WEBVTT

00:00:00.000 --> 00:00:00.000
there we go!

00:00:00.000 --> 00:00:07.000
And Michael, you can start maybe one or 2 min after

00:00:07.000 --> 00:00:23.000
Alright. so recording has started

00:00:23.000 --> 00:00:42.000
And I am starting the webinar now.

00:00:42.000 --> 00:01:12.000
Webinar started. See the participant number climbing

00:01:15.000 --> 00:01:43.000
Alright, we're gonna wait just another minute or so so people can find their seats

00:01:43.000 --> 00:01:54.000
Alright, the the participant number is kind of leveling off a little bit, or at least it's not changing as passage that was changing before.

00:01:54.000 --> 00:01:59.000
So we're gonna get started. so hi everyone welcome to our first size.

00:01:59.000 --> 00:02:05.000
Distinguished lecture for the 2,022 2,023 academic year.

00:02:05.000 --> 00:02:10.000
My name is Michael Litman. i'm the division director for information and intelligent systems.

00:02:10.000 --> 00:02:21.000
I joined Nsf this past july as a rotator from Brown University, and I am delighted to introduce today's speaker, who has done some truly paradigm breaking research.

00:02:21.000 --> 00:02:29.000
I first met Cynthia Ruthen, either late in her time at grad school at Princeton, or shortly after, when she joined Mit.

00:02:29.000 --> 00:02:45.000

I think it was the latter, because I remember hearing about her transition from a more classical machine learning theory, kind of research that she was doing to what seemed to me to be more applied work and at Nsf: we like to talk about

00:02:45.000 --> 00:02:48.000
use inspired research, but I think Cynthia was doing some research.

00:02:48.000 --> 00:02:56.000
Inspired use on behalf of a utility service provider, and she came away from the experience, a changed person.

00:02:56.000 --> 00:03:00.000
So machine learning algorithms that she was using were failing to live up to some of their hype.

00:03:00.000 --> 00:03:11.000
This. This is a thing that happens every once in a while, every decade or so so moving forward, she focused on using her algorithm design prowess to make learning algorithms that solved hard problems.

00:03:11.000 --> 00:03:16.000
But at the same time produced rules that were legible to human beings.

00:03:16.000 --> 00:03:24.000
Now. not everyone, possibly not anyone in the machine learning community at the time quite under understood why that was so important.

00:03:24.000 --> 00:03:35.000
But fast forward about a decade. And suddenly machine learning algorithms were being applied to all sorts of real-world problems, and quite often the results were opaque and brittle.

00:03:35.000 --> 00:03:39.000
If only there was some way to understand what rules these algorithms were learning.

00:03:39.000 --> 00:03:45.000
Researchers lamented, Oh, right right! Cynthia was working on that already.

00:03:45.000 --> 00:03:53.000
So earlier this year Cynthia received the squirrel Ai award for artificial intelligence, for the benefit of humanity, for her work.

00:03:53.000 --> 00:04:04.000
I like to think of it, as ai's noble prize. I'm so glad that Cynthia will be kicking off our series, and so, ladies and gentlemen I introduce Cynthia room thank you Michael

00:04:04.000 --> 00:04:14.000
i'm delighted. to have gotten an invitation to do this and your introduction was exactly right, and it's exactly where I was going to start this talk.

00:04:14.000 --> 00:04:25.000
So yeah, I work in interpretable machine learning, and I started working in this field because I was working with the New York City Power Company in my first job after my postdoc and I was trying to provide power reliability issues and I

00:04:25.000 --> 00:04:36.000
realized that more complex machine learning models were not giving me any better predictions than very simple models, and they were very hard to troubleshoot, And so I thought maybe the problem that I was working on was an anomaly, But

00:04:36.000 --> 00:04:44.000
it wasn't like the same thing kept happening over and over again, and so here's Another example, So back in 2,015.

00:04:44.000 --> 00:04:48.000
We wrote this article called Interpretable Classification models for recidivism prediction.

00:04:48.000 --> 00:04:53.000
In this article we use the largest publicly available data set on criminal recidivism.

00:04:53.000 --> 00:05:03.000
Over 33,000 people released from prison, all in the same year, and we used interpretable machine learning tools to predict whether someone would commit one of a bunch of different kinds of crimes after they were released.

00:05:03.000 --> 00:05:08.000
From present. So these are misdemeanors, violent crime, sexual crimes, property, crimes, drug crime.

00:05:08.000 --> 00:05:22.000
You know you name it, and we what we showed was that you don't need a complicated model to predict recidivism and race is not actually a useful variable in predicting whether someone will be arrested for a crime and we were

00:05:22.000 --> 00:05:34.000
pretty surprised when, a few months later, the republica group, came out with this article that said that there was this proprietary model used in the justice system, and that it uses race, and we thought first of all race doesn't even

00:05:34.000 --> 00:05:43.000

predict recidivism. and, second, these types of models should not be proprietary because they determine people's freedom.

00:05:43.000 --> 00:05:49.000
So we were wondering like how accurate is this model Anyway, this compass model that's used throughout the justice system rates widely used.

00:05:49.000 --> 00:05:54.000
How accurate is it so? We can compare the scores from compass from the Republic article.

00:05:54.000 --> 00:06:05.000
We compared the accuracy of those scores to an algorithm that we developed at the time that the data set was released which the algorithm was called corals and corals is a very complicated algorithm.

00:06:05.000 --> 00:06:12.000
But it produces very, very simple models and so i'm going to show you a machine learning model that was produced by corals.

00:06:12.000 --> 00:06:17.000
Corals produces one-sided decision trees. So it took the data from

00:06:17.000 --> 00:06:30.000
Barrack County, Florida, from the Republica article, and it was just basic information about people's criminal history and their age, and some other demographic information, and it produced a model like I said It was so small that it

00:06:30.000 --> 00:06:40.000
fit in the corner of a powerpoint slide so the model, says if you're 19 to 20 years old in your mail predict arrest within 2 years of your compass, gar calculation, else if you're 21 or 20

00:06:40.000 --> 00:06:42.000
2, and you have 2 to 3 prior offenses. Then predict the rest.

00:06:42.000 --> 00:06:47.000
Within 2 years of your compass score calculation Also, if you've more than 3 priors predict the rest. otherwise particular arrest.

00:06:47.000 --> 00:06:53.000
And we thought, Okay, that's a really simple formula could that really be as accurate as compass.

00:06:53.000 --> 00:07:03.000

And surprisingly. it was so this is results from tenfold cross-validation, and you're seeing out of sample each color is a different out of sample fold and you're, seeing that they're about equally

00:07:03.000 --> 00:07:12.000
accurate And so it's not clear to me why, we need proprietary models in the justice system here determining people's freedom.

00:07:12.000 --> 00:07:18.000
And then we thought, Okay, we're not getting any more accurate than you know.

00:07:18.000 --> 00:07:26.000
With these with coral so you know Why, don't we just throw the whole machine learning arsenal at this problem and see if we can get any more accuracy.

00:07:26.000 --> 00:07:37.000
And we could not. So these are some of these are really complicated black box, you know, models like boosted decision trees or support vector machines with radial. basis function kernels.

00:07:37.000 --> 00:07:41.000
And then, you know, on on the other extreme corals, this whole model is right here in the corner.

00:07:41.000 --> 00:07:50.000
Okay, Now there was this huge debate about the algorithmic fairness of compass, but I think it was all just completely misdirected.

00:07:50.000 --> 00:07:54.000
I really think the truth is that we just don't seem to need compass at all.

00:07:54.000 --> 00:07:59.000
So anyway, like I said, this is not the only data set where simple models perform well.

00:07:59.000 --> 00:08:06.000
In fact, it happens very very often, and i'm going to show you another example, and this time i'm going to choose a really high sticks.

00:08:06.000 --> 00:08:15.000
Example in medicine. So let's say that you have an aneurysm, and it bursts so that you have a hemorrhage in your brain.

00:08:15.000 --> 00:08:21.000
So this is blood leaking into your brain, and At that point you are in pretty serious trouble.

00:08:21.000 --> 00:08:29.000
So you would go to the hospital and get emergency surgery and be placed in the intensive care unit, where, eg. monitors would be put all over your head.

00:08:29.000 --> 00:08:36.000
Detecting for is your like activity for the possibility that you might have a seizure, because these seizures are common and critically ill patients.

00:08:36.000 --> 00:08:39.000
About 20% of patients get them seizures cause brain damage.

00:08:39.000 --> 00:08:44.000
They cause death, and the only way to detect seizure like activity is with E. G.

00:08:44.000 --> 00:08:49.000
So you know it's not like patients are shaking this is all you know inside your head.

00:08:49.000 --> 00:09:01.000
Now, if this was you there's a reasonable chance that at that point doctors would score your risk for seizure, using the 2 helps to be score that we created which is so small that if it's in the corner of a

00:09:01.000 --> 00:09:06.000
Powerpoint slide out. the name of the model comes from so it's 2 helps to be 2 H.

00:09:06.000 --> 00:09:13.000
E. Lps, and then 2 points for the B. which allows the doctors to memorize the whole model just by knowing its name.

00:09:13.000 --> 00:09:20.000
And this model. This is actually a machine learning model just so you know it's.

00:09:20.000 --> 00:09:26.000
I know it looks like a rule of thumb that someone made up, but it's actually the full-blown machine learning model.

00:09:26.000 --> 00:09:28.000
All of these thresholds like the 2 Hertz.

00:09:28.000 --> 00:09:31.000
And then this 2 points, and all the one points, all in the selection of variables.

00:09:31.000 --> 00:09:36.000
All of that was done by a machine learning algorithm this algorithm.

00:09:36.000 --> 00:09:43.000
This model is just as accurate as black box models for this data set it doesn't force you to trust it like a black box.

00:09:43.000 --> 00:09:47.000
Doctors can decide themselves whether they want to trust it that's the benefit of interpretability.

00:09:47.000 --> 00:09:53.000
It allows you to decide whether you trust something it's led to in a validation study.

00:09:53.000 --> 00:10:07.000
Back in 2,018. It led to a substantial reduction in the duration of Eeg monitoring per patient so, and that allowed the doctors to mind quite a few more patients than they could before which according to the doctors helps reduce

00:10:07.000 --> 00:10:21.000
brain damage and save lives and it's optimized using a very sophisticated algorithm that again, it's It's a it's a hard combinatorial problem to design these kinds of models because you have to figure

00:10:21.000 --> 00:10:27.000
out which features are you going to pick? and then what are the thresholds, and how many points, and so on?

00:10:27.000 --> 00:10:41.000
All right now. so it seems that there's no benefit from complicated models for lots of problems, and so here are a bunch of applications that I've worked on, and for none of them did I require a black box model so it could

00:10:41.000 --> 00:10:49.000
be like in the case of compass. that we're using complicated or proprietary models for high-stakes decisions in society when we don't need them.

00:10:49.000 --> 00:10:53.000
Now, in any case, there is a bit of nuance to what i'm saying here.

00:10:53.000 --> 00:11:03.000
So there are really 2 fundamentally different types of problems that we encounter in machine learning, and these are like 2 totally different fields of machine learning, like the whole thought process is different.

00:11:03.000 --> 00:11:12.000

It's like night and day when you're working with these 2 data types like you actually have to change your whole language when you're you know, when you're switching between them.

00:11:12.000 --> 00:11:19.000
Now, tacular data kind of looks like this so it's where you have a good representation of the data, and all the features are interpretable.

00:11:19.000 --> 00:11:25.000
Raw data is like sound waves, or or images, or large amounts of text.

00:11:25.000 --> 00:11:35.000
And the only technique that's working right now for raw data is neural networks whereas tabular data is really different.

00:11:35.000 --> 00:11:47.000
So tabular data with minor pre-processing. If you're willing to do that all the methods tend to have similar performance, and that includes very sparse models like small decision trees, like the corals, model

00:11:47.000 --> 00:11:59.000
I showed you or scoring systems like the 2 helps to be model that I showed you and if you use no networks on these problems, you generally don't see any benefit and you could you could potentially overfit Now, raw

00:11:59.000 --> 00:12:03.000
data is really fundamentally different. It tends to live on very thin manifolds of future space.

00:12:03.000 --> 00:12:16.000
So, for instance, if you think about the manifold of natural images, if you alter one pixel in an image, you actually are no longer on the natural manifold of images, you're actually out of distribution at that point

00:12:16.000 --> 00:12:22.000
so It's really kind of a different it's like a whole different kind of data like tabular data is not generally not like that.

00:12:22.000 --> 00:12:33.000
It's generally not that sensitive. to changes in the data. So if I change, like allergies or exercise, or something like that feature vector could still be realistic.

00:12:33.000 --> 00:12:38.000
Now we were wondering like why is it that we're not getting any benefit from complicated models here. right?

