

January 19, 2022 – 2:00-4:00 PM ET  
NATIONAL FAIR HOUSING FORUM  
Mining the Data: Algorithmic Bias in Housing Related Transactions

---

\*\*\*

Speakers: Demetria McCain, Principal Deputy Assistant Secretary, Office of Fair Housing and Equal Opportunity, HUD, Jacy Gaige, Director of the Compliance and Disability Rights Division, Office of Fair Housing and Equal Opportunity (FHEO), HUD,

Moderator: Cashauna Hill, Executive Director, Louisiana Fair Housing Action Center

Panelists: Kareem Saleh, Founder and CEO of Fairplay, Michael Akinwumi, Chief Tech Equity Officer, National Fair Housing Alliance, Sacha Markano-Stark, Attorney, Relman-Colfax

\*\*\*

---

CASHAUNA HILL [0:00:00]: Good afternoon and welcome. We'll get started in just a few moments..

Good afternoon to those that are continuing to get logged in and join us for today's conversation.

We'll get started in just a few minutes. While we're waiting for everyone to continue logging in, we would like to have you respond to a couple of quick poll questions.

Please take a moment to respond to these questions that have popped up on your screen.

After you select your response option, please be sure to click submit. The first poll question is what are you most interested in learning about today. Comprehending the basics of artificial intelligence and machine. Increasing understanding of how data is used. Knowing more about third party technology companies and services used by the housing industry. Ascertaining solutions for eliminating bias in data and technology. Describing the applicable sections of the Fair Housing Act and regulations to other applicable laws. Recognizing features of key cases challenging tech bias in housing or that relate to housing transactions.

The second poll question is, which HUD region are you located in? If you're unsure of your HUD region, a map of all the regions is located on your screen. As people continue to login, please go ahead and answer the two poll questions that have come up on your screen. If you're unsure of which HUD region you're located in, please do check out the map that should be appearing on your

screen.

As we can see from our poll responses, 13% of today's attendees are most interested in learning about comprehending the basics of artificial intelligence and machine learning. 42% of today's attendees are most interested in learning about increasing understanding of how data is used in rental, sales, housing sectors. 7% are most interested in knowing more about third party technology companies and services used by the housing industry. 16% are most interested in ascertaining solutions for eliminating bias in data and technology. 7% are most interested in learning about describing the applicable sections of the Fair Housing Act and regulations to other applicable laws. 15% are most interested in learning about recognizing features of key cases challenging tech bias in housing or that relate to housing transactions.

We have 7% of today's attendees joining us from HUD region 1. 9% of today's attendees are joining us from HUD region 2. 21% are joining us from region 3. 12% joining us from HUD region 4. 10% from HUD region 5. 7% from HUD region 6. 9% from HUD region 7. 4% from HUD region 8. 7% from region 9. 12% are joining us from HUD region 10.

Thank you all again for joining us today.

Welcome to the National Fair Housing Training Academy forum titled NFHTA FORUM: Mining the Data: Algorithmic Bias in Housing Related Transactions.

My name is Cashauna Hill and I'm the executive director of the Louisiana Fair Housing action center where I lead a teamworking to fulfill the organization's mission to end discriminatory housing policies and practices through litigation and policy advocacy, as well as through providing Fair Housing trainings and foreclosure prevention counseling. It is my pleasure that HUD invited me to serve as the moderator of today's event.

Please note that this forum features information and examples that represent the experiences of the speakers. The comments today do not necessarily reflect the policies of HUD before we get started with today's conversation let's review technical tips and instructions regarding today's event. TJ, over to you.

TJ WINFIELD:[0:01:06]: Thanks.

If anyone does have technical difficulties with audio or video, we recommend that you first sign out of the webinar and then sign back in. If you're having schedule with that, you can request help in the Q&A box located on the Zoom panel section at the bottom of the screen or you can send an email to [nfhta@cloudburstgroup.com](mailto:nfhta@cloudburstgroup.com). We encourage you to ask questions. You can enter your questions at any time by selecting the Q&A button on the Zoom panel. Please note that due to time constraints we may not be able to respond to every

question today.

The webinar is it scheduled for two hours and is being recorded. The recording and the transcript will be made available on the NFHTA website on HUD exchange along with resources that supplement today's conversation.

Back over to you.

CASHAUNA HILL: [0:01:59]:

Thank you, TJ.

I would now like to introduce Demetria McCain who in September 2021 was appointed by President Biden as HUD's principal deputy assistant secretary for Fair Housing and equal opportunity. Prior to serving in this capacity, Demetria McCain spent 15 years, most recently at president, at the Inclusive Communities Project known as ICP, which is a Dallas based affordable Fair Housing organization.

Demetria McCain has been a great friend in the Fair Housing community, as I'm sure you all know and is very deeply committed to this work.

Demetria, thank you for joining us today.

DEMETRIA McCAIN: [0:02:38]: Thank you. Good afternoon, good morning, everyone.

Let me start by sending all of our Fair Housing partners and friends to the movement good wishes for a safe and healthy and productive new year.

I'm told it is too late to say happy new year to folks but I'm going do it anyway!

Listen, I really continue to be proud of your work and our partnership to really embed Fair Housing equity and housing related transactions and systems throughout the whole country. As I said before, it is not just that you engage, it is how you engage.

You show up sharing your passion and experience in this fight for housing equity and I thank you for that.

Now, please, please keep on growing partnerships, friendships and enriching the communities in which you use your unique talent.

This is how we will win the fight for housing equity. I didn't say if, I said will.

Of course, our work must begin with meaningful knowledge, making sure we

grow our skills and capacity, and over the last year a whopping 624 of you participated in 16 four-day NFHTA courses and over 4,000 of you joined 7 National Fair Housing forums to discuss emerging Fair Housing issues.

I encourage you to continue sharing the value of NFHTA's education with your colleagues, especially those who are new and early in their careers. That takes me today's forum on algorithmic bias and housing related transactions. This topic, and I'll take notes, trust me. This topic is a terrific example of how we as Fair Housing practitioners must remain engaged in the evolution of how our housing providers, speakers, how they interact and how we ensure Fair Housing choice in the transaction.

We now live at a time where we are called to not only combat discriminatory presence in print and online media, but we're simultaneously called to understand and combat how artificial intelligence and machine learning may deny whole sections of our community's housing choice because of one zip code, race, sex, income, and other classes.

So it is up to each of us, Civil Rights professionals, to understand the very complex challenge and to discuss the practical ways that industry, government, nonprofit, philanthropy, and others can align to really kind of harness technology's value in such a way that it creates equity rather than depriving people from opportunity.

In this spirit, I invite you to please, please be in touch with our team here at HUD as you identify potentially complex and novel ways of discrimination appearing to be in your very own community. As was stated during the commemoration of Martin Luther King, Jr., we must all be warriors, the struggle continues, and it starts with us.

Again, thank you for prioritizing today's training and our time together. Now I'm going to go ahead and give it back to you.

It is all yours.

CASHAUNA HILL: [0:06:18]: Thank you so much for those remarks and for joining us today.

I would now like to introduce Jacy Gaige, the director of the compliance and disability rights division in the Office of Fair Housing and equal opportunity.

Throughout her time at HUD, Jacy Gaige has focused on systematic discrimination and has led investigations of redlining and lending scams, tenant screening practices, not in my backyard denials of affordable housing, distribution of disaster recovery assistance, nuisance ordinances and more.

Jacy also led the department's investigation of Facebook's advertising platform which led to a charge of discrimination in March 2019.

Thank you for joining us today.

JACY GAIGE: [0:07:04]: Hello. Thank you, everyone, for joining.

