Michael Thompson: Thank you. Thank you for the introduction. As you guys would expect, this original presentation was created in general circumstances for hospital flow. In parts of the presentation we will talk about how it's being used in some of the recent events that we've had, particularly in California. Let's go to the next slide.

Michael Thompson: Our original intention is to make sure that we review what our intentions are for learning the objectives of this presentation. We'll talk a little bit about my organization and the data science team. We'll work through a particular problem we had, originally, when we created this presentation and then review of how we looked at some of the models as well. Next Slide. Now let's go ahead and go to the next slide after this.

Michael Thompson: A little bit about Cedar-Sinai Medical Center. We're located in Beverly Hills, in Los Angeles, California, 958 beds. We have over 50,000 patients and admissions per year. We are an academic center, so we have 437 medical residents and fellows. We have a second hospital in Marina Del Rey, which has 154 beds. We also have an outpatient services in Cedar-Sinai Medical Network of 350,000 patient visits per year and growing. Next slide.

Michael Thompson: I want to give some credit to my team. Although I'm up here speaking about some of the work that we've done, really the work is being done by this team of individuals. I joined Cedar-Sinai about three years ago and when we joined Cedar-Sinai, we decided to centralize the data analytics team underneath one umbrella. One area I have is the data warehousing team. The next is the data sciences team led by Todd Davis. My Data Quality and Data Delivery team. The work we're talking about today was done by the data science team and you can see the list of individuals we have there. They are a wonderful group of individuals and I hope I represent them well. Next slide.

Michael Thompson: One of the things we wanted to work on very early on was trying to improve our data science pipeline. As you would have guessed, any kind of project you're working on predictive analytics or data sciences, you need a wonderful pipeline to get the data, so we are using data directly from electronic medical records. We are grabbing the data from medical devices, consumer devices, online surveys. That data is going into a data lake and also our enterprise data warehouse. We store the data in a variety of different methods. Then, in our data transformation, we'll use open source programming languages and tools. A lot of the work we're looking at today was either used with Python or R or an automated machine learning tool that we use as well. Then, on top of that, we have data visualization software. Starting years ago, data visualization software existed, but the open source and the R with Python pipeline for data sciences did not exist nor did the data lake. Next slide.

Michael Thompson: Okay. Before I begin talking about the project that we are working on, which is obviously hospital census and hospital flow, it was important that we can actually can be in a team of individuals who can work on this together. This was not just a data science project. In fact, the idea for the project came out of our event planning group, and our event planning manager, and she said, "Can we predict census?". Well, as we started working on that project, we realized, in order to make it successful, we need to include a variety of people around the table. That included nursing staff, operations staff, physicians, and case management. As well as patient satisfaction staff. Any project that we take on, we realize now that it can't just be a data scientist. A data scientist is just one of many that sits around the table. Next slide.

Michael Thompson: Let's frame the problem. The problem that we were dealing with at the time was capacity strain. Capacity strain is when our capacity exceeds 80% of our availability of beds. In 2018, 12 California hospital facilities experienced capacity strain. This diagram you see to the right, we can see all the hospitals with beds greater than 300 in 2018. You can see the number of hospitals that are going above 80% in capacity strain. Some hospitals are consistently at 100%. In our case, we were somewhere within the range of between 80% and 90%, so 6% of our days, we were sitting above 80% in capacity strain. In order to solve capacity strain we've gone through a lot of projects, which include trying to reduce length of stay, which in turn creates discharges, but the problem is that our volume kept on increasing. So can we staff and figure out what's going to occur. Next slide.

Michael Thompson: What is the impact of capacity strain? Well, patients get bored in the emergency department. Patients will be sitting there for two hours or more waiting for a bed because a bed is not available. Patients will leave the emergency department without being seen. Patients will have overnight stays in post-op recovery rooms. We'll have ICU readmits within 24 hours which is basically, the patient who is in ICU is moved to another floor, and then had to move back into ICU. We have delays and cancellations of surgeries. We have physicians, nurses, and staff who are overloaded. Then our throughput has decreased, with delays in transferring patients to appropriate units. Next slide.