00:12:38.000 --> 00:12:49.000
It seems like there should be some benefit to adding lots of extra complexity. but it doesn't really happen with tabular data. because as soon as you start adding more complexity, you just over fit and so I have a theory as to why

00:12:49.000 --> 00:12:56.000
this happens, and it's a very simple theory my theory is that there are just lots of good models in tabular data problems.

00:12:56.000 --> 00:13:07.000
So let me explain this. a little bit more and I'm going to go into depth on this paper with Lesia and Ron, and the theory in this paper is called the Rashomon. set theory and the Russian.

00:13:07.000 --> 00:13:09.000
Man set theory is that there are just lots of good models.

00:13:09.000 --> 00:13:19.000
So if you think about the space of all models then the Russian onset theory is that there's just lots of good models like, maybe not like half the models, but like a lot of good models.

00:13:19.000 --> 00:13:24.000
This set should be large enough to contain a ball. Okay, a big ball of good models.

00:13:24.000 --> 00:13:38.000
And then the idea is that if if the set of simple models is a good cover for the set of all models, then as long as this ball is big enough, it's going to contain at least one simpler model, so you have at least one simple

00:13:38.000 --> 00:13:44.000
model. That's also good. And now this idea that simpler models are a good cover.

00:13:44.000 --> 00:13:53.000
I think that's totally reasonable. because for instance sparse decision Trees are a good cover for the set of all trees and trees are universal approximators.

00:13:53.000 --> 00:13:58.000
So I really think it actually makes sense so let's call this set of good models.

00:13:58.000 --> 00:14:03.000
The Rasha onset. right? This is based on this, this Brian name of the.

00:14:03.000 --> 00:14:13.000
He used the name of the japanese movie, because the idea is that there's sort of no single right, you know thing is there's a whole bunch of good explanations right?

00:14:13.000 --> 00:14:20.000
So that's the ideas like the rashoman said is the set of good models and set of all models that have lost.

00:14:20.000 --> 00:14:23.000
That's kind of close to optimal in the data.

00:14:23.000 --> 00:14:30.000
Okay. Now, I claim that this Rashomon set is large and many of the types of problems I consider. Okay.

00:14:30.000 --> 00:14:38.000
Now in this paper that this paper over here we did a lot of very computationally heavy experiments.

00:14:38.000 --> 00:14:46.000
We actually calculated the size of the Rashomon sets, or, like, you know, the ratio of good models to all models.

00:14:46.000 --> 00:14:56.000
We calculated that for decision trees for about 70 different data sets. and then we tried to correlate that size of the Rushman set with lots of different things.

00:14:56.000 --> 00:15:02.000
Okay, and what we found was pretty interesting. All right. So let me show you about the conclusions. From this paper.

00:15:02.000 --> 00:15:10.000
We found that large Rashomon set. so lots of good models are correlated with the existence of simpler models.

00:15:10.000 --> 00:15:20.000
Yeah, Okay, we we thought that would happen. Okay, We also found that large rashmansats are correlated with many different machine learning models. methods.

00:15:20.000 --> 00:15:35.000
Having the same performance. And so when I say different, I mean like really different, like models with different functional forms, like, if you have like support vector machines, you know random forests, and what you know all different functional forums, if they all

00:15:35.000 --> 00:15:39.000
tend to perform well that tends to correlate with having a large rush amongst it.

00:15:39.000 --> 00:15:53.000
Why that? Why does that make any sense Well, if you think about it? If you have many different machine learning models having the same performance, you can think about all of these different machine learning methods having these models with very different functional forms, and they're

00:15:53.000 --> 00:16:02.000
all in the rashomot set well and they're all in the same rashmun set. so the rash month that has to be big enough to accommodate all of these very different models.

00:16:02.000 --> 00:16:10.000
Okay, so that makes sense. And then the third thing we found is that large Rashomon sets are correlated with more label or feature noise.

00:16:10.000 --> 00:16:14.000
And so these are problems, or the outcome is hard to predict.

00:16:14.000 --> 00:16:23.000
So for something like criminal recidivism for instance it's really hard to predict whether someone's going to commit a crime within 2 years of their release from present.

00:16:23.000 --> 00:16:33.000
Like there's just so much randomness in this whole process of that that just creates an intrinsic level of kind of noisiness to the data.

00:16:33.000 --> 00:16:44.000
Now. the implications for this theory if i'm right, and there are just lots of good models for most tabular data problems.

00:16:44.000 --> 00:16:50.000
Then optimizing for simplicity, won't actually sacrifice accuracy.

00:16:50.000 --> 00:17:05.000
Okay, that's the implication and this has huge implications for a lot of high-stakes decisions that are made using data in our society that deeply affect people's lives and that includes criminal justice decisions that determine

00:17:05.000 --> 00:17:16.000
people's freedom. for Loan decisions that determine whether someone whether someone can purchase a home or for medical decisions that determine life or death. Right?

00:17:16.000 --> 00:17:26.000
So if i'm right, then for none of these decisions can we really justify black box models for none of them.

00:17:26.000 --> 00:17:33.000
And so now this theory reveals why we were able to find accurate models for these data sets without losing predictive performance.

00:17:33.000 --> 00:17:36.000
It's because these data sets probably admit many good models right?

00:17:36.000 --> 00:17:43.000
They're all tabular they'll have noise because we're predicting things that are inherently difficult to predict.

00:17:43.000 --> 00:17:49.000
So let me go back for a minute to the results of the fact.

00:17:49.000 --> 00:18:02.000
Paper. and I wanna kind of just zoom in on on this notion that if you have many very different machine learning methods that have the same performance that that tends to correlate with large rashomon sets because that

00:18:02.000 --> 00:18:09.000
result is actually really useful. So what i'm saying here is that you run like a lot of very different machine learning algorithms on the data.

00:18:09.000 --> 00:18:21.000
So you run all these different machine learning algorithms and if you see that you get the same performance, then you can decide whether it's worthwhile to run something more computationally expensive to get a more interpretable model.

00:18:21.000 --> 00:18:24.000
Okay, So let me tell you about a case where we did this.

00:18:24.000 --> 00:18:30.000
And this is what the i'm going to tell you about the data set for the explainable machine learning challenge.

00:18:30.000 --> 00:18:35.000
So Feiko gave us this data set about loan decisions, and they said, Make a black box and explain it.

00:18:35.000 --> 00:18:44.000
And the question is whether we need to do that Okay, So just to give you a little background about this data set about 10,000 loan applications.

00:18:44.000 --> 00:19:00.000
There's a lot of factors. about people's credit history. The best black box accuracy we could get on this data set was around 73%, and I thought, could it be tabular data set where you really need a black

00:19:00.000 --> 00:19:10.000
box. and so I asked my students to do some experiments, and I said, Could you please run lots of different black box algorithms on the data set and see if they all perform about the same.

00:19:10.000 --> 00:19:15.000
And then The students came back to me in 2 days, and they said, Yeah, they all perform the same.

00:19:15.000 --> 00:19:18.000
And so at that point I knew we probably didn't need a black box for the State set.

00:19:18.000 --> 00:19:24.000
Okay, So at the time of this competition we didn't have methods that are as powerful as the ones that we do now.

00:19:24.000 --> 00:19:33.000
So I get the benefit of telling you about these results after developing the algorithms. and I'm going to spare you the details of the first few models we created, and how long it took to do it.

00:19:33.000 --> 00:19:38.000
So you could just see some 2,022 algorithms instead.

00:19:38.000 --> 00:19:48.000
Alright. So the first algorithm i'm going to talk about is fast sparse, and this algorithm produces sparse, generalized additive models which are an alternative to logistic regression.

00:19:48.000 --> 00:19:58.000
And it is fast. And so i'm going to give credit to Ja Chung, Judy and Margo on this project. alright.

00:19:58.000 --> 00:20:08.000
So basically just to remind you, Okay, So the best black box accuracy we can get in this data set is about 73%.

00:20:08.000 --> 00:20:12.000
The best black box. Auc is around Point 8. Okay, So those are the.

00:20:12.000 --> 00:20:21.000

Those are the performance measurements that we want to be able to, you know, maintain after we switch to an interpretable model.

00:20:21.000 --> 00:20:31.000
Now fast Sparse takes less than 20 s to run and It's training and test accuracy are right on par with the best of the black boxes.

00:20:31.000 --> 00:20:34.000
And same with Abc. that's right on par with the best ones.

00:20:34.000 --> 00:20:40.000
And the difference, though, is that for the black boxes I can't write the whole model on a slide whereas for a fast sparse.

00:20:40.000 --> 00:20:47.000
I can show you the entire machine learning model. that It produced because it fits on a Powerpoint slide. all right.

00:20:47.000 --> 00:20:58.000
So here's the model that fast sparse produced It's just right here, and the way you read this is that you get a score for each of the variables and you just add them up okay, and that's the whole calculation you just look

00:20:58.000 --> 00:21:01.000
up the points for each variable, and then you just add them up.

00:21:01.000 --> 00:21:08.000
Okay, so and and that total sum translates into a risk for defaulting on a loan.

00:21:08.000 --> 00:21:13.000
So let's take a look at some of the factors here, so let's look at months since oldest trade open.

00:21:13.000 --> 00:21:24.000
So if all your trades are really recent like they're all within the last 100 months, then you get like a certain number of risk points, you get more risk points. because all your trades are recent then number of satisfactory trades.

00:21:24.000 --> 00:21:32.000
If you haven't. had very many satisfactory trades like, if you've had less than 10 satisfactory trades, then you get more risk points.

00:21:32.000 --> 00:21:40.000
And then this feature seems to be really useful the external risk estimate, and it's particularly valuable between values somewhere around 60 to 80.

00:21:40.000 --> 00:21:49.000
So it's really sensitive to values somewhere in here Okay, So what's really shocking about this model?

00:21:49.000 --> 00:22:00.000
Besides, its simplicity in describing this very difficult benchmark data set is that we created it in under 4 s, and that's That's how fast this algorithm runs on average, it's about 4 s.

00:22:00.000 --> 00:22:05.000
For the status set. Now this competition, as I told you. it.

00:22:05.000 --> 00:22:19.000
They told everyone to create a black box and explain it and he told us that because they didn't know that it was possible to create a model like this at all, and we did it in under 4 s.

00:22:19.000 --> 00:22:30.000
So let me tell you about the machine learning method that created this sparse editive model, and i'm going to put a few equations up.

00:22:30.000 --> 00:22:31.000
And so if you don't like equations don't worry about it.

00:22:31.000 --> 00:22:41.000
They'll disappear in just a minute. but if you like equations hopefully, miss, all tell you kind of enough of the backbone of the algorithm that you'll understand it.

00:22:41.000 --> 00:22:45.000
So I'm going to start with standing sparse logistic regression.

00:22:45.000 --> 00:22:56.000
So this is the logistic loss with l 0 regularization that says, Keep it sparse and keep it accurate. and then we're going to use a linear model which is standard for logistic regression, and then we'll use

00:22:56.000 --> 00:23:00.000
the standard conversion to conditional probabilities for logistic regression.

00:23:00.000 --> 00:23:05.000
Okay. So so far, everything is completely standard. Now, to get a generalized additive model.

00:23:05.000 --> 00:23:12.000
The simplest way to do that is to transform the variables before you do anything as a pre-processing step.

00:23:12.000 --> 00:23:19.000
So, for instance, if you have age as one of your x variables, then you can transform age into lots of dummy variables like that.

00:23:19.000 --> 00:23:23.000
And now you've got a lot more features because you blew up the feature space.

00:23:23.000 --> 00:23:34.000
But now you can create very flexible functions. of age because logistic regression is going to give you a weighted sum of these little step functions, and when you add them together you get you can get something that's pretty flexible, and

00:23:34.000 --> 00:23:45.000
pretty funky. Okay, and like the vast majority of machine learning algorithms to minimize this objective, we're going to use something kind of like coordinate descent.

00:23:45.000 --> 00:23:52.000
Some variation of gradient descent. Right? so coordinate descent handles one coefficient optimizes, one coefficient at a time.

00:23:52.000 --> 00:24:01.000
So one of these w's at a time gets optimized and we're also going to try out lots of feature subsets that seem promising according to the objective.

00:24:01.000 --> 00:24:12.000
Okay, Now we came up with this very sophisticated idea, involving cutting planes and quadratic cuts, and it was very fast for solving this problem, and we were really excited about it.

00:24:12.000 --> 00:24:21.000
But then we came up with way to do it that didn't involve cutting planes, and didn't involve quadratic cuts, and it was really very simple.

00:24:21.000 --> 00:24:27.000
And it was about 5 times faster, and that involves changing the problem slightly.

