



Machine Learning to Save Lives

Case Study

Use Case

Studies show that patients who undergo an unplanned transfer to the Intensive Care Unit (ICU) experience significantly worse outcomes than patients who are directly admitted to the ICU. Patients who are moved to an ICU as a result of a sudden deterioration typically stay in the hospital 8 to 12 days longer and have significantly higher mortality rates than patients who are not rushed to intensive care from another ward. The Kaiser team built models to try and predict the probability of a patient experiencing a sudden deterioration based on available electronic records, demographic information and vital signs and lab results collected during recent admissions. Together with clinicians, Kaiser data scientists developed thresholds and configured the Advanced Alert Monitoring (AAM) system to notify the care team if the patient might crash.

The Challenge

Kaiser Permanente is an integrated healthcare system, where patients receive a variety of services and treatments under one roof. With 10 million members, Kaiser Permanente has vast amounts of data, and is using innovative approaches to make this data work for the benefit of its patients and clinicians.

Multiple studies showed that patients in the surgical ward who get an emergency transfer to an ICU as a result of sudden deterioration have significantly worse outcomes than patients who are admitted directly to the ICU or get moved to the ICU before they “crash”. These patients account for only five percent of all hospital admissions, yet, they represent about a fourth of all Kaiser ICU admissions, a fifth of all deaths in the hospital, and about an eighth of all of the hospital days. In fact, patients who experience an unplanned transfer to the ICU experience two to five times the mortality of patients who are directly admitted to the ICU, and they would stay in the hospital an average of 8 to 12 days more than patients who are directly admitted to the ICU. “So we wondered,” contemplates Patricia Kipnis, the Principal Statistician at Kaiser, “Can we find out ahead of time who these patients are who are likely to crash and be rushed to the ICU, and can we do something about it?”

The Solution

Kaiser has always been on the forefront of using technology – as early as 1970s, the company began keeping electronic medical records. Today, Kaiser members have a variety of methods available to them to interact with their doctors and nurses – using the web,

Industry

Kaiser Permanente is a large integrated healthcare organization with global presence, thousands of physicians, tens of thousands of nurses, 38 hospitals and around 600 outpatient facilities and medical offices.

Challenges

- Working with real patient data in a highly regulated healthcare industry
- Operationalizing the data model
- Combining data analytics with clinicians’ opinions and conclusions to deliver best outcomes for the patient
- Near zero tolerance for failure
- Infrequent occurrence of critical deterioration events among patients

Solution


- Using analytics to calculate the probability of a rapid deterioration in a patient requiring unplanned transfer to ICU
- Flexible, configurable open source tools like H2O to help build and continuously improve models

Results


- Clinicians receive an alert if a threshold is exceeded to evaluate the patient and determine further course of action
- Intervention prior to deterioration creates better outcomes for the patient

e-mail and mobile apps to get the answers they need faster. All this data is captured and can be used to solve specific problems.

Using Data to Improve Intensive Care Unit (ICU) Patient Outcomes



Upon examination of many unplanned transfers to ICU, the Kaiser team noticed several trends: most patients began showing physiological deterioration 6-24 hours prior to a transfer, which went unnoticed by the medical care team. To catch these warning signs early, Kaiser decided to use available data about a patient and build a mechanism that could identify when deterioration



might occur. Kipins likens Kaiser’s analytics systems that help identify and intervene with at-risk patients to a hurricane early warning procedure. “We are calling it Advanced Alert Monitoring (AAM) System,” she said. “It helps identify patients who are likely to crash, and we recognize them 12 hours before they experience the deterioration.”

The AAM has four components: risk analysis and the creation of the score, monitoring and warning, dissemination and communication of the scope, and finally – the intervention. For the risk analysis portion, which is the actual building of a score that provides the warning, they used bed history data: where the patient had been in the past few days since the patient was admitted to the hospital; chemistry: which includes all of the lab work lab work; the vital signs or any vital sign information; any comorbidities: presence of one or more additional disorders (or diseases) co-occurring with a primary disease or disorder by bringing in all their past information into the models; and demographic information: so age, gender, etc. All of this information was used to estimate either the probability of a late transferring to the ICU that is a sudden crash or mortality in the hospital.past information into the models; and demographic information: so age, gender, etc. All of this information was used to estimate either the probability of a late transferring to the ICU that is a sudden crash or mortality in the hospital.