Michael Thompson: Our goal. Our goals was to reduce the need for surge planning. Surge planning occurred when capacity reached above 850 beds and we could be in a group or individuals to figure how we can discharge sooner or more safely, and then how we can handle the new arriving volumes. We need to reduce those surge planning needs. Initially, a surge plan would occur during seasonal periods of the year, and then every so often during the month. We starting seeing surge planning occurring almost every day. We need to prevent the diversions or overcrowding in the emergency department. We need to eliminate waits for surgical procedures. We need to improve staff schedules that match demand and we didn't want to overstaff, where we'd have excessive overtime. We need to increase the number of patients admitted to the inpatient unit. We need to utilize case management strategies, reduce length of stay outliers. We need to improve discharge and bed capacity planning. These are all the goals of our predictive analytics and machine learning project. A variety of goals, a lot of goals. I wish it had just been one of the goals, but all these goals were a target. Next slide.

Michael Thompson: Okay. Before we begin talking about predicting, we want to go through and figure out, what is hospital census? On the surface it sounds pretty simple. Well, how many patients are sitting in a bed right now and can I predict the number of patients that will be there tomorrow? Well, it's really an algorithm of the current patients in a bed, plus new patients requiring a bed, minus discharges. Now, new patients requiring a bed can come from elective admissions, then you have urgent direct admissions, then we'll have emergency department admissions. There's a factor of emergency department arrivals and emergency department discharges.

Michael Thompson: While we were hoping to have one model to predict hospital census, we found out we needed to [inaudible 00:11:20] to predict hospital census. A model to predict ED arrivals, a model to predict ED admissions, a model to predict urgent care admissions, a model to predict discharges. When we created these models we realized that we needed to have weekly, daily, monthly, seasonality involved with the model. We knew [inaudible 00:11:41] those were [inaudible 00:11:43] for elected admissions or elective surgeries, and we needed to identify those patients that we need to bed for elective surgeries. Next slide.

Michael Thompson: Historically, we can go back and look at our tree of census that looks like this. In 2018, we had 55,000 patients requiring a new bed; 20% of those were elective admissions, 80% of those were unscheduled admissions. When we look at predicting hospital census, one of the things we need to look at or the variability and the uncertainty. The uncertainty is in that 80%. Of that 80%, we can look at newborns, labor delivery, admissions for emergency department, transverse and direct urgent admissions, and down the line. As we did this decomposition tree, we realized, we do have a lot of uncertainty and that uncertainty is based upon seasonality and other factors beyond our control. Next slide.

Michael Thompson: Let's start out by talking about ED arrivals. Next slide. Typically, any data scientist will look at ED arrivals when doing a decomposition of a seasonality factors. This is what ours looked like. You can see an overall trend change. You can see changes in the yearly and monthly trends. You can see changes in weekly trends. You can see the effect of holidays, and you can see the residual effects. Now, if you haven't seen one of these slides before, it's really nice to do a decomposition, but if you have seen one of these slides before, you realize we're doing a seasonality analysis. We do see that we have seasonality in monthly, we have seasonality weekly, and have seasonality in holidays. Based upon this, we knew that we could reasonably predict emergency department arrival and visits basing on seasonality factors. We cleared that in our model. Next slide.

Michael Thompson: As we said, one of our models were to create the emergency department admission model. When we looked at those models, we actually do not just run one model, we run a variety of models. Anytime we do a data science project, we realize that some models are good during certain points of the year and some models are better at other points of the year. In this case, we have three different models we ran. You can see the different error rates on those models. The other third component is one type of prediction model. Sarema, or seasonally adjusted regressive moving average; I can't ever say that. That's a typical seasonality model that's used. If you do time series models, you always will use that model, sarema. The other one is the Facebook profit model. Now Facebook profit has done a great job of taking one of their models, which based upon Arima and adding other variables into it. They've opened sourced the model and we found that to be a very valuable model because it does deal with some more recent trends better.