Myself and so many of us at HUD are really excited about today's topic. Big data and the algorithms they power are the present and future of housing access. As said, we in the Civil Rights community need to pay attention to understand and figure out how to solve the problems they present.

Data and algorithms already drive how people find out about housing opportunities, who has access to financing and on what terms and who gets approved as a tenant.

However, algorithms are only as good as the data they are fed, and no data in this country is untainted by a long history of discrimination.

I'll talk briefly about one example to tee up issues our panelist also offer more insight on.

HUD charged Facebook with discrimination in 2019 and that case is currently with et cetera the Department of Justice for litigation.

Facebook shows billions of ads daily, they have hundreds of millions of users in the U.S. alone and they also show ads on thousands of partner websites. You see Facebook targeted ads even if you are not on Facebook. Of course, the business model is to charge advertisers for the ability to target the ads.

The explicit ad targeting options that Facebook offers advertisers received a lot of attention in 2016 when Pro Public published an expose showing researchers could run housing ads, excluding users of specific ethnicities among other troubling categories.

The targeting system is even more troubling when you look under the hood about how those categories are created and how Facebook selects users to see a given ad.

If an advertiser says they want to show their mortgage ad to all persons in New York Facebook does not show the ad to all users in New York. They show it to a subset that they, Facebook, deem most likely to engage with the ad. As Facebook itself explains, it uses --and these are quotes --age, gender, language, purchase behaviors on and offline, interests, activities, pages liked, where you live, the places you like to go, the businesses and people you are near and others to show people the ads most pertinent to them.

While users may voluntarily disclose some of this data about themselves, like name and gender when they create a profile, they disclose most of this data unwittingly through the actions they and the people associated with them take as they go about their daily lives.

Or as my colleague put it, walking back from a meeting on this one day, I think we should throw our phones in the river. The truth is, we can't get away from big data and its biases and we need to take them on. Evidence has shown that online ads are delivered to audiences severely biased on protected characteristics. Facebook is not necessarily unique; all major online advertisers deploy at least some of these practices and similar processes are playing out well beyond advertising.

A number of fascinating studies have punctured the bubble of naive hope that taking people out of decision making would take the bias out too.

For example, in a study out of Berkeley from 2018 researchers compared outcomes from in-person and known lending where race was in the known. They found the profitable form of lending discrimination, charging some buyers more for loans was replicated if not increased in online lending while the unprofitable form, denying creditworthy borrowers was less so.

The AI used an online lending determined which borrowers were less likely to shop around for loans and charged them more. The key takeaway here is that discrimination can be replicated when profitable, no matter how sophisticated the system, or in fact because of the sophistication of the system.

We need to look deeply at what the goals of the machines we are building are knowing that those machines are controlling access to housing.

FHAO is continuing to delve into the issues against tenant screening across the cycle and more and it will take all of our smartest minds to figure out how to counter discrimination on the platforms. We are so excited to hear from some of those minds today.

I'll turn it back to Cashauna Hill.

CASHAUNA HILL: [0:12:13]: Thank you so much for your time today, Jayce and for the insightful remarks. Thank you.

As we move on, I would like to share the learning objectives for today's forum.

Together we will comprehend the basics of artificial intelligence and machine learning ; increase understanding of how data is used in the rental, sales and lending housing sectors ; know more about third-party technology companies and

services used by the housing industry, rental, sales, lending ; ascertain solutions for eliminating bias in data and technology ; describe the applicable sections of the Fair Housing Act and regulations to other applicable laws ; recognize features of key cases challenging tech bias in-housing or that relate to housing transactions.

At this time I'll introduce our panel speakers. We are eager to learn from the experiences of these speakers and you can find all of their bios on the forum page of the NFHTA website.

Joining us today we have Kareem Saleh, Michael Akinwumi and Sacha Markano-Stark. Later in the forum I will ask a few questions of each panelist and then open it up to other panelists for additional comments.

I would like to remind everyone that during today's roundtable discussion you will all have the opportunity to submit questions and we'll do our best to answer later in today's event.

However, please note that not all questions will be able to be answered and we will not be able to address any personal questions.

You can submit questions at any time during today's conversation via the Q&A box.

Also, as a reminder, this event was being recorded and all materials, including the slide deck and the event recording will be available on the forum page on HUD exchange soon after the event.

With that, we will start with our first panelist, Kareem Saleh A.

KAREEM SALEH: [0:14:27]: Hey, everyone. My name is Kareem Saleh, I'm the founder and CEO of Fair Play. Delighted to be here today to give you a conceptual introduction to machine learning. Let's dive in. I'm going to attempt to use analogies and stories to express a little bit about how machine learning works.

I'm hopeful that you all can see my screen and I'm going to dive in.

Imagine that we wanted a machine to make a pizza. Well, in the analogue world, and in the world of early computer programming, we would start with ingredients and a recipe, and our hope would be that by applying that recipe to those ingredients it would yield a pizza.

A recipe in this context may also be referred to as a model. We have the ingredients for the pizza, the recipe is a model for how a pizza is made and how it ought to look. We hope that by applying those ingredients to that model, to that

recipe that we'll get a delicious pizza.

What if we don't have a recipe. What if we just have a lot of ingredients and as sometimes is the case in my household, my domestic partner turns to me, says there is a bunch of stuff in the refrigerator, please make dinner.

I'm responsible for transforming these ingredients, these variables into dinner, but I don't have a recipe.

That is fundamentally what machine learning does. It looks at a set of ingredients, sometimes called inputs. It looks at a desired outcome, like a pizza. It figures out the recipe for turning those inputs into the desired outcome.

If you think about it for a second, this is also how lending has worked. Lending decisions used to be made by humans. You can think of the loan applicant as being the ingredients in the recipe or the variables, and this is a banker, a loan officer, he's in this case trying to apply a recipe to the loan applicants' ingredients to determine if that loan applicant will pay back a loan or if they'll default on the loan.

For many years it was the human loan officers that attempted to develop and apply these recipes to the variables, or the inputs submitted by the loan applicant.

As I'm sure many of you know, you can see the obvious problems with having a human be the one trying to make this recipe up especially on the fly. The human may take into account all kinds of variables that as a society we feel are inappropriate or are not reasonably related to creditworthiness which is a thing that this loan officer is trying to predict. Right. A human loan officer may take a look at a loan applicant and say well, I know you from the neighborhood. You kind of look like me. You know, I know your kids from school. These are all variables that this loan officer may take into account as a part of their recipe to determine whether or not a person should be given a loan. I think we probably all agree that those are not variables we want deciding our futures, both individually and as a society.

In the 1980s, we developed a better way of making lending decisions. We started using math. The idea was that if the math would be more neutral, more objective than humans in making lending decisions. You can kind of see how that might be the case.

Here is an example of one of the very primitive, very early lending algorithm, loan approval recipe let's call it.

You can see here the loan approval recipe starts by ingesting credit history as an input, as an ingredient. Is your credit history good or bad? If it's good, we

then -- the model then looks at the second variable, what is your income. Is your income low or high?

If your credit history is good, and your income is high, you can see that the algorithm recommends that you be approved for a loan in either a big or small amount. Whereas if your credit history is good, your income was low, the algorithm recommends that you be approved only for a small loan, not a big one.

On the other side, you can see that if your credit history is bad, and then we look at the income, your credit history is bad, your income is high, the model suggests that we be approved for a small loan but not a big one. You have good income, bad credit history, approved for a small loan, whereas if you have bad credit history and bad income the model recommends that you neither be approved for a big loan or a small loan.