00:24:27.000 --> 00:24:36.000
So, instead of using the logistic loss, here, i'm going to switch over to the exponential loss which is used in at a boost, the probabilistic model changes a little bit.

00:24:36.000 --> 00:24:43.000
Just adds those 2. So it's almost the same and So now we're doing sparse exponential loss classification instead of logistic regression.

00:24:43.000 --> 00:24:56.000
But like I said it's almost the same but there's one big difference, which is that the exponential loss has an analytical solution for the line search at every step of coordinate descent.

00:24:56.000 --> 00:25:01.000
So what I'm. saying is that you're optimizing one of these W's.

00:25:01.000 --> 00:25:06.000
At each iteration. when you're using features that are dummy variables.

00:25:06.000 --> 00:25:16.000
Then when you're minimizing this thing it's all in one dimension right, every iteration it's all in one dimension, you're here and you've got to get to here and you can do that using a

00:25:16.000 --> 00:25:19.000
formula. You don't have to use an iterative procedure to walk down this thing.

00:25:19.000 --> 00:25:24.000
You can actually get it. You can go directly, and the formula is pretty funky.

00:25:24.000 --> 00:25:28.000
I mean the formula says like you know if something or other equals 0.

00:25:28.000 --> 00:25:37.000
Then do something else, and if it you know it's it's like an if, then kind of logical it's pretty, it's a pretty weird formula, but it's a formula, and it gets you directly from a to

00:25:37.000 --> 00:25:46.000
B. without having to walk and take steps, and so that makes it very, very fast, so we can iterate through these steps very quickly.

00:25:46.000 --> 00:25:49.000
And so at every iteration, we just update one of these W's.

00:25:49.000 --> 00:25:55.000
Using the formula on the previous slide, and we use a priority queue to track, keep track of which W's.

00:25:55.000 --> 00:26:01.000
We want to update in which order and we just keep updating them until we've converged.

00:26:01.000 --> 00:26:05.000
And we get these sparse models really, really quickly.

00:26:05.000 --> 00:26:20.000
And so if I go back to the Fico data set here, the algorithm transformed the data set into 1,917 binary features, and then it iterated through subsets of them and it picked out 21 features

00:26:20.000 --> 00:26:37.000
and picked them out in just under 4 s. So so far i've talked to you about the Rashomon set theory, which is that simpler models exist when there are a large number of almost optimal models and this includes

00:26:37.000 --> 00:26:41.000
the the fact that if you run a lot of machine learning algorithms, they all perform.

00:26:41.000 --> 00:26:48.000
Similarly, it could be because you have a large, rash mindset could be because you have a lot of good models.

00:26:48.000 --> 00:26:59.000
If you try to predict outcomes that are uncertain, you probably have a large rashomon set, and if you do have a large rush onset, algorithms like fast sparse can probably find a sparse accurate

00:26:59.000 --> 00:27:15.000
model, and even on competition data sets we're finding sparse, accurate models that people didn't know existed. and this has huge implications for criminal justice loan decisions and other high-stakes decisions because in

00:27:15.000 --> 00:27:19.000
these cases our theory makes it harder to justify the Use of a black box.

00:27:19.000 --> 00:27:28.000
But i'm not done yet. Okay, so I wanna it's first of all, to point out here that lots of machine learning people do not want to hear about this right?

00:27:28.000 --> 00:27:35.000
They're just not interested. These ideas about producing simple models that's really not what mainstream machine learning has been focusing on right.

00:27:35.000 --> 00:27:40.000
They are, and have always been focusing on building more complex models and mainly for computer vision.

00:27:40.000 --> 00:27:47.000
And they're interested in preventing overfitting whereas what I'm talking about is going in the other direction.

00:27:47.000 --> 00:27:54.000
Right. I I want to know how simple we can go and still maintain performance, and I don't need to prevent overfitting.

00:27:54.000 --> 00:27:59.000
My models are so simple that by statistical learning theory right they don't overfit.

00:27:59.000 --> 00:28:05.000
So it's just really a different perspective on what the goal is for machine learning.

00:28:05.000 --> 00:28:15.000
Then what most people have so it's just the really it's really kind of moving in the opposite direction sort of more complex versus, more simple and trying to maintain performance instead of worrying about overfitting so it's just

00:28:15.000 --> 00:28:23.000
a different way of thinking about things. Okay, So I want to move to the next topic, which is first decision.

00:28:23.000 --> 00:28:30.000
Trees, decision tree algorithms. they've been popular since the very beginning of machine learning.

00:28:30.000 --> 00:28:41.000
And the main problem that's always plagued decision tree algorithms is their lack of optimality, because they've historically been greeting myopic algorithms like cart and C 4.5 and these algorithms construct trees from

00:28:41.000 --> 00:28:44.000
the top downward, and then the greenly print them back afterward.

00:28:44.000 --> 00:28:50.000
And the problem is that if a greedy algorithm chooses the wrong split at the very top of the tree, there's no way to undo it.

00:28:50.000 --> 00:28:55.000
So here, if i'm trying to predict whether i'm gonna get stuck in traffic on my way home from work.

00:28:55.000 --> 00:29:01.000
You know. maybe rain wasn't the first question that I should ask if I want to.

00:29:01.000 --> 00:29:07.000
Small tree, but a greedy algorithm picked it and so now i'm stuck with it, anyway.

00:29:07.000 --> 00:29:17.000
So these greedy methods produce suboptimal treats, but it's hard to improve over the greedy methods, because decision tree optimization is really hard right both both theoretically and practically it's really hard right there's

00:29:17.000 --> 00:29:24.000
a combinatorial explosion in the number of possible trees we could consider, and, in fact, optimal, sparse.

00:29:24.000 --> 00:29:28.000
The problem of finding optimal spread decision traces np-hard.

00:29:28.000 --> 00:29:40.000
It's actually factorial in the number of variables and That's why people have been constructing these greedy trees since the early 1,900 seventys in the beginning of Ai.

00:29:40.000 --> 00:29:52.000
But there's been a lot of Somehow this this area became became a very popular research area lately, for some reason I've been working on it for about 10 years, and the latest method that We've produced is called

00:29:52.000 --> 00:29:56.000
ghost, and it uses dynamic programming with bounds that reduce the search space of trees.

00:29:56.000 --> 00:30:03.000
And I just want to put pictures up of the army of people that we've had working on Ghosts ghost is very fast.

00:30:03.000 --> 00:30:08.000
It's much faster than previous approaches so here I specifically want to mention Hayden and Chudy and Margo Margo.

00:30:08.000 --> 00:30:14.000
I've been working with for a very long time We've done a lot of projects together, and I want to mention Jimmy.

00:30:14.000 --> 00:30:19.000
All these people put a lot of effort into getting this the run the run fast.

00:30:19.000 --> 00:30:24.000
So this is one of the problems that ghost solves.

00:30:24.000 --> 00:30:29.000

And this problem is says, Please try to keep the model accurate, but also try to keep the number of leaves in the tree like small like.

00:30:29.000 --> 00:30:34.000
You want to keep it sparse, but keep it accurate.

00:30:34.000 --> 00:30:38.000
And when we solve this problem to optimality, we get a tree.

00:30:38.000 --> 00:30:43.000
Okay, So this is an example of one of these trees on the Florida rearrest data.

00:30:43.000 --> 00:30:55.000
And so here the tree would say, Oh, if you have more than 3 prior offenses predict arrest within 2 years of your compass, score calculation. Otherwise ask about your age and your number of prayers, and so on.

00:30:55.000 --> 00:31:00.000
Okay. So now ghost is a dynamic programming algorithm.

00:31:00.000 --> 00:31:03.000
And so it deals with lots of subproblems.

00:31:03.000 --> 00:31:08.000
Okay, So let me explain. So to figure out what the optimal split is at the top, You say?

00:31:08.000 --> 00:31:11.000
Well, if I made that split at the top, what would be the optimal split beneath it?

00:31:11.000 --> 00:31:23.000
So you gotta figure out the optimal split beneath it, and then to get that one, you have to figure out the optimal split beneath that and below that and below that, and until you get to a small enough subset of data where you can

00:31:23.000 --> 00:31:25.000
prove that the optimal solution is actually just a leaf.

00:31:25.000 --> 00:31:30.000
Then you pass that information back up to the top, and that helps it.

00:31:30.000 --> 00:31:36.000
Reduce. You know, kind of like how many subproblems it needs to consider.

00:31:36.000 --> 00:31:45.000

So a subproblem is to find an optimal tree for a subset of data represented by a binary vector.

00:31:45.000 --> 00:31:49.000
So we're representing all the subproblems as bit vectors.

00:31:49.000 --> 00:31:56.000
So, for instance, if I think about the very top of the tree where i'm considering all of the data points that is a bit vector.

00:31:56.000 --> 00:32:02.000
Of size n that has all ones in it and so I mean you're, considering all the data points in that subproblem.

00:32:02.000 --> 00:32:14.000
And then, if you make a split, then you're considering, you know half some of the data on one side and the other part of the data on the other side. So, for instance, if I split on some variable I split into 2

00:32:14.000 --> 00:32:19.000
subproblems, And this subproblem has data points.

00:32:19.000 --> 00:32:28.000
1, 3, 4, 6, and 8, and so on, whereas this subproblem has the other data points.

00:32:28.000 --> 00:32:40.000
So. And by keeping all of these subproblems indexed by kind of which data points are in them that allows us to work very quickly with bit vector computation.

00:32:40.000 --> 00:32:54.000
Okay, So we start at the very top ghost starts from this master problem that includes all of the data, and then it constructs a big dependency graph that includes all of the subproblems that it it

00:32:54.000 --> 00:33:04.000
encounters. Okay, So here, at the very top of the dependency graph, it's considering every possible split for the very top of the tree so it could.

00:33:04.000 --> 00:33:09.000
It could turn the whole tree into a leaf that's one possibility.

00:33:09.000 --> 00:33:12.000
Another possibilities that can split on the number of priors.

00:33:12.000 --> 00:33:20.000

It could also split on some age variable, and I could, you know, then try from splitting on priors.

00:33:20.000 --> 00:33:26.000
I could split them priors, and then below, that I might consider all possible things. I would do beneath that right.

00:33:26.000 --> 00:33:30.000
What would I split on beneath that? So these are all possibilities for the very first split?

00:33:30.000 --> 00:33:34.000
And then here's possibilities for the second split and then so on, and so forth.

00:33:34.000 --> 00:33:48.000
Ok. So this dependency graph can get really, really big. But we have a whole bunch of theorems that help reduce the size of this dependency graph, and we also keep track of lower and upper bounds on each of each of these

00:33:48.000 --> 00:33:54.000
subproblems, and then, when something gets changed down at the bottom, here we pass that information.

00:33:54.000 --> 00:34:03.000
We propagate that all back up to the top, and so it can. that allows it to kind of eliminate parts of the graphs where it provably has no optimal solution.

00:34:03.000 --> 00:34:09.000
So these graphs can get really big, but we can print them very, very efficiently.

00:34:09.000 --> 00:34:13.000
So I want to go back to the Phico data set, which is my sort of muse today.

00:34:13.000 --> 00:34:24.000
And this is the decision tree that ghost produces and you can see that external risk estimate which we saw before.

00:34:24.000 --> 00:34:35.000
That's actually a really important feature so here, this tree says if the external risk estimate is like too low or too high, then I can figure out the class like I can get the prediction, and but if it's like

00:34:35.000 --> 00:34:41.000
in between kind of like 67, and 76 then I need to like ask some more questions.

00:34:41.000 --> 00:34:43.000

Okay, so that's like the more subtle stuff.

00:34:43.000 --> 00:34:50.000
So I have to ask about like the percent of trades with the balance and the average months in the file, and so on and so forth.

00:34:50.000 --> 00:34:55.000
And so here that's this is the whole tree by the way this is the whole thing.

00:34:55.000 --> 00:35:01.000
This is 10 leaves, 1, 2, 3, 4, 5 6 7 8, 9, 1010 leaves.

00:35:01.000 --> 00:35:10.000
It's again. The performance is very, very similar to the best of the black boxes, and this was computed in 8.1 s.

00:35:10.000 --> 00:35:22.000
Now, again, just to remind you this is not an easy data set it's a benchmark data set, and we didn't know it was possible to construct a single decision tree with this level of accuracy right in fact, without this

00:35:22.000 --> 00:35:30.000
algorithm. I don't actually know of another way. we could have done this because I didn't know that we could get to this level of sparsity and maintain black box performance.

00:35:30.000 --> 00:35:41.000
And so this is the you know. this is the kind of result. We were aiming for it, and i'm thrilled that, you know, after so many years of working with my team on this project that we actually were able to get

00:35:41.000 --> 00:36:00.000
here alright. so I hope that you're getting closer to understanding the implications of the existence of interpretable models, which really makes it hard to justify using black box models for high-stakes applications But

00:36:00.000 --> 00:36:11.000
There's a whole lot more to this story here and no There's really something that's been eating at me for quite a while, and so I want to tell you about that now.

