2013) describes a successful predictive data mining project specific to diabetes. There are countless more examples, and we have tried to leverage what worked for them, but they tend to be limited to specific patient cohorts and single data points. While digital models for hospitals has been widely discussed, we were unable to find a successful system-wide project in the literature.

3 PROPOSED APPROACH

We will detail the implementation plan and phases for our project in section 5 after describing our specific requirements in section 4. In this section, we will just give a quick overview of the approach.

The general approach is to:

1. Extract patient treatment plans from historical encounters. These treatment plans must separate what is necessary to heal the patient from what happened in the historical encounter due to hospital inefficiency or capacity. We use biostatistical analysis to transform actual events into a treatment plan for that patient.
2. Model the resources and allocation rules based on the business processes of the hospital.
3. Create a discrete event simulation model which places random patients into the hospital according to historic trends and then assigns resources to the patients according to their treatment model.
4. Collect statistics around the encounters as they pass through the model hospital.
5. Repeat the randomization many times to determine a histogram of results to show both the median values for length of stay, volume and other KPIs, and also worst and best case scenarios and the probability of significant logjams.

Once this is done, this model will support adjusting the resources and allocation rules, adjusting the treatment plans, and adjusting the patient flow based on market analysis and trends. Each of these creates a what-if scenario analyzed with the same system above, and multiple changes can be done in the model together to evaluate an entire proposal. This is all detailed much more in Section 5.

4 REQUIREMENTS

This project is a planned solution to a variety of metrics for a hospital system. Before we outline our solution and implementation plan we need to more clearly identify the requirements for the system.

4.1 Measures

The following list of measures and key performance indicators (KPIs) was identified by our Operational Excellence team. The goal is that the solution will be able to accurately predict each of these measures.

- Throughput and Volume
 - » Inpatient length of stay
 - » ED length of stay
 - » Surgical volume
 - » Observation hours
 - » External transfer volume
 - » External transfer acceptance rate
 - » Isolation patient days/census by hour of the day
 - » Unit census by hour
 - » Level of care census by hour
 - » Discharges by hour
 - » PACU (post-anesthesia care unit) census by hour
- Wait times
 - » Boarding time
 - » ED
 - » ICU
 - » Surgery
 - » Patient wait time for hospital admission (from patient perspective)
 - » Discharge order to discharge
- Utilization and Productivity
 - » Specialized transport utilization
 - » Staff time spent working at computer
 - » Telemetry utilization
 - » Staff productivity
 - » OR Primetime Utilization
 - » OR Service Block Utilization

- » Imaging utilization
- » Lab utilization
- » Pharmacy utilization
- » Bed utilization
- Finances
 - » Contribution margin
 - » Activity based costs
 - » Cost per discharge
- Patient and clinical Metrics
 - » Uninterrupted hours when the patient can sleep
 - » Readmission rate
 - » Mortality rate
 - » Wellness scores
 - » Patient experience scores

Each of these measures we will want to look at by specific patient cohorts based of diagnosis or diagnosis-related group (DRG), by unit, by day or time of day, by location, by procedure, by patient age, and other groupings.

4.2 Assets and Resources

The goal is that this project support what-if scenario analysis based on adding or removing resources or changing schedules. The following is the list of such resources determined so far:

- Beds

There can only be one patient in a bed at a time. Hospital support one or more combination of patient class, service and level of care. ED beds support one or more combination of acuity and age.
- Rooms

Rooms include one or more beds. Multiple beds in the room can only be filled if the genders and isolation status of the patients matches.
- Staffing

Staffing includes nurse and physician staffing by unit, anesthesia staff, pharmacy staff, imaging staff including technician and radiologists, housekeeping staff, transport staff, and ancillary staff including labs and pharmacy. Staffing volume can be adjusted by hour and/or unit or location.
- Operating Rooms
- OR hours

The operating rooms support various volumes of staffed rooms at different times and days of the week. Moreover, there are primetime hours set by day of week and location. Non-emergent

surgeries are delayed until prime hours. So the OR hours can be adjusted.

- OR service block allocation

Loma Linda holds operating rooms for specific services according to the service block allocation. This allocation can be changed.
- Imaging resources

Imaging devices such as MRI machines can be added, removed, or given a new operating schedule.