This in theory has some -- this has some real advantages over a human. As we'll see, it also has some real limitations. One thing to note about these relationships between credit history and income and loan amount is that it all seems kind of straightforward. The variables seem like they're connected in a very logical sequence. We look at your credit history, we look at your income, we look at how much money you're seeking for a loan and then we make a judgment about whether you should be given a loan and whether you should be given a big one or a small one.

All in this very kind of linear connected sort of train of rational thought. As we all know, the world is not linear. The world is not stray forward. The world is complex. It has a lot of nuances. You can see on the left here that the old credit models used to assume what's commonly called linearity between the variables, that the relationships between the variables are straightforward and could be understood by humans. By drawing a line through those straightforward variables, you could predict a model of who would pay back a loan, the blue dots, and who would default on a loan, the red dots.

You will see that that model on the left, the one that reflects that decision tree thinking that we saw on the previous slide, it makes a mistake, right. Assuming that all of the relationships between ingredients, variable, it is a linear relationship and attempting to draw a line that classifies good borrowers from bad borrowers, it leaves some red dots on the side of the line where folks are being approved for loans and it leaves some blue dots on the side of the line where folks who are being declined for loans reside. A model that assumes these kind of straightforward linear relationships between ingredients is going to make some mistakes. It will deny some folks who should have been approved and it is going to approve some folks that should have been denied.

The magic of machine learning is its ability to understand those nuances and in effect to draw the line in a way that does not -- draw the line that separates good

borrowers from default borrowers in a much more precise way. It draws its line between good borrower and defaulting borrowers in a way that reflects the nuance and complexity of the world.

That's really the great improvement between these kind of decision trees which could be an improvement over humans, but which also make some mistakes to the new world of understanding the complex ways in which variables and the distribution of good borrowers and bad borrowers may not be susceptible to straight line classification.

Let's dig in how they're able to drive these lines and draw these nuances.

Let's try by building a model, an algorithm of our own.

Imagine for a moment that we're trying to build a model that will predict gender. What if I told you we were going to give you height as an input? Height as an ingredient to that model, to that recipe.

Well, height is somewhat predictive of gender because on average men tend to be taller than women. If we all think about it for a second, height is not perfectly predictive of gender. Right. After all, -- oh, what just happened there?

Sorry about that.

Height isn't perfectly predictive of gender because, of course, as we all know, there is some tall women in the world and some short men in the world. Building an algorithm that attempts to predict gender on the basis of height will be somewhat predictive but not perfectly predictive.

What if I said, okay, let me give you another ingredient into that recipe, another variable into that model.

What if I said in addition to height, I'm going to give you weight to help predict gender? Well, including weight as a variable does enhance the predictive power of our model somewhat because even at the same height men tend to be taller than women due to things like testosterone and bone muscle density. Including weight, in addition to height, it adds predictive power to our model but is still not perfect, every model that predicts gender on the basis of height and weight will assume that every child is a woman. A model that seeks to predict gender on the basis of height and weight will misclassify every child as a woman. What if I said to you, okay, fine? In addition to height and weight, I'm going to give you birthdate as an input to this model.

Now our system for predicting gender is actually looking pretty good. We have got height and weight which are reasonably predictive of gender. Now we have got birthdate which helps us isolate out children, our model for predicting gender

is actually pretty good. If I had told you at the start of this story that birthdate was predictive of gender, I think a lot of you would have told me I was crazy. That's the power of machine learning. The power of machine learning is its ability to understand the ways in which variables interact to encode information that no man could possibly discern on their own.

The power of machine learning is to understand that height and weight and Garrett encode gender. Now, that was kind of easy for us to figure out when we only had three variables. What if we have 1,000 variables. We couldn't have walked through an example like that together with 1,000 variables, and there is no way a human could look at 1,000 variables to understand how all of those variables were interacting to encode information.

A recipe that attempts to map a lot of variables on to an outcome is going to be so complicated that no human could understand it. That's what we sometimes in the machine learning and artificial intelligence world call a black box. We know the algorithm took inputs, reached an output, often a very accurate output but we don't know how it did it. The computations involved and the number of variables involved were so vast that no human could actually understand what was going on and we refer to that as a black box. Why is it really important to understand what's going on?

Let's take a look at an example.

The problem with not understanding what's going on in the black box is that the risk that's seemingly fair variables, it will interact just like height, weight, birthdate in ways that we don't understand, that result in the black box or the machine learning algorithm making decisions on the basis of information that it ought not to have.

Let me give you another example.

Imagine that you are a used car lender. You are trying to build a model, build a recipe, build an algorithm to determine who will pay you back.

Let's imagine that you're looking at the list of variables you want to include in your credit, underwriting model, your credit underwriting recipe and your business colleagues have recommended two variables for inclusion.

The first is what is the mileage of the car that's row posed to be purchased.

You can see why the mileage of the car that's proposed to be purchased would have a reasonable relationship to creditworthiness. After all, the higher the mileage the car, the more likely it is to breakdown, the more likely the car is to breakdown. The more likely it is that the owner of that car will not be able to get to work and those -- you know, the inability to work may have an effect on their

income and it may cause them to default on that car loan.

You can see how including the mileage of a car, the model, is reasonably related to the thing we're trying to predict, and it is not obviously discriminatory. It feels kind of neutral. We're not looking at the characteristic of the applicant who's buying the car, we're just looking at the condition of the car. This variable would normally not pose a problem to a Fair Housing perspective. It kind of – the mileage of the car appears to be kind of neutral, objective imagine the colleagues on the business side also said, okay, the other variable we want to include in the model is whether the applicant is from Nevada. We have noticed that folks from Nevada default at a slightly higher rate than the national average. Let's say there's something, you know, going on economically in Nevada that causes that to be the case, maybe it's Nevada's rich history of hosting the gaming industry, maybe Nevada people are more risk seeking than the average American, but states are really big. There is no obvious redlining concern, it is not like they're coming to you, saying we don't want to lend in this neighborhood in Nevada. Just saying that am captain exhibit a slightly bigger risk of default. It is a big state, no redlining concern, the numbers are what they are.

You have the two variables, both of which appear totally related to creditworthiness, both which feel at least facially kind of neutral or objective. They don't obviously discriminate against any one group.

Let's put them in the model, right. Turns out if you're buying a high mileage car in Nevada, it is more likely you're a person of color. That's why we care about what's going on inside of the black box. We have given variables that seem to be neutral, mileage of the car, you the state you live in, the two variables encode information about race and there is a real risk that the algorithm is made based on race rather than the seemingly neutral factors.

Another issue with algorithms, we would be advised to keep in mind, it is that every algorithm must be given a target. The output that it is trying to produce.

Will this individual pay back their loan? Will this recipe and ingredients yield a pizza and the thing that makes machine learning and algorithms unique is that they relentlessly define their recipe to achieve the best possible outcome. They adjust their recipe over and over again to better accomplish their target or objective.

If you think back to one of the earlier examples, the idea of the ingredient, the recipe, the pizza, maybe the first pizza that the algorithm makes won't be super delicious. Right. Over time it refines the recipe until that pizza is from the greatest pizza shop ever every algorithm must be given a target. We can draw an analogy here, it may be helpful to social media as many you have may know, social media algorithms seek their target, the target they're given, it is to maximize your engagement, maximize the user engagement with the social media platform.

You know that algorithm that seeks to maximize your engagement, it is going to show you engaging stuff to keep you interested regardless of whether or not the stuff that it is showing you to keep you engaged is good for your mental health or good for society. You didn't tell the algorithm to care about what's good for your health or society. You just told the algorithm the objective was to maximize engagement. It is showing you all kinds of stuff to keep you engaged without consideration for the other harm that the algorithm may be doing.