Michael Thompson: The final model we looked at is Naïve. As you can see in here, Facebook profit, in this particular case on our training data, was probably the best performing model. It only had a mean absolute error of 4%, which means it had accuracy of 96%. Now, you may ask, what is the Naïve model? Naïve model is pretty interesting. It's basically just looking at what is the volume we had in the emergency department today and assume what the volume will be tomorrow. That was typically a process that was used before we actually started using these more advance methods, but the Naïve model itself wasn't too bad. A 5.55% absolute error, not so bad. 94.5% of the time, if you just said what the ED visits was a day and ED visits tomorrow, that's okay, but we needed to [inaudible 00:16:01] Naïve well in our figures. Facebook profit looks like it did very well, so we'll use that one for ED. Next slide.

Michael Thompson: This is just a visual comparison. Any time we go and start sharing this data with our round group of people. Remember, our project team of nurses, physicians, case managers, [inaudible 00:16:29], what's more better? As you can see here, the dots are the actual values. As far as actual visits during this period of time. As you can see, Facebook profit is fitting very well to those dots. Although, you do notice there are spikes that occur, and Facebook profit is not catching those spikes in volume, and we should take a look at that later and figure out what those spikes are. Overall, Facebook profit seems to working well. Next slide. Next slide.

Michael Thompson: Okay. I know this might be a little hard to read, but I wanted to keep it exactly the same as what we did our data scientists created. It's something interesting. We started displaying this information and sharing this information with the group. Typically, when we did data science projects in the last [inaudible 00:17:31] seven days is 251, 250, 250, 249. We had to share the [inaudible 00:17:39] that needed it. That's okay, but the problem that we wanted to solve for is we didn't want to create a model and then have it not be adopted. Every time that we start showing predictive values, the most important part of this slide is we had the bottom part, which shows how we performed last week. Every time we give a number to our group of individuals, we're predicting then to understand this variability of that number and can you they trust the number we're giving them? [inaudible 00:18:10] the data scientists just want to give the numbers and not show how well or how [inaudible 00:18:16] my model performed the last week. We can get away with that and buy in a lot better.

Michael Thompson: While the top is saying, "hey, here is the amount we predict for next week", [inaudible 00:18:29] is how well the models did last week. Now, they can go through a different model that changes the different lines in the chart, but as you can see to the right, [inaudible 00:18:41] last week Facebook profit was 3.5% error and some other models [inaudible 00:18:49] have a greater error and we still beat them by Evare. All I take away from here is that we've learned that when we share values with individuals, assuming we predict a value, also show how well you're doing and [inaudible 00:19:04] improve your models. Next slide.

Michael Thompson: [inaudible 00:19:10] To give you an idea of our ED admissions when [inaudible 00:19:23] with ED arrivals, some people go and say, "Well on average, you're 60% [inaudible 00:19:30] or 12% or 10% on my ED arrivals [inaudible 00:19:34] if you want. [inaudible 00:19:40] However, a handful of individuals use the percentage of those [inaudible 00:19:58] 15-16% [inaudible 00:20:00]

Michael Thompson: Okay, Great. We had the ED Arrivals. I just wanted to mention that on this slide and [inaudible 00:26:47] the slide you can see there's different types of models you could use for each one of these admissions. I did find that when we were creating some of these models it's not just one model that rules them all, you'll have to use an ensemble or expanded use of models so although you may use a time series for ED admissions you may use something else for [inaudible 00:27:06] admissions. And that will give us a [inaudible 00:27:09] requiring event. Next slide.

Michael Thompson: This again, reiterates the fact that we start sharing data with people the want us to share with them the scheduled surgeries which is very much a known factor and our predicted arrivals based upon the unscheduled arrivals. This follows the same philosophy that we had before; let's go ahead and share with people what we're predicting. Our philosophy is also share with them how well we did last week. And it does train people, if we go to the very bottom of the slide, what you'll see is an error rate. So the very bottom graph is what we predicted versus what actually happened. On Sunday, we were off by 18 patients, on Monday, we're off by less than half a patient. On Tuesday we're off by 28 patients. Now it's very important that they saw these numbers. Because we didn't want people to think that Monday was the norm. There is variability in how well we'll do, and we need to train people from the very beginning that there is going to be variability.

