

Narrator: Welcome to the healthcare.ai weekly live broadcast with the Health Catalyst Data Science Team where we discuss the latest in machine learning topics with hands-on examples. Here's your host, Levi Thatcher.

Taylor: Howdy. We're joining you here this week. And here's what we're going to do today. Welcome to this week in machine learning.

First of all, a couple of administrative of things, we're going to be doing a lot of screen sharing today including some code tutorials. So if you're on YouTube, on the bottom right, you'll see the little gear. Bump your resolution up so that you can read the screens we're going to show you. That will help you a lot. Also, be sure to subscribe.

I did not introduce myself. This is Mike Mastanduno. Taylor Miller. We're here on the Health Catalyst Data Science Team and excited to be here today.

Just to give you a quick heads up of what we're doing today. We're going to talk about Slack. We're going to talk about some really great news articles. We're going to go over some tips and tricks. And then we are going to discuss Chat Q&A. I can't read the cue cards behind this camera.

Mike: That's all right.

Taylor: What's Slack, Mike?

Mike: Well, Taylor, thanks for the intro. And yeah, it's great to be here. You'll notice Levi is not here. He's on vacation today so we'll do the best we can without our fearless leader.

So slack, what is slack? Slack is a messaging app. And here at the healthcare.ai broadcast, we're excited to announce that we've got a Slack channel and a Slack group where members of this community can participate in discussion all week, every week. So how do you get started with slack, Taylor?

Taylor: So jump onto our website healthcare.ai. At the top, you will see a link that says "Slack". Click on that. Throw in your email address. You'll be sent an invitation automatically. Click on that and you'll be in the slack channel. And we'll be there to discuss with you machine learning and healthcare questions, ideas. We'd love to get some interactivity going.

Mike: Yeah. So the primary goal of the Slack is we want to be able to interact with you during the week. So if you have questions for us during the week, you don't have to send healthcare.ai an email. You can hit us on Slack or you can post data there. You can post questions for the community and we'll hopefully be able to start working through your problems together using that as a medium, so really excited to have that. I think it's going to be great.

Taylor: Yeah, awesome.

Mike: So, in the news. There was a fantastic article this week in the New Yorker about AI in medicine. And so, when I saw that come through my email, I got excited about it. And when I saw who the author was I was even more excited about it.

Taylor: Why is that?

Mike: Well, so the author is one of my favorite writers, Siddharta Mukherjee. I've probably just butchered his name but it's a hard one. And he's known for writing a Pulitzer Prize-winning novel on the history of cancer called the Emperor of All Maladies. He's a physician – practicing, and he's a fantastic writer, so I was really excited to read this article about AI. And, man, it did not disappoint. You read it Taylor, what did you think of it?

Taylor: It was super good. The thing that I like, his writing is very approachable, so it's not intimidating if you don't have healthcare or machine learning background. He tells it as a kind of a series of stories and they sort of unfold. And you get to experience these issues that we're dealing with in healthcare and machine learning.

Mike: Great.

Yeah, so it's a pretty long article but I think it's worth kind of stepping through his main points and kind of how he felt about it. And then you can go read it at your convenience if you'd like. But it talks about how one of the values of artificial intelligence and machine learning is that -- that's the machine learning algorithm that learns “how” versus learning “that”.

And so, in the context of riding a bike, I could tell Taylor that in order to ride a bike you have to sit on the seat.

Taylor: Okay, check.

Mike: And you have to turn the pedals.

Taylor: Got it.

Mike: And you have to be going fast enough.

Taylor: Got it.

Mike: That's kind of like you've learnt the rules of riding a bike but would you be able to ride a bike with that information?

Taylor: Not at that point, so no.

Mike: What? Where's the disconnect?

Taylor: I think the disconnect there comes in the, you know, that if they embody that physical-ness of understanding that that bump might be bigger than it seems as you're going down fast towards it and learning to balance and so forth.

Mike: I think that may have been a dig at me, Taylor, but I'll take it. But what you're getting at is learning how, you know?

Taylor: Yes.

Mike: It's like this innate knowledge.

And we found that the deep-learning algorithms can learn how to do something, kind of the same way a toddler might learn to recognize a dog from a cat. They don't go like, "Yup, pointy ears, furry tail – dog."

Taylor: Yup.

Mike: They just look at it and say "dog". They know it's a dog. And so, that's kind of how pattern recognition works.

There was a study in a hospital where the researchers put radiologists in an MRI and watched their brains as they read lung scans. So if you look at the slide, you can see this image on the left kind of shows you what a lung nodule might look like. And that's trivial for a radiologist to pick out. And in fact, it's the exact same activation in the brain as when you're trying to distinguish a rhinoceros from a cow.

Taylor: Yes.