You can see how this -- this one would obviously have problems in other domains of our lives. Right. This idea of giving an algorithm a kind of single-minded focus on one objective. What if it was -- what if we talked about a self-driving car. What if we gave the self-driving car the simple target of get the passenger from point A to point B.

If that was the only target you gave a self-driving car, the car may get you from point A to point B while driving the wrong way down a one-way road or blowing through stoplights or driving in a manner that causes mayhem for peds as you can see here. It is not sufficient to give the self-driving car a single-minded objective of getting from point A to point B. We have to give it two objectives. We have to give it the objective of getting the passenger from point A to point B while also respecting the rules of the road. Giving the self-driving car, the algorithm powering the self-driving car two goals. Get the passenger from point A to point B and do it safely the good news is, we can do this in financial services too.

We have been talking about financial services model that attempted to predict whether or not you were going to pay back a loan. That's the primary target, will this person pay back a loan. We could give it a second target just like the self-driving car. What if we said predict who is going to give back a loan but do it fairly? Do it in a way that minimizes the disparities for one group relative to another for let's say black Americans relative to white Americans or female Americans relative to male Americans.

Here is why as housing advocates we shouldn't really care about this.

What you're looking at here, it is the State of mortgage fairness to women in America in 2020. We define fairness here, the exact same way that courts and regulators do in terms of something called the adverse impact ratio. This measure asks how often are women approved relative to men. How often is one group approved active to another group. What you can see here is -- well, so, in this instance if women are approved at 90% of rate of men, you see green. If women are approved between 80 and 90% of the rate of men, you see yellow. If women are approved at less than 80% the rate of men, you see red. There is a lot of red and yellow on this map, especially the South and the great Plains. Do you want to see what it looks like for black Americans? In many parts of the

American mortgage department, black Americans are approved less than 70% of white Americans for mortgages. These lending decisions used to be made by human loan officers. Today they're made by algorithms that appear to be replicating down to the county level the discrimination of the past, the redlining of the past.

Could we do better, the answer is yes in part by constraining -- by giving the algorithms two targets as we have recently discussed. Let's give you -- to give you an idea of what the results could be, what you're looking at here, it is a plot of different algorithms and the X axis, the bottom is the accuracy of the algorithms and the Y axis, the part that goes from bottom to top, it is the fairness of the algorithms. What you can see here is that the conventional modeling methods, the decision tree I showed you and the off the shelf AI technique, a neural network, we can see that they are extremely accurate and they're not very fair. In this case, they're approving less -- they're approving black applicants at less than 60% of the rate of white applicants. Now, what if we gave those algorithms two targets like I have told you about. Predict who is going to pay back a loan while also being fair to all groups. What would the result of that algorithm be? I have got good news for you.

The algorithm that has two targets, predicting default but also being fair is slightly less accurate than the algorithms that have a single target but look, their fairness, it is almost double.

We could increase mortgage approval, mortgage approvals for black applicants dramatically while still achieving levels of accuracy that are very, very close to what lenders are able to achieve today not only that, but when we give the algorithm two objectives for pricing those mortgages you can see that it is possible to give algorithms the objective of pricing mortgages both commiserate with the level of risk involved and also more fairly and what you see here, when lenders do that, there are these models that are fairer models that are both more profitable for the lender and more fair to the applicants by reducing their pricing. The other thing that's really interesting from, the algorithms are built using the same ingredients. As you see here, this is a list of the ingredients that these algorithms are taking into account to reach the conclusions and you see the old school mathematics is able to approve the black applicants at 60% the rate of white applicants and those are the variables that's being taken into account. You can see the off-the-shelf neural network, the off the shelf AI technique approves black applicants at a slightly lower rate, 56% the rate of white and taking the virtual same variables into account.

You have the fair AI technique on the far, right. That one is approving black applicants at 96% the rate of white applicants and it is again taking into account virtually the exact same variables. What that tells you is two things, A, that the increase in fairness is coming from the mathematical technique and not from any new crazy data source, wild data source, it is using the exact same ingredients in

its recipe that the old school math was using and that the unconstrained neural network was using.

What is driving the AE takeover of all of the high-stake decisions.

AI has been talked about for years as being the technology that was soon going to be running our lives.

For many, on many occasions that hype proved -- there were things called AI winter, times when things were optimistic about the advance of AI, only to be disappointed by its limitations.

It is very, very -- I'm very, very reluctant to say this time is different, it is always risky to say this time is different. There does seem to me that there are some things that changed in the world that make AIs advance more likely to succeed now. There are the overwhelming amount of data that were all generating -- that we're all generating every single day. There are advanced mathematical techniques that are common for universities and other self-taught researchers who are finding more and more complicated ways of building predictive models. Finally, the computational power that it requires to run these advanced algorithms on a lot of data has become basically free and infinite.

It is the unity of data and math and Cloud computing that make AIs advance more likely to succeed across a bunch of domains and my good friend Michael Akinwumi will now illustrate for you the ways in which AI is taking over housing, real estate, lending, many of the domains we all care about.

Thank you for your time.

I look forward to your questions during the Q&A.

CASHAUNA HILL: [0:45:23]: Thank you so much.

As we move on to the next panelist, a couple of reminders, feel free, everyone, to submit your questions in the Q&A box. We'll try and wrap up this conversation on time so that we can get to your questions this afternoon, we'll just keep going ahead so that we can get to the questions. Please also note that all of the event materials will be made available after today's event. The NFHTA resource page at the HUD exchange. So you will have access to everything that we are discussing today.

We will now move on to our next panelist, Michael Akinwumi, who will speak of the use of data and technology in-housing and related markets, third-party technology companies and services used in housing related transactions. Michael.

MICHAEL AKINWUMI: [0:46:20]: Thank you so much for that introduction.

Good afternoon, everyone. I'm going to continue the conversation from where Kareem Saleh stopped. I would like to clarify that, you know, in my presentation, whenever I say or mention to "this" I'm referring to the machine learning model, the general artificial intelligence. In addition, I will use "AI/ML" several times in my presentation and whenever you hear me say "AI/ML" I'm actually referring to artificial intelligence or machine learning.

In terms of the outline of my presentation, I will walk you through how data and AI/ML are being used in rental, sales, lending housing sectors, and from there move on to, you know, instances or third-party technology companies and the product and the services that they're actually using in the related transactions, and last but not least I'll talk about some solutions for removing bias in data and also algorithms itself.

So as of 2018, the U.S. census bureau showed the median net worth for African-American families is 9,000 U.S. dollars and \$12,000 for Latino families and that compares to -- or contrasted to \$132,000 for the white family. I would like you to pause and think about what median statistic actually means just for a moment.

This means Latinos and African-Americans face uphill battles when deciding what to eat which has huge health indications and also when deciding what quality of education they can afford, which implications are for income earning capacity or when deciding if you own a car and if they should, what type of car they should own. Then you determine how far they want to live from work, what insurance premium they can afford and their contributions to climate change.

So this is choices that are constrained by people's financial stability.

An interesting thing, in the United States, wealth and financial stability, they're linked to housing and -- housing opportunities and Homeownership according to the CEO of the National Fair Housing alliance.

What it means, where people live, what houses they own reflects the quality of food they receive and paying for quality education or driving electric vehicles and making lifestyle choices that could elongate their lives and give back to the communities.

It is no surprise that we're here to pay attention to data and that big companies pay in people of color access to housing opportunities. If we want to ensure that we're not making recipes with that have disparate impact or discriminatory recipes or generating impacting communities of color, we have to have a close look at instances where AI/ML is being used in housing-related transactions.