Michael Thompson: One of the concerns I have is that most recently our models have been very, very accurate. And the CIL walked into my office most recently and said, "You know, we're not sure how your models are doing because they were off by ten patients yesterday". And I was like, "Off by ten patients, that's great!". They got trained to be looking for numbers under 20, under 30, and that's not good because everyday is going to be above 25 or 30 patients, or more off. Training begins at the very beginning. Okay, next slide.

Michael Thompson: Discharges. Discharges is not always the best thing to talk about. How do you predict someone is going to leave the hospital? We can predict expected length of stay, or we can look at the expected discharge rate for average discharges. When you do these models, you need to have the physician in the room because they will challenge what you're saying. It's good for people to understand that we're trying to decide what new beds will be there. Discharge planning is what we did next for predicting [inaudible 00:29:25] explained.

Michael Thompson: You can have basically literature in the stuff that we're found, we have a variety of ways of predicting discharges. You can do a time series model and say, on average we have these run offs during these seasonality periods of the year. You can do an average run off model which is, on average the total population will stay four days or five days. You can actually go through and look at the conditions of patients and say, this patient is expected to stay here three days and they've been here for two days so therefore we expect them one day. Keep in mind that is either a top down or a bottoms up approach we take a look at.

Michael Thompson: The other one, number four. That's a little bit of contention because you're looking at expected discharge date. When looking at expected discharge date you're willing to go back and decide, are the physicians or clinicians entering a date that we feel comfortable with? Are they entering a date that's optimistic? Is the date entered out correctly? We just find that a lot of times the data's not here at all, or when it is entered, it's entered on the first day of arrival but was it entered while the patient stayed? Picking an expected discharge model, you are running into a variety of different factors and that includes the clinical factors and behavioral factors as well. I just want to point that out. Next slide.

Michael Thompson: Here is a histogram of our length of stay between 2012 and 2017. As you can see, the median, which is the red line is somewhere around three and a half, maybe four. And you can see that the mean is somewhere around five. But we do see we have a somewhat normal distribution with a tail to the right, which most hospitals [inaudible 00:31:16] will have. As you would expect, that tail to the right is the thing that gives us a little bit of a trouble, but on average we could use a median or mean length of stay for every patient that comes in and hopefully the wisdom of having a lot of different numbers would be useful. Next slide.

Michael Thompson: Then once we go through and start taking a look in we do see that our lengths of stay have been decreasing. Initially while we want to go through and look at 2009 to 2017, there's been a lot of factors in place that will reduce the number on lengths of stay that we would have. Starting in 2012, we see that there must have been an activity between 2009 and 2012 to reduce length of stay, so we need to make sure that we use more recent numbers in our distribution chart or histogram to decide what we want to do. You can also see we also have a misses increase to. But this is just an expiratory data analysis of length of stay for us. Next slide.

Michael Thompson: Then there's a bit of literature. Before we take on any project we do a lit review, a lit review will tell us is this even a feasible project to do, has someone else done it? And I added this into here because you can see this [inaudible 00:32:40] model that's here to say length of stay. This is something to go back and take a look at, then do use R-Squared numbers in here, I don't like R-Squared, but whatever you want to do for [inaudible 00:32:52]. It's always good that when you sit down and talk to the physicians or talk to the group that you're working to say, yes, this has been done in literature and are we just as good, worse off or better than the literature is saying? Next slide.

Michael Thompson: We did use a variety of different models to predicting expected length of stay. In this particular case, we're predicting expected length of stay based upon patient factors or as [inaudible 00:33:23] we used the Geom blender [inaudible 00:33:25] a random forest-

Michael Thompson: You run a lot of different models and you want to pick the model that is best for you. I would recommend that you run a variety of models, it also does something else, it trains people to realize that one model is not perfect, all models have got their own problems. Next slide.

Michael Thompson: Our final predictive performance. What we have is...

Michael Thompson: Are you guys getting echo?

Brooke MacCourtney: No we can hear you fine still, Michael, you're good.

Michael Thompson: Okay, our final predictive performance is an R-Squared of .87. I just want to point out, the R-Squared of .87, if we go back and look at the literature, we did twice as good as what some literature was saying. This is pretty cool. And you do see that this histogram follows the common histogram that we had in previous slides. Next slide.