And in that study, it's actually pretty interesting. First, just the recursive nature of the study. Okay, we're going to put imaging professionals in an imaging device and image their brains while they look at images - just amazing. So they showed them those radiology images and they flashed also line drawings of animals and line drawings of letters and they were asked to, "Whatever picture comes up, just quickly tell us what it is."

Mike: Yeah.

And so, you know, they saw that's kind of the learned how. That's how radiology works. So we think maybe that could be applied to deep learning as well.

And so, then they go through. And say like—

Okay, this guy at Stanford who ran an artificial intelligence lab, Sebastian's Thrun, he moved on to GoogleX to work on their self-driving car project and a whole bunch of other things in healthcare. But one thing his team did was, they built a deep learning model to distinguish melanoma from other benign skin conditions.

And so, melanoma is a fairly deadly type of skin cancer. But they came up with an image data set of about a 150,000 images and trained it to recognize melanoma. And they found that it gave more correct answers than experts. But then it also was able to find more melanoma than experts could, so it kind of wins on both fronts.

And so, from that, you have these computer scientists who are saying, "With evidence like that, we think that deep learning is going to completely replace doctors because that's exactly what a doctor is doing to diagnose your melanoma."

Taylor: I love this gentleman's quote. This is a Jeffrey Thurn at Computer Science of Toronto. And he said, "Speaking to a group of hospitalists, they should stop training radiologists now." And he argues that in five years, maybe 10, the deep learning algorithms for image analysis will surpass human perception.

Mike: That's crazy.

Taylor: Yeah, it's quite a statement.

Mike: The article concludes with the author kind of following around a very well-respected dermatologist in her clinic where she sees about 50 patients a day. And that comes out to being about 200,000 cases a year. So with evidence like that you might say, "Well, if a deep learning model can analyze 150,000 cases in three months, how is that not going to replace a clinician?"

And the answer comes from following the clinician around. And I think this is where really the light bulb came on for me. And it was it was so cool. In the beginning, we're talking about learning how and learning that. But in the end, we're not talking about how medicine is pattern recognition. That's just a small part of it. The thing doctors are really doing is learning "why". You say, "I have a rash." And they say, "Okay. Well, your rash looks like condition X but let's figure out how you got that and what we can do about it." And that's the part of the job that's definitely not going to get replaced.

Taylor: Absolutely.

Mike: So I think the future is bright for both AI and clinicians.

Taylor: Yeah and the author makes a great analogy to using a good tool. A good tool enables a craftsman to produce better or more work. And that's kind of the end approach of the article, that it was a big takeaway for me is this is not necessarily a competition thing, it is enabling better medicine to be practiced.

Mike: Yeah and that's a great point, Taylor. I think that's definitely a central point of the article so. But check it out. It's on the New Yorker – Mukherjee. His book, it's

also amazing. If we could get them on this show that would definitely make this show more popular. Yeah, I hope you check it out.

So we're on to the meat of the program. What are we talking about today, Taylor?

Taylor: So, today, we're going to go over some tips and tricks for talking about putting your machine learning model in production, some of the troubles we have, some solutions that we've come up with. We'd we love your feedback, and questions, and ideas on those specific points and others.

Mike: Great. So, let's see. Over the past weeks, we've gone through R in our studio. We've talked about how to install healthcare.ai. We've shown some examples from healthcare.ai. So, if my screen is up, we can we can start to step through some of those and just pick up where we left off.

We've been talking a lot about model development. But if your model's going to actually go into the wild, you need to be able to deploy it which is the other big portion of healthcare.ai.

So I've opened up in our studio console. I'm going to load the package.

And tell us on the chat, have you had problems with this installation? Have you been able to do it? What's preventing you from doing it? Let us know. We'd love to chat about it with you.

So we're going to load the package. And as before, we can do a "?healthcareai" to bring up the docs.

And so, in previous episodes, we've kind of been looking through the LassoDevelopment and RandomForestDevelopment.

But let's go into the deployment and see what that does. So this deploys a production-ready predictive Lasso model. And so, you can either have a saved model in your memory or you can train it right now on the spot. I'm just going to grab this example and copy it into a blank R script, which is here. And then we can run it by either highlighting everything and hitting the "run" button or we can go down to the bottom and use control+Alt+B to run everything up to our cursor.

So, I've run all this code which is going to query my local host database. And it's the it's the same diabetes example. It's just housed in a SQL database. We're going to add it to a data frame, build the supervised model deployment parameters, train the Lasso model. And then we're going to actually deploy it.

Let me make that a little bigger. So just as before, we're building that Lasso model. And this is a regression model so we're interested in intercept. And that's

our performance metric. And you can see we still get those same feature importance coefficients that we get out of other Lasso models.