Once we look at how data and AI/ML are used in rental, sales, lending housing transactions, so I'll walk you through a couple of instances and examples in each of these different sectors. There are multiple applications of data and AI/ML in rental market, and we'll focus on these three examples. The first one, when it comes to rent and leasing, for example a data scientist, others may exploit the rental listing to the advantage or to the detriment of consumers. For example, AI/ML may be used to develop a natural language processing model to detect listings where the base range and the concession rate are different and then decide what the user interface had, they could decide that they want to show the base rent or the compensation rent depending on the profit appetite. Imagine a scenario where the data scientist can detect a perspective renter's rate, age, sex using the web traffic data. And for the landlord, the data scientist generates leads that meet a discriminatory condition like -- this is actually a real advertisement. It reads like this. Not to sound racist, because we want to make things more understandable for our younger child, we would like to house white children. That is a real ad, a real advertisement, that a data scientist would decide to feature on the rent or listing and then the web developer may feature an adversarial rate in addition to the discriminatory advertisement which is in the listing and if it shows the end of the block of texts or is listed what the rate is listed, the base rate, the base range, before or after concessions have been populated. That's one example or instance where the data or machine learning could be used.

Another example is online screening.

So the screening process, it is for data miners and others alike, especially during the Pennsylvania. AI/ML screening assistant are trained to learn about user behavior from an existing management system and compare to a human monitor who can only learn so much from their time with a potential renter.

AI can get the historic data from a user profile when they apply for a new proxy. Then it is going to roll a forecast of fallacies to see the risking in the code the potential tenant would be. However, there is a reason that risk is in quotes, it is because it is up to the landlord to decide what constitutes a good versus bad tenant, what data to use to develop an AI/ML model that predicts the likelihood of a perspective tenant.

For example, when it comes to range of market, it is maintenance service requests. So maintenance has always been a headache for tenants. With AI powered chat bots or virtual assistance, the request order could be scheduled at any time. The AI will send your request to the maintenance right away.

While landlords and property owners, we promote the request system as a win/win for everyone involved in the transaction. Our Fair Housing advocates, we would like to think of problems that the tenant will face with the service.

For example, tenants with language barriers, they may be frustrated by the

inability of the chat bots to process their request or its inability to connect them to the maintenance.

This is one of the many issues with voice recognition technologies or natural language processing solutions in general. In addition, tenants without easy access to infrastructure, high-speed Internet or efficient computer, they may be at a disadvantage.

When it comes to sales, I think, you know, coming from the previous speaker, we have seen instance, the Facebook instance for example, where we see how data, machine learning is being used in sales.

We're going to focus on these three different examples, because these are incidents where we have seen machine learning model being integrated into web applications.

When it comes to predicting housing prices, the Washington post article forecast, as trained for 2022, and mentions roles that high bias plays in the real estate industry. So the quotes in the articles, higher buyers, real estate company, they allow consumers to basically buy and sell on demand. They will buy your home for a price -- in this case, I'm referring to the machine learning model, AI or other models, and they'll buy your home for a price that's algorithm, they say it is correct and allowing you the freedom of making a cash like offer. An example of high-buying companies, it is zero. So this is aimed to make offers to the customers based on the predictions of the machine learning models and perform miner APS and they have the profits.

An example that I'm going to cite in terms of how data and machine learning is used in sales transactions, it is automated evaluation models.

So one of the most asked questions from home sellers is what is my home worth. You go to the market to sell your home, that's probably the very first question that will pop up on your mind.

Human agents usually give the answer by experience. It is so hard to find the proof to support that judgment, the judgment.

So a recent review of the uniform standards of the professional appraising practice and the real property appraising criteria by the National Fair Housing alliance revealed historical bias in our practices. So the estimated model, it is a good attempt of using AI or machine learning model to actually predict or estimate the worth of a property.

In this case, AI can take a much larger amount of data than a human can handle, just like our previous presenter actually mentioned and give a fair result of strong data -- when I say fair, it's in quote because I'll touch on that later.

Developing an AI/ML solution on this, on bias, on the data, the racially segregated data will scale appraiser bias to the detriment of the homeowners and by the next example, the instance where data and machine learning model is being used in sales is the real estate family. When it comes to the supply of homes for sale, you know, we have each sale, especially since the start of the pandemic. This makes today's housing market ultracompetitive for real estate agents.

So agents are desperate to find new listings and some of them belong to AI/ML for some help.

So unlike traditional agent who North Carolina on the doors of a lot of homes, all AI/ML, this could be used to find geographic areas where homes are most likely to sell or find individual homes that are most likely to sell in a specified time window.

If one factors in the history of gentrification, racial segregation, appraisal of buyers and redlining in the data being used in this real estate use case, it may be easier to use how AI/ML is scaled to the bias that may be inherent in such data.

The next example that I'll focus on is using data or machine learning model in lending.

You know, we would like to have the marginal picture, a customer joining when it comes to becoming a homeowner. It usually starts from advertising, outreach, and then once the customer reaches the potential lender, then the next thing is preapplication inquiries, and then followed by the loan underwriting and enterprising and -- and pricing and lastly, you have the loan administration. I'm happy to actually tell you that each of the faces actually has use cases for data and agreed learning.

If you start with the first one, it is advertising and outreach, AI/ML uses powerful but questionable and potentially discriminatory algorithms in conjunction with optimization preferences of real estate agents so recommended housing advertisements that meet the customer's needs. For example, the National Fair Housing alliance and three member organizations filed a lawsuit against Facebook in federal court in New York City alleging that Facebook's advertising platform enables landlords and real estate brokers to exclude family was children, women, other protected classes of people, of color to receive the advertisement.

Why the Fair Housing group entered a lawsuit with Facebook, the lawsuit is a clear example of how an AI/ML solution may be unlawful and may become a tool that causes damage to protected groups.

Another example, use case here, it is preapplication inquiries. And the preapplication, use cases that come under the guise of the fraud detection model. That's what we have seen in practice, the fraud detection model.

Natural language processors, you know, competing with loan officers, loan counselors, you know, for example, if the perspective home buyer makes the tying and outreach phase of the mortgage lending process, the usual next step is for the prospect to inquiry about qualifications and also terms and conditions. You know, the subjective phase has been replaced gradually with AI/ML technologies like chat bots. The AI chat bots want to make human conversation by picking up on conversation cadences and can react or respond to both written and spoken prompts.

Google, Alexa, Google, Chrome, it could actually resonate with you.

Based on responses from the prospect and the recommendation models that are prebuilt in the chat bottom, it could discourage the consumer from applying or from predictive tear lenders and to look at the prequalified consumer based on business practice, again, business practices or fraud detection models or the fraud model or advise the submitting the application for mortgage financing after providing a minimal list of documents that are required for the application.

In each of the chatbot's decisions, it is liable to buyers.

Next, when it comes to -- assuming that there is a prospect that makes it past the advertisement and outreach phase and also the preapplication inquiry phase, then followed by, you know, the writing and the pricing, the previous speaker had said, when it comes to using AI or machine learning, to appraise creditworthiness of perspective applicants, so a type of machine learning algorithm used in this phase, it is called as you were advised learning algorithms. I call that out, because I like -- later on I will also speak to the relevance of targets and the profile of people that actually are involved in the defining what the target is, which are in my opinion, the previous speaker actually did a very good job of trying to explain why that is required. Almost any credit scoring solution that relies on data from credit reporting agencies utilizes some form of supervised learning or learner to arrive to applicants. However, if a machine learning solution is not intentionally designed to filter out histories of racial discrimination in the underlying credit data, it may become a weapon of credit disruption or a recipe for discrimination or another detector skied as an algorithmic system. -- disguised as an algorithmic system. Let's see that people of color are going through the loan on the right-hand process that's now being driven by machine learning model.