Michael Thompson: And again, we have the chart, we can go on. Next slide.

Michael Thompson: I want to point this one out, only because we created this dashboard, this dashboard was created because we not only want to go through and say, here's expected ED arrivals, here's expected census volume, here's expected admission. We've found by giving numbers and graphs was probably not enough. And so the data science team created a way to describe the numbers in human readable form. Let's have a look at admission prediction points. 187. Above that you see words, Occupancy is predicted to be at 94% and 93% on May 31st at 7:00am for a surge in that department. These natural language words were produced by our model to describe in human readable form what we are saying. What it did for us; it allowed the individuals using it to feel like a real human was talking to them rather than a machine. So it's pretty cool, and I recommend you do the same. Next slide.

Michael Thompson: Then we have prediction accuracy, we produce these on a regular basis. You can go through the slides and take a look at them, but you get an idea of what they look like. 93% on average over the year. Next slide.

Michael Thompson: As much as we want people to read our graphs, as much as we love creating our dashboards, people don't look at them. So we said, "how else can we do this?". What we did, is we're watching the news at night, and we saw the weatherman who said, "Hey, there's a likely chance of rain tomorrow. There's an 80% chance of sunshine tomorrow." So maybe we should create that kind of thing too? So rather than give people charts and graphs and dashboard, we created emails like the one to the right. And it says, tomorrow, on Wednesday, September 5th we sent Thursday 6th, there's a low likelihood that the census would exceed 850. Why 850? Because 850 is when we start having our census planning and surge planning meetings. On Friday, there is a medium chance. On Saturday, there is a low chance.

Michael Thompson: We found that when we said there's a low chance, a very high chance; People responded. Rather than saying we think this could mean census of 844, here people responded. Next slide.

Brooke MacCourtney: Hey, Michael, real quick. I guess we're getting some reports of some echoing, can you just make sure that you laptop sound is turned down now that you're speaking through your phone.

Michael Thompson: Has that worked?

Brooke MacCourtney: I'm hearing you fine, I just I'm getting reports from attendees that there is some echoing. I'm just trying to figure how to fix that.

Michael Thompson: Let me know if you hear any echoes.

Brooke MacCourtney: Yeah, I think so, I'll let you know. We'll go to the next slide.

Michael Thompson: Okay. This next slide is- Of course I'm getting an echo.

Brooke MacCourtney: I'd just make sure that the sounds down on your laptop now that you're speaking through the phone.

Michael Thompson: Yeah, it is.

Brooke MacCourtney: Okay, we'll just keep going.

Brooke MacCourtney: Again, apologies everyone, we're doing our best.

Michael Thompson: What we do with [inaudible 00:39:34] we have low, medium, high or very high census predictions. We ask the leaders of the [inaudible 00:39:41] what would they do when they have a very high? Well if it's very high they have to discharge patients to holding, they might [inaudible 00:39:50] they might [inaudible 00:39:52] centered on the emergency. So what this was, was a translation between our predictions, and actions that were taking place. Next slide.

Michael Thompson: And finally, what we wanted to do was go through for Ease of Use and Rapid Data Refresh. We actually created a mobile app, this is an example of the mobile app that was created. Next slide.

Michael Thompson: This is the stop where we're at. The census prediction moves into a [inaudible 00:40:34] models. The hospital flow app is when we went through and started looking where people were flowing. Transfer turn app, they start [inaudible 00:40:42] patient being boarded. They start looking at OR optimization. One predict [inaudible 00:40:49] project moved into another. Next slide.

Michael Thompson: Finally results were: We saw discharge lounge [inaudible 00:41:04] used. We saw 50 percentage point increase in discharges before 11:00am. We saw reduces in ED boarding. We saw better staffing levels. We saw happier individuals, and we actually saw people trusting the model. Next slide.

Michael Thompson: Our lessons learned in this hospital flow is not a single department problem. We must involve operational and clinical leaders. Interpretation of the data may vary, therefore you must help them interpret the data. And you must have passionate/incredible champions.