00:36:11.000 --> 00:36:18.000
Something is broken now with the machine learning world is trying to create with these really complex models.

00:36:18.000 --> 00:36:22.000

That is not clearly not what we need, the world, what the world needs for high-stakes, decisions, right?

00:36:22.000 --> 00:36:24.000
We need models that are trustworthy models.

00:36:24.000 --> 00:36:31.000
People can criticize things people can double, check, especially for high-stakes. decisions, Right?

00:36:31.000 --> 00:36:41.000
These overly complicated models they're not going to cut it but guess what my simple models Won't cut it, either because the whole paradigm of machine learning is wrong for these problems.

00:36:41.000 --> 00:36:49.000
So let me explain. So it's kind of this is kind of the universal paradigm for machine learning.

00:36:49.000 --> 00:36:52.000
So you take your training set. We stick it into an algorithm.

00:36:52.000 --> 00:36:56.000
The algorithm minimizes some kind of regularized lost in the training set.

00:36:56.000 --> 00:37:03.000
It produces a predictive model. Okay, that's the way we do things, and that predictive model could be anything. It could be a decision tree. It could be a random forest.

00:37:03.000 --> 00:37:07.000
It could be a neural network, could be a linear model support vector.

00:37:07.000 --> 00:37:10.000
Machine. Maybe kernel regression general is additive models.

00:37:10.000 --> 00:37:15.000
What it could be, whatever right it's a model this is the standard approach data in model out.

00:37:15.000 --> 00:37:22.000
Okay. And we do this for essentially all machine learning applications, even for self-supervised applications.

00:37:22.000 --> 00:37:25.000
You got you're predicting some kind of y from some kind of X.

00:37:25.000 --> 00:37:34.000

But I claim that this whole paradigm is just fundamentally flawed, and we should just reconsider for high-stakes decisions.

00:37:34.000 --> 00:37:41.000
What we're doing and here's Why, okay, so these are all domain experts that I that I work with.

00:37:41.000 --> 00:37:46.000
Okay, now, Dan Wagner, I worked with on Crime Series detection, which is used in New York.

00:37:46.000 --> 00:37:59.000
And this is if Vidas and Joseph, who I work with on computer, aided mammography, Shao, who I work with on understanding heart measurements from wearable devices like your watch this is Ed and Dave and I

00:37:59.000 --> 00:38:07.000
work with them on understanding the reservoir Hiv Reservoir and Hiv patients.

00:38:07.000 --> 00:38:14.000
And then this is Brandon, who I work with, worked with on the 2 helps to be score, and I work with him, still on understanding how to care for critically.

00:38:14.000 --> 00:38:20.000
Ill patients. Now, the thing Now these people are all very different from each other, right.

00:38:20.000 --> 00:38:29.000
They work on completely different things. But the thing that's common to all of these people is that they've all at some point told me that I was wrong.

00:38:29.000 --> 00:38:34.000
Okay. So when you work with domain experts and you go to great lengths to bring them a model, they can criticize.

00:38:34.000 --> 00:38:38.000
They will do it? They'll say I think something's wrong with this model?

00:38:38.000 --> 00:38:40.000
Can you build one that doesn't depend on this variable in this way?

00:38:40.000 --> 00:38:46.000
Or is there another model that doesn't depend on age so much, or can you incorporate fairness constraints into it and maintain performance? right?

00:38:46.000 --> 00:38:48.000
Or can you just tell me what else is out there like?

00:38:48.000 --> 00:38:54.000
They will ask you questions, and they will tell you why you are wrong now.

00:38:54.000 --> 00:39:04.000
The whole idea. The whole premise of the rashomon set is that there were probably lots of good models, and there are probably lots of good, simple models, too.

00:39:04.000 --> 00:39:11.000
And if that's the case, why, should we not let the domain experts choose between them.

00:39:11.000 --> 00:39:26.000
So I propose a new paradigm for machine learning which is something that's more kind of human facing than standard machine learning which is to hand the user the whole. rashomon set just hand the user lots of good

00:39:26.000 --> 00:39:38.000
models. not just one model. Let them choose ok let them choose which model they want, so they can pick something that doesn't just agree with the data, but also agrees with their dummy knowledge or knowledge, You know that's knowledge of

00:39:38.000 --> 00:39:44.000
the problem that isn't in the data set then instead of the algorithm producing one model.

00:39:44.000 --> 00:39:52.000
It produces lots of bottles. So this is my proposal for a new paradigm for practical machine learning.

00:39:52.000 --> 00:39:57.000
But how do we get this to work? So let me show you our attempt to do this.

00:39:57.000 --> 00:40:09.000
So this paper, which was just recently accepted with Ray Takuya Chudi Margo and Jeff actually solves this problem for sparse decision trees.

00:40:09.000 --> 00:40:14.000
It's an algorithm called tree farms and it produces all almost optimal trees.

00:40:14.000 --> 00:40:17.000
It's pretty amazing like even to find one optimal sparse trees.

00:40:17.000 --> 00:40:23.000
Np-hard. and this thing finds them all in minutes, and sometimes seconds, and it's implemented in ghost.

00:40:23.000 --> 00:40:27.000
So it Leverages ghost's way of representing subproblems as bit vectors and its dependency graph ideas.

00:40:27.000 --> 00:40:34.000
But it also has a really interesting way of keeping track of subproblems like it stores all the trees in an implicit way.

00:40:34.000 --> 00:40:40.000
So it doesn't actually enumerate them but it tells you how to kind of combine different parts of different trees to kind of produce the whole rush amongset.

00:40:40.000 --> 00:40:49.000
So you can loop through it. if you want to so even if the rashoman said, is absolutely huge. you can still store it and work with it.

00:40:49.000 --> 00:41:05.000
So We've been working with this brilliant young human computer interaction expert called Jay. and he actually wrote a lovely interface to tree farm that I'm going to i'm going to actually show you so i'm

00:41:05.000 --> 00:41:10.000
going to stop sharing this screen, and i'm going to share a different screen over here.

00:41:10.000 --> 00:41:19.000
Okay, So the the interface that jay road is called timber track and timber truck is loaded here with the compass data set.

00:41:19.000 --> 00:41:26.000
And so what it is doing is it's showing me all of the trees in the rashoman set for the compass data set.

00:41:26.000 --> 00:41:39.000
So let's say that you and it allows you to kind of walk along a branch of a tree, so let's say he wanted, like a tree with where the top split is where the number of prior crimes is greater than 3 Okay.

00:41:39.000 --> 00:41:56.000
So then you can click that, and then it now it's limiting you to just looking at trees where the number of prior crimes greeted in 3 is at the very top of the of the of the tree okay and like you

00:41:56.000 --> 00:41:59.000
can go through and look at all the trees if you want.

00:41:59.000 --> 00:42:03.000
If you find a tree that you really like you know let's say I like this tree, then I can store it.

00:42:03.000 --> 00:42:08.000
You know I could save it and then like let's that I want my next split to be it's less than 21.

00:42:08.000 --> 00:42:21.000
So I want to separate age less than 21 out, and then maybe I want to have people with no juvenile crimes, and then I can look at all the trees that remain here. and you know if I like a particular tree I can say I like

00:42:21.000 --> 00:42:24.000
this tree, you know, and I can write a little note to myself.

00:42:24.000 --> 00:42:29.000
This tree is cool and do it like that, and then I can also.

00:42:29.000 --> 00:42:38.000
I can also visualize the tree in different ways. So here, if I click this button, then it's showing me how many data points are going down, each branch of the tree.

00:42:38.000 --> 00:42:43.000
So here it's showing me that about 2 over 3 of the data are going this way, and one over 3 is going this way.

00:42:43.000 --> 00:42:49.000
And then here most people don't have any juvenile crimes, if you do, and so on, and so forth.

00:42:49.000 --> 00:43:01.000
And so if you decide that you don't like that branch you can. If you don't, you know want trees that look like that, you can go back up to the top and say you know what I think I want to look at trees

00:43:01.000 --> 00:43:13.000
where age is the first split, like I can split on age less than 23, and then I can look at those trees and pick out trees from there that I like and store them as well.

00:43:13.000 --> 00:43:24.000
Okay, so it really gives you kind of a bird's eye view of what the what the rasha months that really looks like And it's It's been a lot of fun to kind of like play with the rash mindset and figure out

00:43:24.000 --> 00:43:36.000

what the trees actually are that are in there, so and we can hopefully, you know, provide something that's more useful to practitioners because they can actually look through and pick out the trees that they want.

00:43:36.000 --> 00:43:45.000
And so. Yeah. So I you know, I think this is going to be more useful to practitioners, and this is where I think the future of Ai should be focusing.

00:43:45.000 --> 00:43:49.000
I think it should be on these more kind of human facing human facing questions.

00:43:49.000 --> 00:44:00.000
For Ai all right. So to summarize, I talked about the Rashomon set theory.

00:44:00.000 --> 00:44:17.000
I talked about fast, sparse and ghost, and the fact that this theory, the rational set theory and the existence of these maths, have huge implications for criminal justice, loan decisions and other high-stakes decisions.

00:44:17.000 --> 00:44:23.000
Because they make it very, very difficult to justify the use of black box models.

00:44:23.000 --> 00:44:31.000
I think it has huge implications, particularly if we can provide users the freedom to choose between models.

00:44:31.000 --> 00:44:41.000
Thank you very much

00:44:41.000 --> 00:44:51.000
Fantastic talk. Thank you Cynthia, so so one of the fun things about. I guess this mode of giving talks or or holding talks is that you know, usually as a host of a talk I'm.

00:44:51.000 --> 00:44:56.000
I'm spending a lot of the talk trying to think what would be a good question to ask at the end.

00:44:56.000 --> 00:45:00.000
But we've got 18 questions that have already been asked and so I was thinking to.

00:45:00.000 --> 00:45:03.000
Maybe we could just kind of go through them. I could.

00:45:03.000 --> 00:45:06.000
I could share with you what's what's going on in the Q. A.

00:45:06.000 --> 00:45:12.000
And you can respond. sure. did you want me to I know you wanted me to talk a little bit about my background.

00:45:12.000 --> 00:45:16.000
At some point that's true thanks for being so on the ball.

00:45:16.000 --> 00:45:25.000
So. so. One of the things that has been a structural part of this talk series in the past is that we have our speakers talk a little bit, you know.

00:45:25.000 --> 00:45:31.000
Talk about their research, which We're very excited about but also talk about their own personal journey, how they came to be where they are.

00:45:31.000 --> 00:45:38.000
And and you know, what do they like outside of of this this kind of research perspective?

00:45:38.000 --> 00:45:42.000
And Yeah, I would love it if you if you could tell us a little bit about yourself.

00:45:42.000 --> 00:45:49.000
Sure. So yeah. So I you know I didn't always want to do this right.

00:45:49.000 --> 00:45:56.000
This is not not where I envisioned I would ever be so it's kind of a journey getting here.

00:45:56.000 --> 00:46:00.000
So I majored in kind of mathematical physics and music theory.

00:46:00.000 --> 00:46:18.000
Those were my 2 majors. my French music at the turn of the last century. and unfortunately there's not too many people interested in modern composers of old music, and so that had to kind of go by the wayside if you want to listen

00:46:18.000 --> 00:46:23.000
to some of that music. you can go online. I had a few people.

00:46:23.000 --> 00:46:30.000
I had one of my friends played one of the pieces, then another random person on the Internet.

00:46:30.000 --> 00:46:37.000

I posted a score and I didn't have anyone to play it, and somebody random scientists had stumbled upon it and computer scientists and played it.

00:46:37.000 --> 00:46:41.000
And so I put the recording up so you can listen to some of that French music.

00:46:41.000 --> 00:46:48.000
If you want Well, It's not French me because i'm not French, but you know, whatever it's close enough just a i'm just a Francophile.

00:46:48.000 --> 00:46:52.000
I guess, and then So I I wanted to do applied math.

00:46:52.000 --> 00:47:03.000
So I went to applied math for grad school, and then I met a very energetic young scientist working on machine learning, and I realized these guys were. You know they were trying to predict the future from data.

00:47:03.000 --> 00:47:12.000
And I was like, Oh, this is really cool. I wanted to do this, and so I read, you know, statistical learning theory, all this kind of stuff, and learned about it.

00:47:12.000 --> 00:47:20.000
And you know, met a bunch of people as as a grad student that are now, you know, the like you know, i'd walk into Yann Lecoons office and ask him questions.

00:47:20.000 --> 00:47:29.000
About support vector machines, which is pretty funny in retrospect because obviously it's always working on neural networks.