4.3 Treatment or process adjustment

We would like our solution to support adjusting specific treatment plans, clinical processes or operation processes, in addition to adjusting resources. For example, if we are implementing a stroke initiative that has been shown to reduce stroke length of stay by 10% at other hospitals, we can use our model to predict our measures given that all stroke patients need 10% less treatment time. If we have an initiative to write discharge orders or lab orders earlier in the day, we can analyze the potential effect of this change, and the appropriate alteration in staff schedules. If we have a program to reduce red blood cell transfusions by 50%, we can analyze the likely effect of this result on ancillary results such as length of stay, readmissions, hospital-acquired infection rate, cost, and lab productivity.

4.4 Use case examples

In this section we will list use cases we have been given as potential questions this solution would be able to answer.

1. Do we have the capacity necessary to achieve the annual projections for growth by service line?
2. Do we have the capacity necessary to handle the increase expected from expanded primary care coverage?
3. Do we have the capacity necessary to handle the increase expected from CIN and ACO arrangements?
4. What would be the impact on patient flow if we:
 - a. Reduce room exit time?
 - b. Move beds from MICU to SICU
 - c. Open an operating room in Saturdays for non-emergent cases
 - d. Perform more triage in the ED waiting room
 - e. Reduce the number of outpatient OR cases performed at the acute hospitals.
 - f. Increase transfers from the ED to auxillary campuses
 - g. Reduce readmission rates

- h. Reduce OR case duration variation for same procedures
- i. Reduce consult response time for consults ordered in the ED
- j. Reduce turnaround time (close-to-cut) for OR procedures
- k. Increase hours of operation for ORs
- l. Reduce the number of urgent admits that arrive in the ED
- m. Reduce PACU Boarder time
- n. Reduce the number of patients that are discharged the next day even though they are clinically ready to go the night before
- o. Expand shifts for care givers
- p. Increase number of post-op patients that transfer to auxiliary campuses
- q. Increase or decrease the number of outpatients in inpatient beds
- r. Reduce the number of observation patients
- s. Reduce the number of isolation patient days
- t. Reduce the number of avoidable days
- u. Increase the number of uninterrupted hours of sleep and when they occur
- v. Decrease the number of pokes (lab draws, images, medication administrations, line placements)

4.5 Goals

As we look at this vast number of potential measures, use cases and what-if scenarios, it is important we focus on the strategic goals. They are to optimally assign resources and create processes so that we:

1. Increase transfer volume and acceptance rate
2. Increase surgical volume
3. Reduce cost
4. Reduce length of stay
5. Reduce boarding times
6. Reduce readmissions
7. Reduce mortality
8. Increase patient satisfaction

The last three are highly clinical and difficult to establish clear causal relationship for and predict, so we will concentrate on goals 1-5.

5 IMPLEMENTATION PLAN

5.1 Phases

This project is broad so our goal is to create a model that can be expanded throughout several phases.

Phase 1: Proof of Concept

In this phase, we will create the model specific only to beds and surgeries. We will calculate the volume, throughput and utilization measures excluding staffing and block utilization. We will concentrate on a single use case: What will happen when we experience growth by service line according to projections?

Phase 2: Adjusting assets and treatment plans

In this phase we will build on phase 1 by allowing beds to be reassigned, operating rooms to be added and treatment plans to be changed. We will focus on the complex scenario: Predict the measures given that we reassign X number of beds from the surgical ICU to the medical ICU, we reduce surgical ICU length of stay 10% and we open one operating room on Saturdays for non-emergent cases. This shows how scenarios need to include multiple factors to truly allow what-if analysis, and the ability of the model to support this.

Phase 3: Adjusting operating room schedules

Currently, some operating rooms are pre-assigned to specific services on specific days. This is a virtual floating operating room and does not specify the physical room where the surgery takes place. In this phase, we will support what-if scenarios where the blocks are reassigned and we will include the block schedules in assigning OR rooms and staff to encounters. We will also add the measure of OR block utilization.