Once the data is collected and cleaned, the team worked on feature engineering and then applied a variety of training and testing datasets. “The part where we are different is the validation of the models,” added Taposh Dutta Roy, Data Leader & Science Manager, Decision Support at Kaiser. “When you are talking about saving lives, you cannot be off by even a little, so we go through rigorous responses and validation.” The team explored each parameter and its linkage, got the clinicians

involved to look at every case, and then performed a characteristics review and threshold analysis to make sure that a threshold made sense.

Open Source is the Right Way to Go

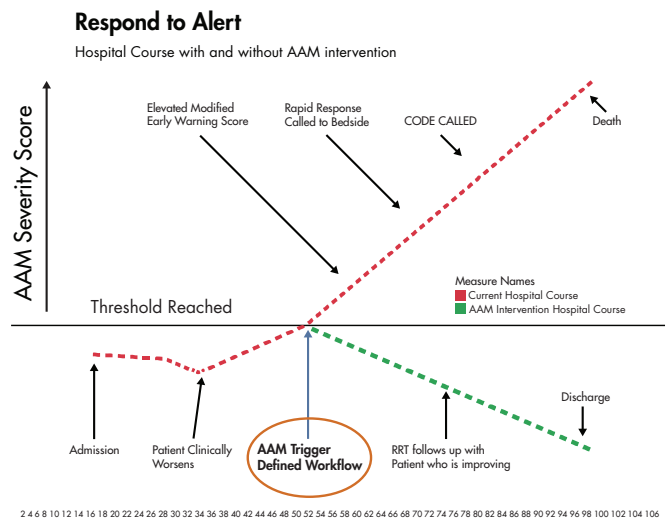
Before the system could be operationalized, the team underwent several proofs of concept -- to improve the infrastructure, processing time, experiment with different Machine Learning algorithms and big data tools. “We learned a lot,” continues Roy. “We saw how faster processing and faster access to data can help. We decided that we don’t need to sample data because we have these new advanced tools available to us. And, finally, we decided that open source was the right way to go, as long as we have support.”

“We saw how faster processing and faster access to data can help. We decided that we don’t need to sample data because we have these new advanced tools available to us. And, finally, we decided that open source was the right way to go, as long as we have support.”

– Taposh Dutta Roy, Data Leader & Science Manager, Decision Support at Kaiser

Results

The results are currently available every six hours, but AAM can be configured to calculate the likelihood of critical deterioration on an hourly basis. If the probability exceeds predetermined threshold, the team is notified to come in and evaluate the patient and make a decision on any further actions. “If we are able to intervene prior to deterioration,” adds Kipnis, “the team can either send



2 4 6 8 10 12 14 16 18 20 22 24 26 28 30 32 34 36 38 40 42 44 46 48 50 52 54 56 58 60 62 64 66 68 70 72 74 76 78 80 82 84 86 88 90 92 94 96 98 100 102 104 106

a patient to ICU earlier or take action that would avoid the transfer altogether, and definitely help avoid the worst outcomes.”

The Kaiser team plans to continue testing and improving their models, adding new features and refining the monitoring and warning process. “At Kaiser, we are

focused on outcomes,” concludes Roy, “It takes a village to get these things done, and everyone from IT engineers to statisticians and data scientists are dedicated to this project of using data to save lives.”

About H2O.ai

At H2O.ai we see a world where all software will incorporate AI, and we’re focused on bringing AI to business through software. H2O.ai is the maker behind H2O, the leading open source machine learning platform for smarter applications and data products. H2O operationalizes data science by developing and deploying algorithms and models for R, Python and the Sparkling Water API for Spark. Some of H2O’s mission critical applications include predictive maintenance, operational intelligence, security, fraud, auditing, churn, credit scoring, user based insurance, predicting sepsis, ICU monitoring and more in over 5,000 organizations. H2O is brewing a grassroots culture of data transformation in its customer communities. Customers include Capital One, Progressive Insurance, Zurich North America, Transamerica, Comcast, Nielsen Catalina Solutions, Neustar, Macy’s, Walgreens, Kaiser Permanente and Aetna.