00:47:29.000 --> 00:47:35.000
And you know, just was still answering my questions about non-neural network topics.

00:47:35.000 --> 00:47:38.000
So that was that's pretty cool and I believe I did meet you as a graduate student.

00:47:38.000 --> 00:47:44.000
Yeah, and then Yeah. So I started working in machine learning.

00:47:44.000 --> 00:47:54.000
And then, after I you know, I was working with Rob on convergence of adipus rupture, period, and convergence of addaboos, and then I switched to doing this, very applied work and then I got very depressed

00:47:54.000 --> 00:47:58.000
for a while, because none of it actually worked right. and all the stuff I learned about in grad school.

00:47:58.000 --> 00:48:02.000
None of it worked, and I was like, okay, you know what not doing any of this anymore.

00:48:02.000 --> 00:48:09.000
I'm gonna try to design methods. that are more easy to troubleshoot, because the power power company data was really messy.

00:48:09.000 --> 00:48:12.000
And so then I thought you know I just can't I can't troubleshoot. this stuff.

00:48:12.000 --> 00:48:17.000
It's just you know it's not getting me any better performance, and I can't troubleshoot it.

00:48:17.000 --> 00:48:25.000
And there was one kind of very embarrassing incident where we told the Power Company to go into a particular manhole, and they said, there's nothing wrong with that man hole.

00:48:25.000 --> 00:48:27.000
You guys are crazy. And we were like, Okay, what happened here?

00:48:27.000 --> 00:48:35.000
And it turned out that, like there were a lot of problems with the way we set the problem the way we had set it up.

00:48:35.000 --> 00:48:39.000
And you know it was just it wasn't we weren't getting targeted predictions.

00:48:39.000 --> 00:48:43.000
The way we could if we had understood every variable in everything.

00:48:43.000 --> 00:48:55.000
And then there was one time where we said, to the power company you know the number of neutral cable seems to be a really useful factor, and we don't understand why and they said there's something wrong with your model something wrong with your

00:48:55.000 --> 00:49:08.000
model. And then we were like, okay, and it turned. out that there was something wrong with the data like the data that they gave us was a snapchat of data from 1 point in time, and it didn't have have like if it If it was from

00:49:08.000 --> 00:49:15.000
multiple points. In time it would have been different. but we had some like leakage of information, because it was from only 1 point in time.

00:49:15.000 --> 00:49:23.000
And yeah, so it was. It was really a problem. So when we got rid of the number of neutral cables predict much better.

00:49:23.000 --> 00:49:26.000
It was actually really good at predicting what manhole events would happen in the future.

00:49:26.000 --> 00:49:30.000
I think they're like fires and explosions that happen in New York City.

00:49:30.000 --> 00:49:35.000
And so from there I just started working on interpretable machine learning, and then realized that a lot of our stuff was useful for health care.

00:49:35.000 --> 00:49:40.000
So I started working in healthcare in criminal justice.

00:49:40.000 --> 00:49:43.000
Yeah. fantastic to Spitz yeah, it's been a really interesting journey.

00:49:43.000 --> 00:49:52.000
And and you know the field is better for you having gone through what you went through here. Your shared your knowledge with everyone. So let me let me try to hit some of the questions.

00:49:52.000 --> 00:49:58.000
I'm not sure exactly how much time we have for that but there's so many interesting questions flowing in

00:49:58.000 --> 00:50:00.000
So this one relates to something I was wondering as well.

00:50:00.000 --> 00:50:12.000
So Alyssa? the Shenko asks: have you looked at, for example, the overlap in individuals for whom recidivism is predicted, using different models in the ration on set so like is it the case

00:50:12.000 --> 00:50:14.000
that the the same people are flagged by all these models?

00:50:14.000 --> 00:50:19.000
Or is it possible that for any given individual you've got a good model?

00:50:19.000 --> 00:50:22.000
But it actually disagrees with other good models okay so that's the magic question.

00:50:22.000 --> 00:50:29.000
So let me let me go to answer that here. So so I happen to have this slide.

00:50:29.000 --> 00:50:35.000
This is slide immediately after my Thank you, Slide. and this is showing you a different views of the Rashomon set.

00:50:35.000 --> 00:50:43.000
And so you can take that since that, rashomon set it's so many models, and you can think about these models as sort of representing, you know it.

00:50:43.000 --> 00:50:50.000
Just it. You can think about these, these, the whole set of these models, and then you can think about distances between models in different ways.

00:50:50.000 --> 00:50:55.000
So, for instance, to get from one tree to another tree, you could think about edit distance like you know you.

00:50:55.000 --> 00:51:08.000
You get rid of one node, and you put in a different node and then if you look at the distance between all trees in the rashomonset, and then you can project it down using a dimension reduction technique, so you

00:51:08.000 --> 00:51:18.000
take this graph of the trees and how they relate to each other. and you're just projecting it in a way that's trying to preserve the neighborhoods and the and the global structure. of this data.

00:51:18.000 --> 00:51:22.000
And you can. You can actually see interesting things in the in the Russia month set.

00:51:22.000 --> 00:51:35.000
So you can see that like the colors are for the top split in the tree so you can see that there's a lot of trees in the rash mindset that have different top splits, and then here instead of using edit distance we

00:51:35.000 --> 00:51:46.000
used prediction set different distance. So if 2 what you're looking at is that the distance between 2 trees is how many predictions are different between the 2 trees?

00:51:46.000 --> 00:51:56.000
And so when you project it down, you see that the trees can have very different predictions from each other, and that trees with the same top split, tend to have fairly similar predictions to each other.

00:51:56.000 --> 00:52:10.000
And then you can also look at the feature set distance, which is sort of like how many features are similar between 2 trees, even if the features are used in the different ways you just project down to like which features are used in the tree and

00:52:10.000 --> 00:52:22.000
then you can. You can view the view the Russia on set that way. And then these little blue circles are just showing you examples of the treat like this is the tree that lives here in this part of the space.

00:52:22.000 --> 00:52:28.000
And so you can see that these trees are similar to each other in that they have the same top split which is holding sword.

00:52:28.000 --> 00:52:40.000
So Yeah. So you can basically what i'm saying is that questions like that, like you know, you can answer that by looking at the Rashomon set just looking at it and projecting it down using this using these kind of dimension

00:52:40.000 --> 00:52:55.000
reduction tools. that give you kind of a really nice broad view of of everything about the Russia monette, which variables are used. How the predictions are different, all that stuff I feel like part of the question though, is is what are the implications of

00:52:55.000 --> 00:53:10.000
that like if it's the case that that we're using these models to make predictions or not not make predictions, but actually make decisions about real people and different simple models, make different recommendations as to what to do how how troubled should

00:53:10.000 --> 00:53:17.000
we be by that? and is there any way to mitigate it?

00:53:17.000 --> 00:53:21.000
That's that's what we're trying to that's what we're trying to get people to think about.

00:53:21.000 --> 00:53:25.000
Yeah, Yeah, right? and it's just when you've got these overly complicated models.

00:53:25.000 --> 00:53:32.000

You can't even really have that conversation so it's like the computer told me, and and that's there's not much you can do beyond that.

00:53:32.000 --> 00:53:37.000
This this next question. touches on me I never really thought about before, but it's really interesting.

00:53:37.000 --> 00:53:43.000
So Andrew Bell says: In many of the comparisons you made between black box and simpler models, accuracy is being used as a performance.

00:53:43.000 --> 00:53:48.000
Metric are the observations you've made about similar performance robust across other metrics.

00:53:48.000 --> 00:53:57.000
So Auc is one that you mentioned f one score, but most specifically or most importantly specific metrics like precision at K.

00:53:57.000 --> 00:54:01.000
So. how how much, how precise is it If you look at the top?

00:54:01.000 --> 00:54:10.000
K. and in Andrews work he observed that black box model black box models often tend to perform better than tabular on precision.

00:54:10.000 --> 00:54:18.000
Okay. Okay, So I I used my language very carefully to avoid other metrics here, but actually, ghost is generalized.

00:54:18.000 --> 00:54:24.000
Optimal stress, decision trees, and the way the reason it's called generalized, is because it can handle a wide variety of loss functions.

00:54:24.000 --> 00:54:33.000
So, whereas for something like cart, if you try to use cart that is like it's splitting criteria is kind of optimized for accuracy.

00:54:33.000 --> 00:54:43.000
So it just can't handle kind of like f one score or something like that. whereas ghost can optimize pretty much any of these like reasonable loss functions directly.

00:54:43.000 --> 00:54:49.000

So you can optimize, for instance, like you want to minimize the false, positive rate, subject to a constraint on the false negative rate.

00:54:49.000 --> 00:54:52.000
Guess what you just tell it to do it and it'll do it.

00:54:52.000 --> 00:54:53.000
If you want to optimize f one score, you can do that.

00:54:53.000 --> 00:54:57.000
If you want to optimize Ac. you can do that directly.

00:54:57.000 --> 00:55:03.000
If you want to optimize weighted accuracy or balanced accuracy, you just tell goes to do it, and it'll optimize that.

00:55:03.000 --> 00:55:07.000
Then we also in the Tree Farms paper, where we produce the whole ration on set.

00:55:07.000 --> 00:55:17.000
We showed that if you take the Rashomon set for accuracy, and you increase it a little bit, if you make it a little bit bigger like you increase the parameter that governs how many trees you recover like the

00:55:17.000 --> 00:55:26.000
accuracy of trees. You recover. Then the bigger set there includes the Russia on set for F one score, and the Rushman set for balanced accuracy.

00:55:26.000 --> 00:55:35.000
So you're gonna actually get all of the good f one models and all of the good A Uc. or balanced accuracy models.

00:55:35.000 --> 00:55:38.000
You can get them all within the rashoman set for accuracy.

00:55:38.000 --> 00:55:48.000
As long as you collect a few more models, than you could before, and we related how these we actually showed how the parameters relate to each other for the different metrics.

00:55:48.000 --> 00:55:55.000
So that's a unique thing about these new methods and that there, you know, they directly optimize for these metrics. right?

00:55:55.000 --> 00:55:59.000

But so that's that's very cool the the question that was asking about precision at K.

00:55:59.000 --> 00:56:08.000
Which feels maybe different than some of these other metrics, and whether or not the same kinds of the Rashman argument kind of seems to hold there as well.

00:56:08.000 --> 00:56:14.000
Precision. a k need to go back and look at precision at case.

00:56:14.000 --> 00:56:23.000
Specifically i'm not sure I can answer that for precision at Cape, and i'm guessing based on all all these other performance metrics like f one and balanced accuracy and weighted accuracy that you're

00:56:23.000 --> 00:56:27.000
gonna get very similar performance gonna be hard to be different on that one specific metric given.

00:56:27.000 --> 00:56:30.000
How similar all the other ones are yeah no it's it's interesting.

00:56:30.000 --> 00:56:43.000
But it yeah that's that's really cool we have noticed that cart and c 4.5 just don't perform well in those metrics. So that could be what what andrew is seeing as well okay alright that's yeah it's

00:56:43.000 --> 00:56:52.000
worth kind of revisiting chinadu Ella asked a question which I think is on a lot of people's minds, which is okay?

00:56:52.000 --> 00:57:02.000
Well. so you leaned pretty hard into trees, and definitely some of the trees you showed us are just, you know, any of anybody can look at those and get a feel for it very quickly.

00:57:02.000 --> 00:57:09.000
But really how explainable and in terms of explainability, how simple is a decision tree model.

00:57:09.000 --> 00:57:14.000
So Are they really simple enough for people it's It depends on the tree.

00:57:14.000 --> 00:57:17.000
It depends on the problem. It depends on the tree and it depends on who's looking at it?

00:57:17.000 --> 00:57:20.000
So interpretability is just inherently subjective.

00:57:20.000 --> 00:57:27.000
And So it could be that I show you a tree and it's very tiny, and you might say that doesn't make any sense.

00:57:27.000 --> 00:57:34.000
Your tree is not interpretable, because that variable doesn't make any sense in this model, and so even a very scarce tree might not be interpretable to you.

00:57:34.000 --> 00:57:44.000
So I think you know sparsity is kind of a good it's. It's not just a proxy for interpretability, but it you know it's a requirement in some it's it's a requirement in some cases

00:57:44.000 --> 00:57:48.000
not all cases, some cases. okay. So so big trees big anything is never going to be helpful.

00:57:48.000 --> 00:57:53.000
But small things. Some of them could be helpful yeah so it's you know.

00:57:53.000 --> 00:57:55.000
I mean, humans can handle 7 plus or minus 2 cognitive entities at once.

00:57:55.000 --> 00:58:02.000
So when you start building these really big models that people can't even keep in their their head, then you know you're just losing.

00:58:02.000 --> 00:58:05.000
You're just sacrificing something that you don't need to sacrifice.