Phase 4: Anesthesia and ancillary staffing

In this phase, we will add measures around staff productivity, we will include staff capacity in our model and we will support what-if scenarios of changing staff levels specific for anesthesia, labs and images. We are not including nursing or attending physician staffing to reduce complexity. Nurses are generally staffed to bed occupancy anyway. To add staffing to our solution, we will bring in data from human resources and from the time and attendance tracking system.

Phase 5: Cost

In this phase we will add cost data to get an understanding of the financial impact of the what-if scenarios from the previous phases.

5.2 Team

The best opportunity for success in this project comes from building a diverse team of experts in different complimentary disciplines. Our team includes:

1. A computer science PhD specializing in data warehouse analytics models for healthcare
2. Two biostatisticians specializing in advanced statistics and predictive analytics
The director of operational performance improvement for the hospital
3. A professor of data mining and big data from a leading university
4. A clinician who is a registered nurse and has a Masters in Health Informatics and specializes in clinical details for data warehouse analytics
5. Two quality nurses from the hospital patient safety department
6. Experts in visualization tools
7. Data architects specializing in databases and data warehouses
8. Two technology experts specializing in utilizing our chosen technology

5.3 Technology

We have investigated using both commercial and open source digital twin products, but our plan involves instead leveraging the technology the hospital already has invested in. This includes Microsoft SQL Server and an advanced enterprise

data warehouse which will support both calculating the KPIs for the current environment and providing the historical patients and treatment plans needed for our model. This tight integration allows us to continue to support our enterprise data definitions and “single version of the truth” in both the actual data warehouse and the new virtual model. We also utilize Cisco DV, a data virtualization platform that allows us to quickly integrate multiple data sources and manage information independent of physical source. This is important as we bring in human resource data, cost information and time and attendance records. Cisco DV platform can also be used to enhance our system with auxiliary data such as avoidable delays or hospital-acquired infections from our epidemiology system. Additionally, we are leveraging SAP HANA (Farber 2012), an in-memory data analytics appliance which will provide exceptional performance and has a built in predictive analytics library. We are also using R an open source language specifically targeted to advanced statistics and predictive analytics. System R, originally a IBM product, is an open source suite of such tools and algorithms. “R has also been identified by the FDA as suitable for interpreting data from clinical research.”[38] Finally, we are using the dashboard visualization tool Tableau. Tableau is already in place at our hospital as the delivery mechanism for dashboards and data discovery, and it supports both HANA and R. This allows us to create a common set of tools, data definitions and interfaces rather than introducing new technology and learning curves. Figure 1 shows our system architecture.