And then finally, it tells us that the SQL server insert was successful. So what does that mean? Let's go check SQL server and take a look. So if I open SQL server and connect to my local host environment, go down to the databases. I wrote it to the SAM database. Let's look at that table in there. So, regression base, that one is called. We will select the top thousand rows. And we're just going to order - "order by column number three" which is the date it was written, in descending order. So even, we'll get a little SQL this week. Good for us.

So if I execute that query, we see that on April 6th at 07:19 UTC, which is right now-- sorry, 7:11 which was a minute ago, we see that we wrote some things to the table. So this is kind of the way you would test whether or not your model is taking those new rows generating predictions with the prediction probability number, the factor 1, 2 and 3, most important features to get that number. And that's where we are.

So, how do we make sure our performance in the wild, in this in this SQL environment, is just as good as it is when we're in development and kind of playing with stuff locally? And so, the first one is that the production environment is different. And so, predictions can be made multiple times for the same patient. If Joe is an inpatient and he's there for four days, he's going to get a new prediction every day, if that's how you set up your model. So you need to think about which of those predictions you want to show to the clinicians. You need to think about how you want to report that in your performance metrics later on.

Taylor: Right.

Mike: Do you use the most recent prediction? Do you use the first one?

Taylor: Sometimes—

Mike: Go ahead.

Taylor: Yeah, maybe it's helpful for a clinician to see a trend of a particular risk as it's gone through the day. And then maybe that could be correlated with some kind of either intervention or changing the patient's status and so forth. We'd love you're your thoughts on how you would surface multiple predictions over time to a clinician that would make sense without bogging them down with extra information that's going to slow them down.

Mike: Mm-hmm, definitely.

And we're planning on doing a whole episode on visualization so I won't to steal Taylor's thunder on that one.

But another thing that's different is that in the production environment your data sources could be really different. What are the complications of that, Taylor?

Taylor: So this is a problem that we've seen quite a few times when we go to deploy a model. When we're building the model we have retrospective data. Typically, we have pretty good quality data and most of the data is there. So when we build our model, we tune, we create it into our feature engineering. We tune our model. We've got a decent AUC. We ship it out into the wild. And guess what, the accuracy and AUC - the performance goes down.

So we've looked at this a lot and would love ideas here. One of the problems that we've found is that, very often, fields are not populated when you would expect them to be - particularly compared to your retrospective data, which typically is fully baked.

Mike: Yup.

Taylor: So one of the things that we've done is to build a feature in healthcare.ai, both into the Python and soon-to-be-released the R packages, that we call the Feature Availability Profiler.

Do you have a picture of what this looks like?

Mike: Yes, so I have a picture of what comes up when you use that.

Taylor: Yes, so it's an exploratory tool that you run on your data set. It's typically run against a production data set - it's read-only, and it grabs a snapshot of the existing data. And, particularly, you want to filter it to patients who are still in the hospital. So it's a snapshot and then we're going to look at as a patient is in the hospital longer - theoretically, we should see more of those fields be filled out.

Mike: Ah, okay.

Taylor: And what this allows you to do is if you can see some problems—

Mike: So do you mean like as a patient gets lab work done—

Taylor: Yes.

Mike: --those fields would get populated?

Taylor: Correct.

Mike: But they're not going to be there right at the beginning?

Taylor: They may not be. And with retrospective data, it's hard to tell if that's the case.

Mike: Ah, okay.

Taylor: So this is a tool that allows you to look at that and assess your data and maybe make some choices. "Okay, do I want to make/suggest some kind of intervention?" If this field is very important to the clinician, "Hey, we really can make a better prediction if we have this data quicker" which is hard. Or "maybe, I can drop this field out of the model" which might make it less accurate.

Mike: Yeah.

Taylor: Or another—

Mike: If that field's not getting put in there, until two weeks after the patient leaves, no point using it, right?

Taylor: Yeah, no point.

Or another approach is to delay the prediction. Maybe, instead of making that prediction in the first six hours when, if you look at this feature profiler you're going to find out that the data is typically there by 12:00, push that prediction back until 12 hours post admit, and then you can rely on that being a better prediction.

We'd love your thoughts on that kind of a problem. And let us know how this tool works or doesn't work for you.

Mike: Great.

So then, the last the last tip and trick we had-- and these are all huge topics--

Taylor: Certainly.

Mike: --we could we could talk for hours on each one of these.

The last one is just, "How do you make sure your model is being adopted and that you are improving outcomes?" because that's the whole goal of the thing. So how do you make sure your model's being adopted?

Taylor: Well, at some point, you've got to surface that prediction or that data to the clinician. There's lots and lots of ways. That can be done through visualizations, through data points, through color schemes. You need to think about that in your clinical workflow and come up with something that makes sense.

Mike: I think that's a great point. Workflow is so important because if it's not in the workflow, clinician doesn't have time to use it, you know?

Taylor: Yes.