00:58:05.000 --> 00:58:14.000
So it's it's really you know we're building something that the whole idea of the rashomans said, Is that just the one tree I give you is not going to be interpretable for you So you need to choose between

00:58:14.000 --> 00:58:26.000
them. Gotcha alright. So this is a question maybe not for you. but I think it's worth answering it publicly, which is chile Song asked Great talk.

00:58:26.000 --> 00:58:30.000
Would it be possible to share the recording i'm currently teaching Ml.

00:58:30.000 --> 00:58:35.000
Would love to share this with my students. So I think that maybe Blaine or somebody else can answer this question.

00:58:35.000 --> 00:58:41.000
How how can people get access to the recording after after we're done?

00:58:41.000 --> 00:58:52.000
I put the link in the chat. you can just visit our website and a setup of slash events, and you will see a recording of our lecture today.

00:58:52.000 --> 00:59:00.000
Pop up within a few days of the lecture. Perfect thanks very much alright.

00:59:00.000 --> 00:59:09.000
And Abida Zacor. says, does the Rashman set theory also apply to more complex data like images and computer vision?

00:59:09.000 --> 00:59:13.000
So you talked in your talk that there really are these maybe maybe they should even be the same field.

00:59:13.000 --> 00:59:18.000
There's these 2 different kinds of machine learning what can you tell us about that other branch of machine learning.

00:59:18.000 --> 00:59:24.000
Yeah. So a few years ago we published a paper on interpretable neural networks, and we've been using it for the computerated Mammography project.

00:59:24.000 --> 00:59:29.000
And you know I was wondering about that. I was wondering like.

00:59:29.000 --> 00:59:32.000
Is it really true, like, Do you need to have a black box for computer vision?

00:59:32.000 --> 00:59:42.000
But what does interpretability even mean for computer vision? and So the way we set up our interpretable neural network is that we forced it to do case-based reasoning.

00:59:42.000 --> 00:59:55.000
So it just it tells you like I think this part of the image looks like this part of this other image. And so it's doing these comparisons of parts of this test image to like parts of training images where we know what's going on and so for computed

00:59:55.000 --> 01:00:04.000

computerated mammography it's taking the breast lesion, and it's breaking it up into different pieces and comparing each piece with known cases and saying, Well, I think this looks like that and this looks like that So

01:00:04.000 --> 01:00:07.000
it's not keener's neighbors but it's like K.

01:00:07.000 --> 01:00:11.000
Nearest parts of prototypical cases so it's kind of get that K.

01:00:11.000 --> 01:00:16.000
Nearest neighbor's e case-based reasoning feel to it.

01:00:16.000 --> 01:00:26.000
So it provides a lot of you know the kind of explanation that a human would explain to another human, and it at least gives you a sense of whether the neural network is reasoning properly about the image.

01:00:26.000 --> 01:00:38.000
And I think the reason we were able to incorporate constraints into the neural network and still maintain black box performance is because the thrasherman set theory I still think there's a Russia onset for computer

01:00:38.000 --> 01:00:48.000
vision. I was recently reading a paper by Michael. Bazani, who has well and other people, but he I've mentioned him because he had my job number of iterations to go.

01:00:48.000 --> 01:00:51.000
So he was kind of on my mind, but he was saying that for explainability of images.

01:00:51.000 --> 01:01:03.000
It's really useful to not just point at things but to label The things you're pointing at, and it sounds like what you're describing is a somewhat automated way of doing that to say this piece of the image actually relates to

01:01:03.000 --> 01:01:09.000
this other thing that we've seen before now giving them actual like human, interpretable labels would be even better.

01:01:09.000 --> 01:01:13.000
But that seems like a fantastic way of getting off the ground.

01:01:13.000 --> 01:01:19.000

Yeah, I think getting the human interpretable labels is quite difficult for something like mammography, because the words don't exist.

01:01:19.000 --> 01:01:25.000
Yeah, Yeah, it's like so the human has to kind of figure out how this is similar to that.

01:01:25.000 --> 01:01:32.000
But that's fair, though. yeah, domain experts I feel like they spend a lot of their training putting names to things right really getting getting familiar with yeah, something.

01:01:32.000 --> 01:01:41.000
And then one of the ways to own it is to give it a name right, and that that way you become, you become the powerful thing that the the name giver?

01:01:41.000 --> 01:01:48.000
Yeah, Nathan, Crosby said, asks, Are there certain conferences that are more receptive for interpretable models than others?

01:01:48.000 --> 01:01:54.000
Do I need to leave out references to explainable Ai or other trigger words that the X Ai community cannot get passed right now.

01:01:54.000 --> 01:02:00.000
So what's been your experience with the the sort of sociology of research in this area?

01:02:00.000 --> 01:02:03.000
And do you have any advice for for other people who are contributing?

01:02:03.000 --> 01:02:10.000
You know i'm really glad you asked me that because I forgot to put that into my bio, and you asked me to tell about my background, which is that almost my whole career.

01:02:10.000 --> 01:02:13.000
People have been telling me how stupid this whole research area is.

01:02:13.000 --> 01:02:23.000
I mean you were not one of those people, Michael. You were one of the few who was like, Okay, but like although I mean, I would give talks, and people would walk up to me and yell at me like, why do we get this?

01:02:23.000 --> 01:02:29.000
Interpretability stuff, because the whole idea of machine learning was that the the algorithm takes care of everything that you just give it the data.

01:02:29.000 --> 01:02:33.000
And it tells you what tells you the predictions. and you know I was.

01:02:33.000 --> 01:02:39.000
I was trying to tell them like you know I don't trust the data set like That's therefore I can't trust the model that you build from the data set.

01:02:39.000 --> 01:02:44.000
But like I didn't have the words back then and also like just people were not into it.

01:02:44.000 --> 01:02:49.000
It was just such a negative, horrible thing to work on interpretable machine learning.

01:02:49.000 --> 01:03:01.000
So it was just. It was just really difficult, to get anything accepted. There's a squishiness to it that I think the field at that time was not ready to embrace right, because like Well, people people are Oh, people.

01:03:01.000 --> 01:03:10.000
Are busy we don't have equations for them we we can't think about this, but it's So is it different now? And if so, where like where should people be trying to publish this kind of work?

01:03:10.000 --> 01:03:14.000
Where where's what's more where are people more receptive I don't know.

01:03:14.000 --> 01:03:18.000
So I also I mean, I also have trouble getting applied work published.

01:03:18.000 --> 01:03:28.000
That's another one so you're actually doing something that's going to benefit the world. if it's not scientifically notvel, then you know some ai or viewer doesn't deem it scientifically novel

01:03:28.000 --> 01:03:32.000
it's very, very difficult to publish it and that I think, is just a blemish on the field.

01:03:32.000 --> 01:03:39.000
I think we're. an embarrassment to the world by not allowing, you know, high quality applied papers to be published in the field.

01:03:39.000 --> 01:03:44.000
So you know I'm. obviously doing my best to work on that through being on the Acm.

01:03:44.000 --> 01:03:50.000
Sig Kd executive board and trying to renovate Kdd and doing a bunch of other stuff.

01:03:50.000 --> 01:03:59.000
But it's really a problem. it's really a problem, because you know, nurps, icml ai stats ui!

01:03:59.000 --> 01:04:07.000
They don't accept apply papers. they just don't and even the places that apply, except applied papers like you know Kdd or or Iii.

01:04:07.000 --> 01:04:20.000
They want it deployable which basically rules out all of science. So if there's really no there's really no good solution to that, and I kind of search around for where I can go to send things.

01:04:20.000 --> 01:04:22.000
I had a pretty decent luck with Dmkd. for what? For a while.

01:04:22.000 --> 01:04:24.000
So. Dmk: Do you spell that more slowly? Yeah.

01:04:24.000 --> 01:04:32.000
Dm. Kit is the data mining and knowledge discovery It's the top journal and data mining, and they seem to have been able to handle some of my applied papers in the past.

01:04:32.000 --> 01:04:41.000
But there have been cases where you know i've had to write to the editor and say you know your scope in this journal says that you take a apply papers, but you don't you know because here you you just

01:04:41.000 --> 01:04:44.000
rejected my applied paper because it was applied.

01:04:44.000 --> 01:04:59.000
So I really do think that there needs to be a major So major surgery done on on our, on our, on our machine learning in Ai worlds to accept a flood papers and papers on interpretability, thanks thanks for

01:04:59.000 --> 01:05:03.000
bringing that up hopefully. Some of the people listening we'll take that to heart.

01:05:03.000 --> 01:05:13.000
Jim, please? asked I think a classic question about machine learning, which is, How do you decide when marginal improvement is worth the complexity?

01:05:13.000 --> 01:05:23.000

Right to trading off those 2 things. A small improvement in high-stakes decisions could also be incredibly practically significant. and it seemed like that's going to vary from problem to problem.

01:05:23.000 --> 01:05:30.000
I think it does vary from problem to problem I mean there's there's some things in say criminal justice that are, you know.

01:05:30.000 --> 01:05:40.000
It makes me think that it doesn't it's not really worth it, because, for instance, even to check your own data like we've had all these cases, there were articles published in the New York Times, where people couldn't calculate their

01:05:40.000 --> 01:05:51.000
own. they they their risk score was miscalculated, because their criminal history data was entered wrongly into the model, and they couldn't check it in the pearl board just denied, their pearl like every once in a

01:05:51.000 --> 01:05:58.000
while I get letters like there's a letter over there in my office, from a prisoner in some prison, and you start reading it.

01:05:58.000 --> 01:06:06.000
And you think, does this person? Why are they writing me and Then you realize that they know all this very detailed mathematics about norm groups?

01:06:06.000 --> 01:06:13.000
And you're like How do you know this stuff and the answer is because that determines their freedom, and they shouldn't have to know that you know it.

01:06:13.000 --> 01:06:20.000
Just it. They should just know that it should be a very simple formula, and they shouldn't have to know what their norm group.

01:06:20.000 --> 01:06:27.000
You know what all this you know what I mean it's just embarrassing that our that our system works like that.

01:06:27.000 --> 01:06:30.000
So I think just being able to troubleshoot in a lot of like medical and criminal justice cases.

01:06:30.000 --> 01:06:41.000
Just that alone. You really need the interpretability for Well, said I think it's about the the stakes matter. the stakes matter.

01:06:41.000 --> 01:06:46.000
Yeah, yeah, right right, right, it's like Oh, there's there's, you know, 99% accurate.

01:06:46.000 --> 01:06:51.000
But the 1% that makes a mistake on is you like it matters alright.

01:06:51.000 --> 01:06:55.000
Let's see what this is scrolling on me Harry.

01:06:55.000 --> 01:07:02.000
Sorry, Alice Schwartz asked. is it straightforward to extend ghost to decision forests, or does that defeat the purpose of searching for simple models?

01:07:02.000 --> 01:07:06.000
So how do? How do decision for us fit into your worldview?

01:07:06.000 --> 01:07:10.000
I think it possible that you could create interpretable decision.

01:07:10.000 --> 01:07:14.000
Forests like you know, you have a vote of a few decision trees that's just not something that we've done.

01:07:14.000 --> 01:07:25.000
So I think it's. I think it's possible to extend it. it's just not something that we've that's on our agenda it doesn't feel completely foreign to because you had at least that one model

01:07:25.000 --> 01:07:30.000
with the little step scoring thing that went down where each variable had its own little score profile.

01:07:30.000 --> 01:07:34.000
Yeah, that seems, you know, decision or whatever for us are a little bit like that.

01:07:34.000 --> 01:07:36.000
There's just a whole bunch of things that are contributing to the answer.

01:07:36.000 --> 01:07:41.000
As long as there's not too many of them in each one individually is gracable.

01:07:41.000 --> 01:07:50.000
It might not be so terrible just to just to Remind everybody This is the slot that we've got for this is an hour and a half, so we still have 22 min left.

01:07:50.000 --> 01:07:53.000
If you need to go because you'll allocate it an hour.

01:07:53.000 --> 01:08:02.000
Thanks for joining us but i'm gonna i'm gonna keep. we've got a ton more questions, and I'm just going to keep going through and hopefully find interesting ones for Cynthia to respond to Well, this

01:08:02.000 --> 01:08:06.000
next one is definitely near and dear to my heart. Harry Dan Quitz asked.

01:08:06.000 --> 01:08:10.000
What can be said about modeling repeated actions rather than single shot?

01:08:10.000 --> 01:08:23.000
Classification tasks. So you know something closer to reinforcement learning, I mean, I don't I only have one paper on a interpretable reinforcement learning and i'm just and that was that was you know my

01:08:23.000 --> 01:08:31.000
colleagues who are very very smart you know they they came up with an interpretable policy for mazes, so I don't.