After the algorithmic solution decides that the applicant and a search recommends that a loan be approved and other type of machine learning algorithm, the unsupervised learning algorithm may be used to recommend terms

and conditions for the loan. This usually involves using the historical data to group credit consumers based on credit attributes and the lender may then choose to use average terms and conditions in a group the prospect belongs to as the recommended terms and conditions.

Again, the type of algorithm used, the representative of the underlying data and the business policy used to act on the algorithms will determine if the loan is predator or not. predatory or not.

The last use case I will give when it comes to lending, loan administration. The lender may choose to use AI to streamline their loan servicing process. If the borrower defaults on the loan, the lender may choose to use AI to recommend a loss mitigation strategy that meets its own overall tolerance level.

The lender may train a classifier, just an example of supervised learning algorithm on the historical delinquency and the default records to choose from forbearance, repayment loan, short sell, loan modification as a loss mitigation strategy.

This choice may further be augmented by a loss given default model. What it means, that the customer eventually defaults, what is the exposure, and they could bring in an AI or machine learning model that will predict the loss of what the default will look like.

This is based on traditional statistical models or even supervised learning algorithms.

It was explained, instances, use cases for AI or machine learning or data in the different sector.

I think we should try to spend time to look at third party technology performance and what services or products they actually bring to us in market.

It may be instructive to have the technology needs involved in transactions by the services and products provided along a customer's own ownership journey. So I mentioned in the previous slide, it starts with advertising, generation, outreach, so with the purpose of clarity, I'm going to exclude this category, the categories of advertisement, generation and outreach.

When it comes to machine learning, I really like the analogy of ingredients, recipe, also output. You want to think of the ingredient as data. We have a group of components that actually focus on data services and there are other credits and three national credit reporting agencies, so they want to make sure that they actually use what is available on credit reports of a particular applicant to develop that.

They will also have a group of companies that the business value proposition is just developing, focused on the model that leverages it's the data that's available to the applicant and examples of FICO, you also have the network compass and other similar start-ups.

We have companies that, you know, they look at -- assuming what you actually have, the ingredient, the poisonous ingredient, and you also have the discriminatory model, and you want to make sure that the customers -- they're not eventually exposed to outcomes that would actually impact them.

I group this on the fairness of the services, and that's where the companies here, they're common. Last, in this category, loan services. We have start-ups and even established companies that focus on using AI to manage the relationship once the loan is actually operational. As for capacity, actually that is covered.

When it comes to explained gateways to bias, discrimination, prejudice in AI, machine learning, as well as the data, they're not just talking about what are the problems or potential issues are. I would also like to offer based on the research that we have done, what are the tech solutions available out there and also to build on some of the responses that NFHTA has given to some requests for information.

Generally, there are at least three gateways for bias and discrimination to manifest in an algorithmic solution. Those gateways are data. Data, that's the ingredient, the algorithm, it is the recipe.

Also, the outcome, which is the decisioning on the outcome of the model.

So, it makes sense to actually categorize tech solutions that tend to remove bias or discrimination from machine learning pipeline into three pockets. The three buckets, they're processing, in processing and generally, the processing techniques, the focus, it is on modifying the data that comes into the pipeline in order to overcome any inherent underrepresentation in the data. Whereas in-processing methods focus on constraining the machine learning algorithm to meet a specific training and the post-processing methods, focus on removing the disparate impact from the model output.

In addition to the tech solutions, the National Fair Housing alliance has advocated for the hybrid approach that involved consumer groups, consumer advocates and also data scientists. Merits of this collaborative approach are documented in the responses to request for information for financial regulators, National Institute of Science and technology and most recently, the Office of Science and Technology Policy.

This collaborative approach is very important, for example, when it comes to deciding what the output should be.

For example, should we actually involve consumer -- should we not involve consumer advocate groups or consumer groups in general when it comes to deciding what constitutes bad versus good loan, or deciding what features or ingredients should be used in the data.

For example, when it comes to data, right, we all know that gender is not just about male versus female, but if a data scientist who is not really aware of the social context comes into using the data blindly, some of the conversations will probably not happen and they'll handle the developing algorithms or models that would actually cause disparate impact or treatment to the consumer.

With that, I'm going to pass it on to Cashauna Hill and I hope that will be more time for me or us to take any questions that you have.

Thank you.

CASHAUNA HILL: [1:13:28]: Thank you so much, Michael.

As we prepare to hear from Sacha Markano-Stark, I do want to ask you a quick follow-up question, Michael, if that's okay. I think this is a great time to just clarify that was vocabulary question that came up.

We had a couple of attendees asking about the term concessions rate saying it is new to them, certainly new to me as well. So just wanting to know if you could clarify what that is and what's the difference between the concession rate and the base rate.

MICHAEL AKINWUMI: [1:13:56]: Thank you. That's a good question.

When I refer to concession rate, I'm actually referring to a discounted rate. You have the base rate or the actual rate, and then due to, you know, for example the pandemic, the landlords may use their discretion to offer discount on that rate. That's what I mean.

Generally consensual rate, I'm referring to the discounted rate.

CASHAUNA HILL: [1:14:23]: Great.

MICHAEL AKINWUMI: [1:14:24]: Does that answer the question?

CASHAUNA HILL: [1:14:25]: Thank you. That's very helpful. Thank you so much, Michael.

Okay. We'll now move on to our final panelist, Sacha Markano-Stark, who will speak to the Fair Housing laws that are applicable to data and tech bias as well

as give us insight into other laws that we should consider as we think about these topics.

SACHA MARKANO-STARK: [1:14:54]: It would help if I unmuted myself.

Thank you for the introduction.

Hey, everyone.

My name is Sacha, I'm an attorney at Relman Colfax, a Civil Rights firm in Washington D.C. and we bring and advise businesses, including financial institutions oner fair lending issues, including issues having to do with bias. There are four areas I plan to talk about today.

The first is an overview of the sources of law that could be applicable to algorithmic and technical discrimination you may come across.

Second is a discussion of types or theories of discrimination, mainly disparate treatment and impact, both of which are relevant in the areas that are being discussed today.

Third, will briefly discuss HUD's disparate impact rule and its current state.

Fourth, I will very briefly describe four cases that have to do with technology and discrimination.

I will do my best to get in my 20 minutes without speaking so quickly that no one can understand me -- quickly.

First, sources of law.

There are six areas of law that I want to briefly highlight when it comes to technological discrimination and housing. Then we'll go into depth on the next slide on three of them.

The first source of law of course is the Fair Housing Act. Technology including algorithms, AI, machine learning, all could in principle raise discrimination issues under the Fair Housing Act and provisions, kind of depending on the technology being used for.

Second, there is the equal credit opportunity act, and to the extent that these type of technologies are used in the credit sector including for housing credit, the provisions there could be implicated for discrimination.

The third area of law to mention is UDAP or unfair or deceptive or abusive acts and practices.

Recently there has been an increase in interest as a means to counter discrimination and to increase access to financial services. On the federal level, it is really a tool for regulatory enforcement, not private lawsuits but on the state level, UDAP laws could vary a lot.

If people are interested in this, these issue as applied to discrimination, a couple of my colleagues have actually published an article in this area and I'm happy to share that as a resource.

The fourth source of laws are state public accommodations laws. The details of these laws also differ from state to state and jurisdiction to jurisdiction.

Some of the laws explicitly apply to websites and Internet companies, some explicitly provide for disparate impact liabilities, and some protect a really broad range of different classes of persons, even over and above the classes of persons that are kind of standard under federal law.

If you are finding yourself litigating technological bias and discrimination issues and you haven't yet done so, it is definitely worth giving your local public accommodations law as look to see what they may contain.