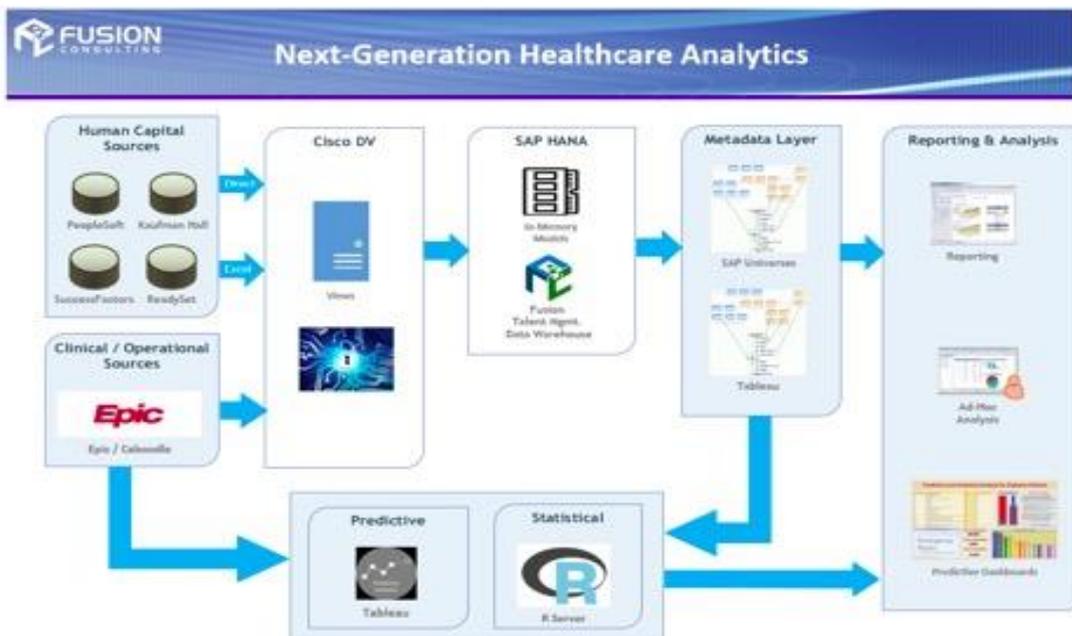


Figure 1: Data analytics system architecture

5.4 Specific Implementation Plan

We have a complete execution plan for the first two phases of our project and the first phase is under way, with an expected 14 week timeline.

5.4.1 Calculating the KPIs

For the first phase, we are targeting the KPIs around discharge, transfer and surgical volume, and around length of stay, census, bed utilization, isolation and boarding times. Most of these data points are in the enterprise data warehouse, but we need to aggregate them appropriately and add a few fields around the transfer center and isolation status.

5.4.2 Data Preparation

We need to create our historical patient encounters. We do this as a multiple stage process. First, we extract from our data warehouse and build an event log for each hospital encounter. This log includes timing information concerning all patient movement. This will include arrival and triage time, each bed movement, each order which changes the patient's status or level of care, each level of care or service change, and the admit and discharge orders. Additionally we add events surrounding surgeries: when they were scheduled, when they were performed, how long they took, etc. For phase 1, we are not considering imaging, labs or medications yet. This level of specificity and these resources will be added in future phases.

We then use our statistical tools and biostatisticians expertise to convert our historical event logs to treatment plans. For example, consider a patient who spent 1 day in inpatient ICU status but physically still in the emergency department and two days in ICU status in a ICU bed. Does their treatment require two days of ICU care or three? The answer is probably somewhere in between because they were receiving care while in boarder status in the ED, but not the same care as they would have received in the critical care unit. The goal of the statistics in this phase is to determine the relationship of boarder time to length of stay. Once we have this information, we store with each historical encounter a mini treatment plan. This plan consists of:

- The length of time in ED (not including waiting room) prior to decision to admit or discharge.
- The length of time in each patient status/level of care/service combination.
- Each surgery, when during the treatment plan above it was requested, and its duration and type of urgency.

5.4.3 Modeling the assets

For phase 1, the assets we care about are limited to just beds and operating rooms. For each bed outside of ED, we generate a bridge table which shows every type status/service/level-of-care combination supported by that bed. We also map the beds to rooms (which is already in the data warehouse). Now we have a list of resources which we can use to support events on our encounters and mini treatment plan.

For ED, we have six types of beds: rapid intake, adult, express care, pediatric, diagnostic and surgery. ED is modelled somewhat differently. The only deciding fact is acuity, age, and time of day (express care is not open 24 hours a day).

Operating rooms we load the list of operating rooms/locations, the prime time hours and the staffed hours for emergent cases. We are not concerned with service block schedules in the first two phases.

5.4.4 Building the allocation agent

We build a set of coded procedures which support the treatment plan of patient encounters by:

- Assigning the patient to an appropriate bed when they are in boarder or waiting room status
- Adding a hold to the other bed in a room (as based on gender and isolation status)
- Moving patients around to support holds (for example if two males are in separate rooms with empty beds in a unit and a female needs a beds, the males can be moved together.
- Assigning an operating room and executing a surgery at a specific time

5.4.5 Additional Statistics

We calculate additional statistics to support our model. These include:

- Histograms which tell us how many patients to expect at specific times
- Transfer center statistics to show us the relationship between the wait time for a bed and the cancellation rate of requested transfers
- ED statistics to show us the relationship between delays in treatment and patients leaving prematurely
- Discharge processing statistics to show us the relationship between discharge order times and volume to actual discharge times.

5.4.6 Virtual Model Execution

To execute our model, we place patients in the hospital (via ED arrival, direct admit or transfer

request) according to our histograms from the statistical analysis. We then loop every “15 minutes” on our virtual clock and add new patients, discharge patients, and attempt to move patients according to their treatment plan. Any time a patient is in boarder status, we will attempt to move them to an appropriate bed if available. Each time a patient leaves the hospital in our virtual model, we record statistics about their encounter to match our KPIs.