01:08:31.000 --> 01:08:43.000
I don't really know the answer. to that one you know like it could be that it could be that these decision trees are really useful for writing down interpretable policies.

01:08:43.000 --> 01:08:48.000
But you know i'm not maybe a little bit beyond me to kind of answer.

01:08:48.000 --> 01:08:53.000
Maybe you could answer that one, Michael I mean I think I think I think there's a place for it.

01:08:53.000 --> 01:08:57.000
I think that having I mean So what? what? reinforcement learners try to learn our policy?

01:08:57.000 --> 01:09:01.000
So decisions, action decisions that you make as you're interacting with the environment.

01:09:01.000 --> 01:09:09.000
It's a little tricky you can't just turn it into a classification problem, because potentially one small little mistake can cascade.

01:09:09.000 --> 01:09:13.000
So it's not enough to just say Okay, Well, it's 99% accurate on individual decisions.

01:09:13.000 --> 01:09:20.000
And then you run it. and it's like 2% accurate, because one of those decisions that it made a mistake on is actually critical all the time.

01:09:20.000 --> 01:09:22.000
So I think you need to. maybe re-weight things a bit.

01:09:22.000 --> 01:09:33.000
But but having having interpretable policies, for lots of problems would be really valuable. We, we don't understand what a lot of these these, these reinforcement learning systems are actually deciding.

01:09:33.000 --> 01:09:41.000
So you have a a program that's playing. it's either being a self-driving car, or maybe it's playing a board game or something like that, and it seems great.

01:09:41.000 --> 01:09:52.000
But sometimes you can actually poke at them. and discover that they've got these weird holes in them, and that would be easier to tell that if we had some description of what it was that they were doing Yeah, you could probably

01:09:52.000 --> 01:10:02.000
also. So i'm not a big fan of explaining black boxes, because i'd like people to try to create interpretable models rather than just being satisfied with an explanation.

01:10:02.000 --> 01:10:12.000
But you could try to interrogate like if you had a black box, you know method, you could try to interrogate it and try to figure out how to design an interpretable model from that.

01:10:12.000 --> 01:10:21.000
You know, like, if you take the policy that the black box is using, and you can try to created a decision trade that would mimic that policy and see where they disagree.

01:10:21.000 --> 01:10:28.000
For instance, and see whether It's could give you any information yeah that makes sense.

01:10:28.000 --> 01:10:35.000
Haranath Garro doardi dot dot orgy asks for the healthcare applications.

01:10:35.000 --> 01:10:42.000

Have experts considered anomalies in addition to averages, or maybe another way to to ask it is,

01:10:42.000 --> 01:10:54.000
Can we learn anything from these models that that that influences how we think about the the problem as opposed to not just, you know, providing a better classifier?

01:10:54.000 --> 01:11:01.000
So in other words, so the way i've been building mine is like, for you know, computer aided decisions.

01:11:01.000 --> 01:11:05.000
So these are humans in charge, and then you know you're getting a little bit of information.

01:11:05.000 --> 01:11:19.000
But the humans actually making the decision. so we you know it's possible that you could use it for scientific or medical discovery right? You could It's just not what i've been working on Gotcha gotcha you don't have any

01:11:19.000 --> 01:11:21.000
anecdotes of of the experts looking at the trees and going.

01:11:21.000 --> 01:11:31.000
Oh, wait a second. Well, I mean, we are using these kind of decision Trish type models for materials discovery.

01:11:31.000 --> 01:11:40.000
And we have noticed that certain patterns, that appear in the material like these are amount of materials, so they're kind of like a mixture of 2 different materials, a soft material and a hard material.

01:11:40.000 --> 01:11:46.000
And we've been finding certain types of geometric patterns that lead to certain kinds of band gaps.

01:11:46.000 --> 01:11:51.000
And so the experts look at it and go. Oh, yeah, if you have this kind of star-shaped pattern that it has this kind of band gap.

01:11:51.000 --> 01:11:55.000
So I guess you could you know you could think about that as being an example.

01:11:55.000 --> 01:12:03.000

Yeah, that's that's really neat I just feel like in general, just the more quality time you spend with your data, the more that that you do deeply understand it.

01:12:03.000 --> 01:12:13.000
And so some of these tools, especially that tree browser man that's that is wicked cool just the the ability to go, and and just like explore this set this really how complicated set.

01:12:13.000 --> 01:12:22.000
And look for interesting stuff in there? it's it's super informative another question for blame, which is har enough oops?

01:12:22.000 --> 01:12:27.000
Garrardi asked, Where do I see the questions from?

01:12:27.000 --> 01:12:30.000
The other participants. So i'm reading these questions out can other people see them.

01:12:30.000 --> 01:12:33.000
I think I think it comes in the Q. and a button which should be at the bottom of everybody's screen.

01:12:33.000 --> 01:12:39.000
Certainly the bottom of my screen Is it? Is that right? It should be in the Q. A. button.

01:12:39.000 --> 01:12:49.000
Fred, are you able to confirm that so all the people answering this question? Are the people who are have a different interface from everyone who wants to know the answer?

01:12:49.000 --> 01:12:52.000
So we're having a little trouble verifying it but if there is a Q. A.

01:12:52.000 --> 01:12:57.000
Button at the bottom. Push that, and you can scroll through and see all these amazing questions that people are asking.

01:12:57.000 --> 01:13:06.000
I am I am skipping a bunch of them and not because they're not great questions all of them have been fantastic, but i'm trying to probe different i'm like the tree browser i'm trying to probe different parts

01:13:06.000 --> 01:13:12.000
of you know. sort of cynthia's knowledge space t to give people the the most complete picture that we can. next question.

01:13:12.000 --> 01:13:17.000
I was gonna ask is from ken wang who's a division director.

01:13:17.000 --> 01:13:22.000
Sorry the program director in my division. what are the subtleties in estimating how large the Rashomon set is?

01:13:22.000 --> 01:13:30.000
So when you survey the space of good, simple models, do you come across good, simple models with contradictory interpretations?

01:13:30.000 --> 01:13:36.000
I guess we kind of touched on that a little bit like labeling the same point differently.

01:13:36.000 --> 01:13:41.000
Different models will label the same point differently. but what about how so?

01:13:41.000 --> 01:13:46.000
So so you you frame the rash, mindset thing as kind of a theory.

01:13:46.000 --> 01:13:51.000
Is there any way, is it possible to actually verify or or confirm the theory?

01:13:51.000 --> 01:13:58.000
Is, or or the all the the magnitude of these things that you're asking about are just too large to to answer these questions.

01:13:58.000 --> 01:14:03.000
Well, we've we've we've been poking at it from different directions.

01:14:03.000 --> 01:14:15.000
So what we what we have we have some kind of very small something that's very close to a proof that Ron and Lessey and I were working on with like 2 Gaussians in one dimension so if you 2

01:14:15.000 --> 01:14:22.000
Gaussians in one dimension if those gaussians kind of overlap, so that you have some noise in the wat in the labels right?

01:14:22.000 --> 01:14:32.000
If the label, if there's enough label noise that these Gaussians kind of overlap with each other like that, y equals one Gaussian and the Wyles 0 Gaussian, or whatever they Overlap with each other, then you

01:14:32.000 --> 01:14:35.000
can prove that the Rashomon said is larger, or you could almost prove it.

01:14:35.000 --> 01:14:46.000
There was a little tiny bit of math that we couldn't quite solve. but we know the answer numerically, so that's as close as we got to a proof on on showing that logically proven theorem Well, yeah, but

01:14:46.000 --> 01:14:57.000
it's I mean it's like. you know you could see it on the plot, like as long as these 2 numbers, you know, meet each other, which they do, they do. So we can at least prove that you know if if we we can at least

01:14:57.000 --> 01:15:09.000
prove that if the data are kind of gaussian ish that that there's which a lot of data are Gaussian ish multi, you know, multi-class problems, or the classes are gaussian

01:15:09.000 --> 01:15:17.000
ish that that you have a lot of noise and then in those cases you actually can probably have a large Russian onset.

01:15:17.000 --> 01:15:25.000
So We're getting closer to kind of proving it but I think you know the part of the issue.

01:15:25.000 --> 01:15:29.000
Is that like? What does it mean? You know what? What are the what distance metric?

01:15:29.000 --> 01:15:37.000
Do you want to use to create the rational mindset and that's? that's kind of a complicated question that we I mean, I've been trying to just be very conservative in what I call a hypothesis

01:15:37.000 --> 01:15:44.000
right some of it, you know we've actually computed the number of model like we actually have these numbers.

01:15:44.000 --> 01:15:54.000
So if you compute it this way, then, we already have that empirical proof for these data sets, but it's it's hard to say anything in general, because just because you show it on 70 different data sets doesn't mean

01:15:54.000 --> 01:16:00.000
it's true, for all the data sets so you know yeah, Right?

01:16:00.000 --> 01:16:05.000

That would be impossible to for all possible data sets. alright.

01:16:05.000 --> 01:16:11.000
So So A couple of people asked versions of this question, but I like Patricia Francis Lyons version, so she says, fascinating.

01:16:11.000 --> 01:16:17.000
Thank you Do you know I think what's happening is when I read the questions.

01:16:17.000 --> 01:16:24.000
Probably Blaine is is marking them as being answered, and that causes the displayed to jump.

01:16:24.000 --> 01:16:27.000
And so in the middle of it I lose where we are.

01:16:27.000 --> 01:16:32.000
All right it It's still there. yeah but I just Can't I have to find the one I was reading fascinating.

01:16:32.000 --> 01:16:37.000
Thank you. Do you know how reproducible the results of this process are from domain expert to to domain?

01:16:37.000 --> 01:16:47.000
Expert in the context of how do we correct for bias of the field of expertise which is historically rampant in human health and behavior?

01:16:47.000 --> 01:16:50.000
So we're still are we do we do we have an answer to the question.

01:16:50.000 --> 01:17:00.000
Why can we get around the fact? that people are biased or are we still stuck, having to deal with that some other way? Now you gotta deal with it some other way.

01:17:00.000 --> 01:17:06.000
I mean we're we're even injecting more uncertainty in this process, right?

01:17:06.000 --> 01:17:11.000
So now you have to deal with it some of the again. the hope is that it's more.

01:17:11.000 --> 01:17:22.000
It's more discussable right because it's a little bit more out in the open, though I guess if it's if it's if if the expert is the one using the tree browser to pick one particular, tree, from this huge

01:17:22.000 --> 01:17:26.000

space. Well, yeah, that's worrying right because we have no idea what went into that decision.

01:17:26.000 --> 01:17:31.000
But in general it's at least you can at least you can see it, and and interrogate it right.

01:17:31.000 --> 01:17:40.000
That's a good word. Well, it'd be nice if you know, instead of the doctors having to drop write these models by hand, which is what they had to do in the past, they could like, you know.

01:17:40.000 --> 01:17:45.000
Look at all the models and then work together to pick a model that agreed with the data.

01:17:45.000 --> 01:17:49.000
But also their expertise kind of together, right that's how all of our metal.

01:17:49.000 --> 01:17:54.000
So all of our medical scoring systems essentially were developed by teams of doctors right.

01:17:54.000 --> 01:18:02.000
And so the idea is now you're you're just figuring out a way to insert data into that process in a way that works with it.

01:18:02.000 --> 01:18:12.000
Fantastic. Chris Bill has asked in my experience, Clustering and other unsupervised models can give widely different results.

01:18:12.000 --> 01:18:16.000
Does the theory of Rashmon sets apply in the unsupervised setting?

01:18:16.000 --> 01:18:26.000
I'm sure it does we haven't been working on a lot of rash mindset problems for unsupervised data, but i'm sure it does i'm sure it's even worse for

01:18:26.000 --> 01:18:31.000
unsupervised data, yeah yeah that's that's my sensible.

01:18:31.000 --> 01:18:36.000
But, as you said, a lot of those unsupervised models are often solved by reducing them to supervise learning.

01:18:36.000 --> 01:18:39.000
Right. So you you you tear and you take the data, the unlabeled data.

01:18:39.000 --> 01:18:45.000

And you say well use this part of it as a label, and then it sort of becomes just like all the other stuff.

01:18:45.000 --> 01:18:49.000
Yeah, you can think about a semi-supervised approach.

01:18:49.000 --> 01:18:56.000
Or sorry self-supervised yeah I I don't know the answer to that one because we haven't been working on.

01:18:56.000 --> 01:19:07.000
I think you know, or unsupervised you there's so much debate about like what even the performance metric is so for something like clustering right there's like 20,000 different performance metrics I mean that's true

01:19:07.000 --> 01:19:11.000
for every you know, for every machine learning problem there's like a ton of different performance.

01:19:11.000 --> 01:19:15.000
Metrics, but I think for unsupervised problems the performance metrics are even more complicated.