Fifth, there is a growing body of legislation at the state level, pending legislation, a lot of it, that explicitly focuses on issues of algorithmic and other technological bias. It remains to be seen whether and how the legislation is going to be enacted, interpreted, and forced -- enforced, I'll talk through a couple of examples of these kinds of laws.

Finally, this is not a statute, but procedural due process and procedural due process issues can arise when there is the potential for erroneous deprivation of one's private interest by the government.

It is not a concern that is centered explicitly around discrimination necessarily but grounded on the requirement essentially that the government has to be able to understand and explain the algorithms or models that it is using and reaching its decisions, and those types of concerns can certainly be implicated by the types of technology that have been discussed today and can be implicated by issues of bias and discrimination.

For today, I'm going to focus on the FHA, ECOA and the state AI legislation.

To talk about this in more depth in a moment, but both the FHA and ECOA cover disparate treatment claims as well as disparate impact claims. Both of these statutes are key and critical statutes for addressing issues of discrimination and housing and the coverage overlaps. In particular, it overlaps with the FHA prohibitions on discriminations in residential real estate transactions. Although

any of the FHA's substantive antidiscrimination provisions could be implicated as I mentioned a moment ago sale, rental, financing, appraisals, the provision of the brokerage service, all of these areas have the potential to be effected by technology and then in the credit sector in addition to the sort of broad coverage of the Fair Housing Act ECOA can often supplement FHA claims, in particular with additional protected clauses that exist under ECOA but don't necessarily exist under the FHA.

When it comes to technology and AI specific legislation, the current action is really on the state level.

There was a federal bill that was introduced a couple of years back, but that really hasn't seen too much traction.

These state bills take a variety of approaches from relatively narrow in scope, applicable only to government actors or government contracting and that kind of more narrow approach is ex everybody identified by the California bill 13 which imposes requirements on government agencies to inventory what the bill is called automated systems that are being used.

On the other hand, bills are quite broad and potentially applicable to a large range of entities doing business in the given jurisdiction.

Those are more along the lines of for example California AB2269 which has failed in the legislature at this point or by the D.C. stop discrimination by algorithms act. That D.C. bill was introduced this past December and if it is enacted it would prohibit the algorithmic discrimination and "important life opportunities" which the bill is defining to include credit, employment, housing, public accommodations and insurance.

That bill would also require testing and auditing on the basis of a number of protected classes, some of which are quite common to test and audit for historically, that's been broadly done, for example the basis of race. Some really are newer to test and audit and could potentially be difficult to do depending on the availability of data and the type of protected class that's in question.

That D.C. bill in particular is I think worth keeping an eye on in this area. If it is enacted, it has the potential to really change the landscape of the types of laws that we see in this area.

Different states take different approaches, and you may want to check to see if the state that you're in as any bill in this area that's currently pending.

A couple of times I have mentioned disparate impact and treatment.

We will discuss those theories of discrimination and how to proof them up in a little bit more depth.

Starting with disparate treatment, it is the more classical, stereotypical picture that many people have of discrimination and that arises in a case where a defendant either explicitly or implicitly is using protected class characteristics in its decision making.

It doesn't have to be done for reasons of negativity, you can run into disparate treatment issues when people are considering protected class characteristics, also in order to benefit members of that protected class, which favor that protected class members. It is particularly true if you are in a context that's not part of approved, appropriate affirmative action program.

As far as demonstrating disparate treatment, there doesn't need to be a super smoking gun, explicit statement of intent to discriminate or intent to render decisions on the basis of protected class. You really can prove these cases up using circumstantial evidence. For those of you that are familiar with employment discrimination litigation, some of those types of content and types of evidence and types of proof will transfer over into the housing context.

You can look to a history of evidence, stories of discrimination and you can also run comparator analyses where you look at a member of the protected class and look at a member of -- look at someone not a member of the protected class and look at someone not a member of the protected class and compare outcomes for them and see if it is race or a protected characteristic that is driving a difference in what has occurred.

Again, the core element of a disparate treatment claim is intentionally treating members of a protected class differently from other people.

In the technology space, the relevant form or theory of discrimination often but not always will be disparate impact. Disparate impact means that there is a facially neutral policy, practice, technology or model, so something that's not doing disparate treatment, but nevertheless, the outcomes of that policy or practice or model are disproportionately distributed.

This could be an underwriting model that produces disproportionately negative outcomes for black or Hispanic borrowers on the basis of facial neutral criteria, which some examples we walked through earlier today would be of that type.

Or to take a super simplified version of the case which I'll discuss at the end of the presentation. It could be a tenant screening process that disproportionately prevents minority applicants from accessing rental housing. If you have a screening tool, it operated in a way that it is recommended against approval tenants with an arrest record, that might lead to a disparate impact because

people of color are disproportionately likely to have an arrest record.

When it comes to showing an impermissible disparate impact, sort of a discriminatory effect, we do have to do more than just show a mere disparity, and the key element is that the disparity is not justified so the justification of disparate impact is analyzed using a burden shifting framework. In the first step of that analysis, a plaintiff would provide evidence of a disparity and assuming that there is evidence of that disparity, the defendant has to put forward a legitimate business interest for the model, what have you, is serving.

If there is a legitimate business interest, then the third question in the analysis is whether there is a less discriminatory alternative that would nevertheless adequately address or adequately serve the legitimate interest that's been put forward.

It is continuing with the tenant screening example, you know, assuming that there is a legitimate interest that's being served, for example, preventing certain types of crime in a housing complex, could that legitimate interest be served by a screening tool that considers charges but doesn't consider arrests or that considers convictions, not charges, just a subset of convictions as opposed to all convictions, et cetera.

At each cut, where that becomes more narrow, restrictive, they're still likely to be some disparities, but they will probably reduce as that mechanism becomes narrower and they still may serve what is ultimately the legitimate interest of the entity that's using that screening tool.

Specifically, when it comes to AI and models, disparate treatment and disparate impact generally has to do with types of variables that a model is using. If a model takes a protected characteristic or a proxy for a protected characteristic directly as an input to the model then you're in the land of disparate treatment that's intentional discrimination.

When I talk about proxy, what I mean, it is a variable where the predictive power of that variable and the model is really just due to its relationship to protected class. It is not doing anything sort of independently by nature of itself that is helping the model work.

For example, you can think about something like language preference, which is closely related to national origin and other protected classes. If you have for example a credit scoring model that's using Spanish language preference as an input that could really suggest a disparate treatment issue because whatever explanatory value that variable is giving in the model, explanatory power it is giving could very well be just by virtue of the proxy nature with some protected class.

On the disparate impact side, like I was discussing on the last slide, some variables and some types of information are just likely to be disparately distributed across the population. You will probably have a pretty reasonable intuition as to what many of the types of variables would be, criminal background, like we were discussing, hyper local geography that often correlates with racial and ethnic makeup of an area that can be used, it is distributed or again depending on context, something like language preference could all lead to disparate impact coming out of the model.

When it comes to the final step of the disparate impact analysis the search for less discriminatory alternatives, it could be difficult as an outsider to definitively say whether there is a less discriminatory alternative, what it would be, just because you don't have insight necessarily into the models, or what the variables are, but if you're coming up against this issue in litigation, you hopefully have the benefit of some discovery, some expert assistance to help evaluate that question. There are also ways to go about just kind of asking for general information of that type outside of the litigation context as well.

Continuing with disparate impact I want to spend a moment on HUD's disparate impact rule.

The upshot is that in 2013 HUD issued a rule that formalized the prevailing legal consensus that disparate impact claims are recognizable under the Fair Housing Act. In 2015 the Supreme Court also confirmed that fact in a case called inclusive communities.