Mike: So let's say, we've built this great model and we've taken all these tips and tricks and our model performance in the wild is really great, and it's being used to improve outcomes. That's awesome. So what's going to happen to the model performance over time if it's being used to improve outcomes?

Taylor: So this is really perplexing and counterintuitive.

So let's say we've got a beautiful predictive model. It's predicting some intervention that will help patients with some outcome. Maybe we're trying to reduce infection and we're suggesting an intervention that will help the clinicians decrease infections. So if we predict that patient A is going to have an infection and an intervention is then done on that patient preventing that infection, then that means inherently predictions will get worse as the interventions get better.

Mike: Okay.

Taylor: So we have this very tricky balance of "we want the model to be very accurate" but that's not our goal. Our goal here isn't to make predictive things accurate. Our goal is to help people get better medical care.

Mike: Well, we want to do both, right?

Taylor: We do but that's the tool to get better care.

Mike: You're right, exactly.

Taylor: So at some point, I think in the future as we, as a community, figure this out, we're going to have to fold the interventions back into the model as another predictive factor or we're going to have to come up with some way to re-train over time.

Mike: Yeah.

Taylor: Very, very tricky.

Mike: And one of the hot topics on our plates during the week has been "How do we measure our outcomes improvements based on the machine learning tools we're creating?" And the best one we've come up with for that is to run kind of the equivalent the equivalent of a clinical trial where we have a control group and a test group. And the test group is getting machine learning intervention. The control group is not. When I when I say it like that, it sounds pretty simple but there's a lot of pretty difficult considerations we have to take in mind. There's actually going to be a blog post on it coming out within the next couple days so you could check that out at healthcare.ai/blog.

Taylor: Yes.

Mike: And let us know your thoughts on Slack or in the YouTube chat.

Taylor: We've got a couple interesting comments in the chat right now. Paul is suggesting that "Oftentimes, the EHR interfaces need improvement to help the clinicians get data into that EHR quicker." And that's a great point. I mean, EHRs are not known for their amazing design and that'd be a great thing to work with your clinical technology staff - to impact that. Another point that Paul makes is that, "Oftentimes, a lot of that data is buried in clinical notes." Buried in unstructured data.

Mike: Yeah.

Taylor: These aren't discrete fields you can point your model.

Now, that's where we would love to be eventually with NLP solutions. That's natural language processing, where you're able to extract information from those unstructured notes. We are definitely nowhere close to that.

Mike: So that's kind of like-- an example would be, in a clinician's notes, having a machine be able to tell the difference between patient's father is a smoker and the patient is a smoker.

Taylor: Right.

Mike: Right. You can't just search for smoker.

Taylor: Right.

Mike: So that's really interesting.

Taylor: It is.

Mike: I've kind of been toying with the idea of like, "Could we use machine learning to pre-populate fields for clinicians? So that instead of having to select the drop-down menu and find the right order set, it's kind of like they're sorted by how useful they would be—

Taylor: Yes.

Mike: --kind of thing. And I think that's just kind of operational machine learning that's really interesting.

Taylor: It is.

I'm kind of a pie-in-the-sky person. Eventually, I can imagine some beautiful future where -- and I think, where there isn't necessarily a UI where you're pointing and clicking through all these terrible things, terrible radio buttons and check boxes everywhere. But it's more of a conversation. And I think things like

Siri, and Google Now, and Alexa are showing us that it might be possible. It's a small glimpse into that direction.

Mike: Definitely. Yeah, we're totally moving towards that. Can you imagine just a clinician wearing a headset, just kind of dictating the patient's history and--

Taylor: Yup.

Mike: Oh, man.

It's a great field to be in, Taylor. It is a great field to be in.

Taylor: It is. It's exciting.

So, let's see, just checking the chat here to see if there's anything else interesting. Paul's mentioning we should check out dynamic machine learning. We will. That sounds very interesting.

Let's see, anything else we need to address today?

Mike: Yeah, we just wanted to give you another reminder to subscribe to our channel. And we want you to join the Slack channel if you're interested as well. And one thing we're really excited about using this, at least the slack channel for, is getting consensus on what you guys want to hear about because, I'll be honest, we keep topics maybe a month in advance but if there's demand for us to give a show on something that seems interesting, we'd be happy to do it so.

Taylor: Absolutely, absolutely.

Mike: Please give us feedback on how we can improve what we can talk about.

And, next week, we're going to be talking about open healthcare data sets. So you've seen the model development. You've seen our studio. You've seen the model deployment. And now, we're going to talk about where we can find great data sets to start tinkering with healthcare data.

Taylor: And that's really exciting because healthcare data is notoriously difficult to get.

Mike: It is.

Taylor: So this will be an exciting topic.

Mike: Yeah, I can't wait.

And please join the community.

Thank you so much for watching. And we'll see you next week.

Taylor: Thanks so much.