Michael Thompson: Let's talk about the positive effect of COVID-19. We, like many hospitals, are experiencing patients [inaudible 00:41:54] California. And we've actually had to go through and figure out, what our expected volume is going to be for COVID-19 patients combined with the volume that we're going to have in normal demand. So we've used our prediction models to figure out, what will the census be tomorrow and the next three days. Oddly enough, while we were planning for a very high census because we were running very high earlier on, what we have figured out is that we've shut down the surgeries. We've had less patients arriving because they've been staying away from the hospital [inaudible 00:42:33] admission. But we still have a very stressed environment, but at least we're still able to more accurately predict the number of patients we expect to have in census tomorrow and the next three days. Next slide.

Michael Thompson: One of the lessons learned that we have is that you should not overestimate the amount of [inaudible 00:43:04] that our physicians and clinicians have. When creating a predictive analytic model involve them from every step of the process. Have frequent meetings to keep them involved, and let them use their insight into your models. And our belief is that datanalytics can assist humans, but doesn't replace humans. And we've found out that data science is not a singular event.

Michael Thompson: So when we go through and start creating our predictive analytics model, a lot of data scientists will put the model in production and let it run. We've found that we're constantly changing environments. Therefore, lets take a look at our models, tweak them as current events are occurring. So for example, we have the models run everyday, and then we compare the error rates of the models. Now if the error rates are getting too high, the data scientist has an email that goes out to a group of people on our team to say we must intervene with the model.

Michael Thompson: So for example, this week, we expect to have everything with normal, a census in the high 800s, maybe low 900s. Strong models were predicting that, but we knew because we had canceled surgeries that were going to occur in the future and we knew that we had patients that were staying away, our models, although they were trying to adjust, didn't adjust well enough. And so we had some error rates happen on the weekend, on Saturday an Sunday, where our models were over predicting. It sends an email to the data scientist, the data scientist then takes a look at the model, and intervenes before the data goes out. Not all models are correct, all models are wrong and all models will change in accuracy. They have a data scientist tend to care and feed and watch the model. Next slide.

Michael Thompson: The final slide is that one predictive model will feed into all predictive models so although our original thing was predict census, then we had ED arrivals and we had unscheduled visits and we had expected discharges, which were an ensemble of models. You'll find that this is a combination of a variety of models that can form together. And you'll notice we also have patient experience updating and wonder, "What's patient experience got to do with it?". Well, we found that ED[inaudible 00:45:43] has the impact on patient experience, and that's feeding into when a patients being boarded, remember one of our models was predicting [inaudible 00:45:51] boarding. But also [inaudible 00:45:53] on likelihood of patient satisfaction.

Michael Thompson: So I know this was a strange Webinar, we had some technical difficulties and quite frankly I'm dealing with a lot of different on-demand [inaudible 00:46:06] projects. I appreciate for those people who have stayed with me through this conversation, you're welcome to send me an email and talk further. And you'll have this line, but I'm going to turn this presentation over to [inaudible 00:46:17] now.

Brooke MacCourtney: Thank you Michael, thanks for sticking with us with all those issues we were dealing with, and we appreciate everyone staying on the line. We will send the slides out, you'll be able to look through those, and we can help answer questions through Michael if you have questions. We'll see if we can send you this recording and we'll see if we're able to give you a recording, at least you can hear some of it if you missed some.

Brooke MacCourtney: So now I'm going to turn the time over to Holly Rimmasch, she is our Chief Clinical Officer at HealthCatalyst. And she is going to be sharing some information about the COVID-19 analytic solutions that HealthCatalyst is developing so I'll go ahead and turn the time over to Holly.

Holly Rimmasch: Thanks Brooke. If you could just go to the next slide. We have been working internally to figure out how we can help our providers and our clients we work with. We just wanted to share a couple of ideas.

Holly Rimmasch: I will start with saying that we are doing whatever we can to help our clients [inaudible 00:47:21] even in our last comment we have helped clients build dashboards around supplies. We've helped them build dashboards around different flows. And actually, looking at using some of our trackers like leading wisely to help create transparency for leadership that is at the site. Another one I heard this morning is we're using one of our population help tools which identifies cohorts to actually start tracking patients that are even tested for the COVID. I think there's a lot of ways we're thinking about it.