01:19:15.000 --> 01:19:21.000
And so this even more debate. Yeah, and sort of less settled on what's the reasonable way of looking at a given problem?

01:19:21.000 --> 01:19:27.000
One of the attendees asked: well, the question is, have you Have you have your package in Matlab?

01:19:27.000 --> 01:19:30.000
But I guess the question more generally is, Is your code available?

01:19:30.000 --> 01:19:34.000
Is any or any of the projects available for other people to use.

01:19:34.000 --> 01:19:39.000
Yes, they are. Let me just go. Okay, so let me just share my screen.

01:19:39.000 --> 01:19:42.000
First see the see if I can do this in real time.

01:19:42.000 --> 01:19:47.000

So timber truck. I actually just went to this website So this is just a website, and I just went to it.

01:19:47.000 --> 01:19:57.000
So there's not there was nothing you know nothing hit in there, and then there's different data sets here that you can, and you can use it if you you know there's Yeah, you can use it if you want i'm

01:19:57.000 --> 01:20:04.000
just going to see if I can get to my own website So I'm going to go to my code website.

01:20:04.000 --> 01:20:11.000
And here is the this is ghost. This is the ghost code which is over here.

01:20:11.000 --> 01:20:16.000
We are very proud that it now works on windows in addition to working on other things.

01:20:16.000 --> 01:20:19.000
There's so there's that oops I think I brought that up.

01:20:19.000 --> 01:20:27.000
But you couldn't see it, but anyway, so that's the ghost code, and then this is the fast sparse code, which is right here.

01:20:27.000 --> 01:20:30.000
Here are some of its predecessors, but this fast sparse here.

01:20:30.000 --> 01:20:36.000
We actually have publicly available code for dimension reduction, and then for some other problems.

01:20:36.000 --> 01:20:40.000
And then, yeah, all of I mean the tree farms code is in the Tree Farms paper.

01:20:40.000 --> 01:20:44.000
Since that's a very recent paper you actually have to go to the you have to.

01:20:44.000 --> 01:20:47.000
You have to go to the paper and then there's a link within the paper to get it.

01:20:47.000 --> 01:20:54.000
But yeah, it's all it's all everything that we do is public. all the code, except except if we're working with like medical data, We can't make that public.

01:20:54.000 --> 01:21:01.000

But everything else is public Devin, Martin said. Thanks for the talk.

01:21:01.000 --> 01:21:11.000
If the rashoma offset is so big that you can let people pick the model that makes the most sense to the problem, then doesn't that indicate that human understanding of the problem is irrelevant it seems like this is tacking on

01:21:11.000 --> 01:21:17.000
expertise after the fact. So I think that's How would you how would you reframe that question?

01:21:17.000 --> 01:21:20.000
No, I I disagree because the thing is that data is finite.

01:21:20.000 --> 01:21:26.000
The data spinning implemented. Yeah, and so because that data doesn't have everything you need to know.

01:21:26.000 --> 01:21:35.000
You need to main expertise to figure out which of the models in the rashoman set is the rate, you know, this is one that you would feel comfortable, using right?

01:21:35.000 --> 01:21:41.000
So it also when you were talking about this stuff, it also kind of made me think about, you know.

01:21:41.000 --> 01:21:44.000
Are there implications? The fact that there seems to be okay alright.

01:21:44.000 --> 01:21:50.000
So there's simple problems that have simple solutions and then there's really complex, noisy problems that have simple solutions because they're noisy, right?

01:21:50.000 --> 01:21:55.000
Because it's it's not really possible to do super well,

01:21:55.000 --> 01:21:58.000
What is does that have implications to artificial intelligence?

01:21:58.000 --> 01:22:02.000
In the in the in the original, broad sense of the term.

01:22:02.000 --> 01:22:08.000
You know, how does, how does, how, how would artificial intelligence be created?

01:22:08.000 --> 01:22:11.000
What is intelligence. Maybe that's what i'm asking does this say like human intelligence?

01:22:11.000 --> 01:22:22.000
Are we taking advantage of this as people like we're we're living in a really noisy world. Maybe a lot of what we're doing is simpler than it appears or that we're somehow leveraging the fact that things are simple to be

01:22:22.000 --> 01:22:29.000
able to control our environments. Better like have you thought about the the Ai implications, I mean.

01:22:29.000 --> 01:22:41.000
I think what you're saying is true that but it could be because the cognitive capacity of humans is limited that you know, like you go into the store and you're like How do I pick a How do I pick a television

01:22:41.000 --> 01:22:51.000
and then you just limit it based on all like a decision tree you're like I I want to. I want a bigger screen. So I'm going to go this way, and then I want a you know this remote So I go this way.

01:22:51.000 --> 01:22:57.000
Or I want. You know this pixel december I don't know, so I do.

01:22:57.000 --> 01:23:05.000
I do think humans kind of, you know, kind of develop their own sort of simple rules in the in the way that they navigate the world, because otherwise it could be.

01:23:05.000 --> 01:23:12.000
You know it could be quite overwhelming there's there's a whole research program in psychology on I think they call it fast simple.

01:23:12.000 --> 01:23:20.000
Hereuristics, which again feels kind of kind of related like that people do seem to find these simple rules that they can apply broadly.

01:23:20.000 --> 01:23:26.000
They somehow learn them. they somehow find reasonable rules things like if you don't know if you're trying to name the capital of the Us.

01:23:26.000 --> 01:23:32.000
State, and you don't know it. name of a City that has a high population in that State that often works, and and you find that people do that right.

01:23:32.000 --> 01:23:38.000
They A lot of people think Philadelphia is the capital of Pennsylvania, because they they know Philadelphia.

01:23:38.000 --> 01:23:49.000
So. So yeah, there does. There seems to be something going on in the way that people navigate the complicated world that that is echoed by the the kind of stuff that you're doing.

01:23:49.000 --> 01:23:54.000
Yeah, I mean there's also a bunch of causal questions as well that you could ask

01:23:54.000 --> 01:23:59.000
So I almost always get some kind of question about causality so everything that I'm doing is predictive here.

01:23:59.000 --> 01:24:04.000
It's not there's nothing causal but I mean there.

01:24:04.000 --> 01:24:16.000
There is a lot of work on observational causal inference, where, if you wanted to use, say if you wanted to develop a decision tree, for instance, that estimates conditional average treatment effects, That's something that you can do So what you

01:24:16.000 --> 01:24:20.000
can do is like match someone who didn't have the treatment with someone who had the treatment.

01:24:20.000 --> 01:24:30.000
And then you make a whole bunch of causal assumptions, and you have at least match groups, and then you can run a You can run ghost on that set of match groups, and you can actually get conditional average treatment like a tree

01:24:30.000 --> 01:24:36.000
of conditional average treatment. effects and that's we found that to be quite helpful for things like trying to dose medication.

01:24:36.000 --> 01:24:51.000
So that's that's something we're working on Yeah. cool, Gabriel Burnier Coleborne asked. Given the noisiness of both data and modeling should ml practitioners refuse to do certain

01:24:51.000 --> 01:24:56.000
tasks that have serious, real-world impacts like pre predicting recidivism.

01:24:56.000 --> 01:25:00.000

I don't think so because if you don't do it.

01:25:00.000 --> 01:25:06.000
Then these sneak oil companies will do it, and the justice system will pay for that.

01:25:06.000 --> 01:25:12.000
And then our justice system will be like it is Now continue to be like his.

01:25:12.000 --> 01:25:16.000
Now. yeah, I mean to be honest, the compass is very unusual.

01:25:16.000 --> 01:25:20.000
Most of the risk. scores used in criminal justice are actually simple in public.

01:25:20.000 --> 01:25:26.000
But that particular one is widely used and it's not it's proprietary.

01:25:26.000 --> 01:25:38.000
So yeah, I think if you don't work with all of you know, I noticed this, that people were refusing to work with like people were refusing to work with the police or work with you But the problem is if you don't

01:25:38.000 --> 01:25:43.000
help. somebody else will go and do it, and they will. They will not make the world a better place.

01:25:43.000 --> 01:25:54.000
So and I feel like I feel like it's important that we're willing to that We should be willing to look into things and study them, but we should

01:25:54.000 --> 01:26:01.000
We should be I don't know open and public about the shortcomings and the dangers, and and as opposed to like.

01:26:01.000 --> 01:26:03.000
We're not gonna touch this it's it's too much of a hot wire.

01:26:03.000 --> 01:26:11.000
I think we should. we should touch it. We just have to be really open about what the the dangers are and and try to help avoid them.

01:26:11.000 --> 01:26:15.000
I think that's I understand there's differences opinion on them.

01:26:15.000 --> 01:26:18.000
We have problems also, not listening to domain experts, which I find very annoying.

01:26:18.000 --> 01:26:31.000
So, for instance, like I, I found out once we started doing this computer aided mammography stuff that you have all these like people just trying to predict malignant versus benign on lesions from images like that's

01:26:31.000 --> 01:26:40.000
a standard machine learning thing to do but that's not what the that's not what would benefit radiologists, because if there's even a 2% chance that the things malignant you'd give the person a biopsy so what you need

01:26:40.000 --> 01:26:47.000
to predict is whether you should give them a You need to help them predict stuff that will align with their thought process.

01:26:47.000 --> 01:26:53.000
That helps determine whether to give a biosy so we tend to do that in machine learning just so that we can test our methods.

01:26:53.000 --> 01:26:56.000
But we're not actually helping anybody unless we work with their mean experts.

01:26:56.000 --> 01:27:09.000
That makes a lot of sense. Yeah, I like to say sometimes the the problems that are hard enough that we actually need computer help with are generally hard enough that they can't be solved with just computer help that we actually need to work together as a team

01:27:09.000 --> 01:27:14.000
people and and the machines Louis or tease I don't know if it's the louis cartes, I know.

01:27:14.000 --> 01:27:21.000
But louis or teas asked can you say a bit about your experience with the myth of a dental identifiability in social science and statistics.

01:27:21.000 --> 01:27:26.000
So I don't know that Phrase do is that something that's familiar to you.

01:27:26.000 --> 01:27:38.000
I don't know what the myth of identifiability is, but I can guess which means that the Rashomon set contains exactly one model, and I can tell you that it does not so even if you lowered the

01:27:38.000 --> 01:27:43.000

rashoman set down to all the way at the bottom, like even if you only permit one, you know, optimal model.

01:27:43.000 --> 01:27:49.000
Guess what with decision trees there's not just one optimal model, you can have many optimum, equally optimal models.

01:27:49.000 --> 01:27:53.000
So it's it's really identifiability in the sense of like.

01:27:53.000 --> 01:28:00.000
There is There's the identity of like truth and we can identify it, using just the noisy data that we have. It's not.

01:28:00.000 --> 01:28:06.000
It's not at least the the the studies that you have shown indicate that that's just not the case that if you're just trying to find a model.

01:28:06.000 --> 01:28:12.000
That does Well, there's heaps of them there's heaps of them. and that's the whole idea behind the movie Rashamon.

01:28:12.000 --> 01:28:20.000
Right that it's a Japanese movie and it's like some horrible murder that happens in the woods.

01:28:20.000 --> 01:28:34.000
And you get that story told by 4 different people are voices or something one's a ghost. but you get you get it told by in 4 different ways, and you just figure out like there's just no truth like there's just

01:28:34.000 --> 01:28:37.000
no, like everybody just perceives it differently. and so there is no no truth.

01:28:37.000 --> 01:28:48.000
So yeah, not into, not into that. You believe that the myth of identifiability is, in fact, a myth. I think, for a lot of I mean it might maybe I I don't know it depends on the problem.

01:28:48.000 --> 01:28:57.000
You can't just say something in general, so like whatever but But in the problems I've worked on like the rush month that really exists, and I can tell you that because I can find it fantastic. alright.

01:28:57.000 --> 01:29:02.000
So i'm being told by i'm gonna say my producer, cause it's cool to say that you have a producer.

01:29:02.000 --> 01:29:08.000
The blade Blaine is telling me that we're basically out of time, and I should before I thank you.

01:29:08.000 --> 01:29:17.000
I should remind everybody at the Nsf people that we have a One-hour office hour with Cynthia at 4.

01:29:17.000 --> 01:29:20.000
And so people who want to follow up and ask, and I did.

01:29:20.000 --> 01:29:23.000
I skipped a couple questions from my program directors who these are great questions.

01:29:23.000 --> 01:29:27.000
So hopefully. they'll be able to come and talk to you this afternoon.

01:29:27.000 --> 01:29:30.000
But Nsf people are welcome to that at 4 for everyone else.

01:29:30.000 --> 01:29:48.000
Thank you so much for being here Cynthia. just thank you for your work, and thank you for sharing it with us today. It's just it makes a big difference, and and i'm so glad you do it delightful to talk to