We load our patients completely randomly. We take a number between 0 and 1 and multiply it by our number of patient encounter treatment plans to choose one. To make up for this randomization, we run the same patients through with different random variations thousands of time to calculate the median KPI values and the histogram of each measure. This allows us to not only predict the measures but also show where the potential for challenges or logjams is.

5.4.7 What If Scenarios

The scenario we are supporting for Phase 1 simply involves adjusting the patient encounters according to given growth projections by service line. We can achieve this with our statistics which define patient encounters to add to virtual hospital model.

For phase 2, we support three types of adjustments: beds (add/delete/reallocate), operating rooms (add/delete/change hours) and treatment plan.

5.4.8 Testing

Before we use our model to make decisions, it is obviously important to be certain our predictions are accurate. To achieve this our plan is to a traditional training and testing process. We will train using historical data from encounters 6-24 months ago and then test by attempting to accurately predict the last 6 months for the hospital.

6 CONCLUSIONS

We have proposed a staged solution to allow hospital to create “what-if scenarios” and predict system results from such scenarios over a large number of important measures. We have created an achievable targeted plan which delivers value in each of several short phases as it builds the entire model. Our model leverages the EMR and enterprise data warehouse, well-known data mining techniques and biostatistical algorithms. It uses biostatistics, discrete event simulation and business process modelling in tandem. It is achievable, measurable and flexible. We believe it will create a solid foundation for predictive analytics for our hospital system.

REFERENCES

- Glocker, G., Jaekle, T., 2017. Eclipse Ditto. Viewed August 2017. <https://projects.eclipse.org/proposals/eclipse-ditto>.
- Tuegel, E.J., Ingraffea, A.R., Eason, T.G. and Spottswood, S.M., 2011. Reengineering aircraft structural life prediction using a digital twin. *International Journal of Aerospace Engineering*, 2011.
- Grievies, M., 2014. Digital twin: Manufacturing excellence through virtual factory replication. White paper.
- Fishman, G.S., 1978. Principles of discrete event simulation.
- Varga, A., 2001. Discrete event simulation system. In Proc. of the European Simulation Multiconference (ESM'2001).
- Howell, F. and McNab, R., 1998. SimJava: A discrete event simulation library for java. *Simulation Series*, 30, pp.51-56.
- Jun, J.B., Jacobson, S.H. and Swisher, J.R., 1999. Application of discrete-event simulation in health care clinics: A survey. *Journal of the operational research society*, pp.109-123.
- Ng, K., Ghoting, A., Steinhubl, S.R., Stewart, W.F., Malin, B. and Sun, J., 2014. PARAMO: a PARAllel predictive MOdeling platform for healthcare analytic research using electronic health records. *Journal of biomedical informatics*, 48, pp.160-170.
- Jamei, M., Nisnevich, A., Wetchler, E., Sudat, S. and Liu, E., 2017. Predicting all-cause risk of 30-day hospital readmission using artificial neural networks. *PloS one*, 12(7).
- Bardhan, I., Oh, J.H., Zheng, Z. and Kirksey, K., 2014. Predictive analytics for readmission of patients with congestive heart failure. *Information Systems Research*, 26(1), pp.19-39.
- Bradley, P. and Kaplan, J., 2010. Turning hospital data into dollars: healthcare financial executives can use predictive analytics to enhance their ability to capture charges and identify underpayments. *Healthcare Financial Management*, 64(2), pp.64-69.
- Maguire, J. and Dhar, V., 2013. Comparative effectiveness for oral anti-diabetic treatments among newly diagnosed type 2 diabetics: data-driven predictive analytics in healthcare. *Health Systems*, 2(2), pp.73-92.
- Gartner, 2017. Gartner's Top 10 Strategic Technology Trends for 2017. <http://www.gartner.com/smarterwithgartner/gartners-top-10-technology-trends-2017/>.
- Färber, F., May, N., Lehner, W., Große, P., Müller, I., Rauhe, H. and Dees, J., 2012. The SAP HANA Database--An Architecture Overview. *IEEE Data Eng. Bull.*, 35(1), pp.28-33.