Holly Rimmasch: I'm going to share three different idea that we wanted to share with you so that you understand some of the things we're working on. And over time, as you learn more, we want to actually help share information across clients and the general community in being a good partner.

Holly Rimmasch: So the first area that we really put some resources here is the Patient Contact Tracer, and this is really a very simple report built on identifying locations and healthcare facilities where patients who test positive have been and actually understanding the flow. We're seeing a lot of manual work being done for if a patient has tested positive actually going back and finding where they've been in the facility. We recognize that many patients who are being tested coming into the door but we also understand that even after this initial traunche, and we were not sure how much we were going to turn this curb, we'll want to track and monitor this over time. So the ability to create a very simple report that does that.

Holly Rimmasch: The next part is the Staff Contact Tracer which really after we've identified someone who has been tested positive, we actually are creating a tabular report where it identifies the staff members that have documented on that patient. Realizing this is not perfect, but it really helps decrease the manual process. In addition to that we're bringing in order sets, things that LAPs have been done. Image unique has been done and folks who've even gone through the ED, the process is using ED, so the idea is to really help them understand how to start tracking back and understanding we are doing it to minimize their risk.

Holly Rimmasch: The next slide, actually is centered around-

Holly Rimmasch: Could you move the slide forward.

Brooke MacCourtney: Yeah I did, are you not able to see?

Holly Rimmasch: Oh, sorry. We got it.

Holly Rimmasch: Is really centered around starting to think about surveillance. We have a patient safety monitor tool that is built for syndromic surveillance. The idea here is that we actually have triggers that are going in the background and understanding, they're running at about 200 triggers currently, that are looking at things like flu and infection and [inaudible 00:50:28] and it's around patient safety. What we're pretty close to being done is actually building these triggers that are more specific to the syndrome around COVID-19. Which means that you've got continual monitoring of all your patients in the facility that we start looking at that. It creates a list with key indicators about why it created the list. It gives the opportunity for your epidemiologist, your clotting nurses, your infectious disease doctors to look at the syndrome. To see, one, We can also pick up if they're tested, but if they haven't been tested, should we test this? And it also helps us understand the syndrome about the incidents of how it's happening.

Holly Rimmasch: Now this can happen for COVID but it also could happen for flu or other types of things. We recognize, in the beginning of this that there is a lot of patients that are in overload, but as we think about the next few months we will want to continue to serve out patients that are coming in the hospital.

Holly Rimmasch: The next step is actually taking more of a public health view where locally this tool allows you to actually hotspot of geomap either the symptoms like high fever, or it could be the whole syndrome to look by Zip code and putting it on a map where you can start to see this area is having the higher incidents of this and is there anything we should be doing?

Holly Rimmasch: We've talked to some of our clients, this is something they feel is really important, and the ability to build this in a tool that has an ongoing ability to look at this. I think particularly where many experts are saying this could be seasonal as we think about houses coming back. In the summer it could drop and in the winter it might pick up again. These are tools that really help you in the ongoing way of managing and understanding the symptom as it goes through and potential incidents of COVID.

Holly Rimmasch: And the last area I wanted to talk about is Staff Augmentation Support. These are areas where we can provide additional trained analytic data science domain expert staff to respond to increased demand. We're seeing this at a couple of our clients where they have a lot of reports to develop and so we're stepping in to do that. We had Dan Burton our CEO send our a press release yesterday that describes these in general, and one thing just to share with you all on the phone is for our current clients these are part of the package and we will include them and work really hard to get any of the things that we know to help in their organizations. And second of all, if it's something that we need to expand services to do some of this work, we're actually discounting our rate through the end of 2020 to be a good partner.

Holly Rimmasch: These are just briefly some of the areas that we're focusing in. As we learn more we will continue to try to share through the other patient safety collaborative that we're sharing through and trying to keep close tabs with our clients who are incredibly busy hospital healthcare systems and recognize that we don't want to do anything that causes undue stress but we want to be helpful in this process. So that's basically the overview, and we'll turn it back to you.