In 2020, this was somewhat controversial at the time, HUD revised its disparate impact rule in a way that effectively would have made it much harder to bring disparate impact claims under the Fair Housing Act.

That revised rule was more or less immediately stayed by courts so as of now we're still operating under the original 2013 rule.

Last year, 2021 HUD proposed to officially reinstate that 2013 rule and as of now we're still awaiting a formal codification of that the 2013 rule includes a three-step burden shifting analysis that's broadly similar to what we discussed a couple of slides ago. It says that a policy has a discriminatory effect when it has a disparate impact unless the policy is necessary to achieve one or more substantial legitimate non-discriminatory interests and there isn't a less discriminatory alternative.

The 2013 rule doesn't explicitly mention models or algorithms, but it does provide a principal framework that does apply to models and algorithms.

Last thing I'm going to do is very briefly just draw your attention to four cases that are broadly in technology and housing discrimination case.

The first, CFCH versus Core Logic, it is the case that serves as the basis for the tenant screening example that I have been using over the course of this presentation.

That case is headed for trial actually I think in March of this year. As of the last time I looked at the docket. It involves a product called Crem space which looked at tenants on their history, the central part to highlight here, it is the claim that the product has an impermissible disparate impact based on race because of the disparities in criminal justice involvement that exists in the background and that sort of are reflected in the data that really power the product.

There have been arguments in the case about the legitimacy of screening for criminal history of various kinds and the possibility of meeting legitimate interests and less discriminatory ways and it really is a good distillation of a lot of the concepts that are applicable in this area.

I consider it to be sort of the leading case in the area of algorithmic bias and the housing space.

The second case that I'm going to touch on briefly is NFHA versus Facebook. That case settled a few years back and, in that case, NFHA said that the advertising platform allowed advertisers to exclude groups based on protected class in violation of the Fair Housing Act prohibition on discriminatory advertisement.

As part of the settlement for that case, Facebook created an ad portal for housing employment and credit that is separate from other types of ads and that is supposed to limit targeting options in order to avoid the issues.

The third case is a bit of a counter point to NFHAs have Facebook and I apologize for pronouncing this name wrong, it is Opiotennione versus Bozzuto and that was in the district of Maryland in the summer of 2020 and was dismissed in the summer of 2021 and is on appeal in the fourth circuit court of appeals.

The claims in that case were brought under Montgomery County and D.C. law. The court ultimately found that the plaintiff lacked standing to bring the claims because the court wasn't convinced that had the court had been hurt by the restricted ad targeting in part because of the information that was contained in those ads, the court reasoned was otherwise available to the plaintiff.

That's not the only case we have seen that ultimately got dismissed on standing grounds.

Another case that's in that category, it is Vargas versus Facebook in the northern

district of California, and that's in the appeal circuit and for those of you bringing or who are interested in bringing these types of cases, I do think it is worth looking at the opinions and the cases sort of as potential pleading pitfalls and also I know I am going to be really closely tracking the appeal in the Bozzuto case that has potential impact in this area.

The fourth case I have listed up here, the last thing I'll talk about, it is NFHA versus redfin, the claims center on the availability of Redfin's brokerage services, listing services, real estate services and allegations that the services are not available in all areas and that essentially, they're restricted on redlining grounds, redlining borders.

The allegation is that the availability of the services is determined in reference to the minimum of home price policy which effectively makes the services unavailable in communities of color.

The case isn't explicitly about AI or ML, but it is a good example of litigation based on the use of technology and the housing sector and it really wouldn't surprise me if on Redfin's back end there were some sort of algorithmic process that determines which homes do and do not qualify for services under the minimum home price policy.

As of the last time I checked the docket in that case, it looks like that case is pretty close to settling and it will be interesting to see what terms of settlement are made public if any of them are.

Those are four examples of litigation in this space and there really aren't a ton of published court opinions in the area of AI and ML and housing.

It will be an exciting couple of years I think to see where things go.

With that, I will turn it back to Cashauna Hill.

CASHAUNA HILL: [1:39:20]: Thank you so much.

As we move into the Q&A portion of the conversation and as we welcome in addition to all of our panelists Jacy Gaige back into the conversation, I do want to start with a quick question for you, Sacha Markano-Stark, how would someone who doesn't have expertise sort of around these areas, how would someone be able to figure out if a piece of technology or an online service for example imposes a discrimination concern.

SACHA MARKANO-STARK: [1:39:55]: It is a great question, and it can be hard as an outsider to really know what's going on in the background of technology or a website or an app.

You know, there still will be public information or a public interface that exists for people to make use of.

Just yesterday I was asked to kind of poke around and check out sort of a financial services adjacent consumer facing application. I started by Googling it, trying to see if I could get a sense of how the technology worked.

I went to the FAQ page that the app puts out about itself and its own functioning, I think pages like that often go into a surprising level of detail about how the product works.

Then I ended up signing myself up for the app and taking note of the information that it was asking me to give about myself and the information it wasn't asking me to give about myself.

Is it asking me for income, is it asking me for plus four, is it asking me for information that could be used to verify my employment, my education, criminal history, things like that?

Is it sort of more broad, not asking for information as a Civil Rights situated person, sort of immediately may make the hairs on the back of your neck stand up just a little bit?

Then I poked around the advice that it was giving me, after I gave it my information, sort of financial related consumer facing advice and tried to think about the information I had given it, the content of the advice, if the advice that it had given me seemed like it could pose a discouragement risk under ECOA, other types of risks, so that's all by way of saying, you know, there is a lot of public information, a lot of the abilities that you can poke around and your instinct as a Civil Rights advocate.

Often it will lead you to something, you know that's fairly helpful. There are a couple of other specific things that I wanted to point out as well. If you are looking at consumer accessible things, there is mystery shopping, there is testing, you will often see that show up in the background on investigations for cases in this area.

A little aside from that, if you're dealing for example with a tech start-up, you can sometimes find on the open web investor facing materials and those will often go into a little more depth, a little less sort of -- I hate to say dumbed down, but dumbed down for the broad consumer audience, sort of technological workings of the technology, and it can also give your insight into what the business that's offering the product, the value add of that product and the way that it is positioning itself in the market can kind of give you potentially again some instincts on what it may be doing in the back end.

In the credit space, you can look at adverse action notices, those can contain information that's really helpful and really leading in terms of the information that's being considered by someone whose involved in a credit making decision.

You can always try to interact with them. There was a question in the Q&A, like can you just ask people for variables, you can. It may or may not work, you may or may not get a ton of detail about what's going on.

If you are a housing provider, a landlord, you have a friendly relationship of the housing provider or a landlord, those are oftentimes the end consumers for a lot of the types of products.

The folks that are developing them may be motivated to provide some details, to provide some information about the workings of the technology, in part because of the landlords and the housing providers have their own legal obligations and have to do verification to make sure that the third-party services that they're engaging are leading to a discrimination risk.

We certainly have very mixed results when we ask people for the special sauce. You can get some idea about the level of test people are doing, the level of awareness a business has on potential discrimination risks.

I have yet to investigate a case where I didn't start on Google for three, four hours.

CASHAUNA HILL: [1:45:04]: Thank you. That's very helpful.

We'll get into that question that you referenced, and this was a question that came in while you were speaking, Kareem Saleh, and I think any of the panelists are welcome to answer. The question is, as Sacha alluded to, it is asking the lender or the screening company what the inputs, the variables going into the algorithm used to analyze the applicant data, is that something that an individual regular person, consumer can actual ask? Would that be something that the lender, screening company would be available -- would be able to answer.

If this is information available to the general public, where would the person looking for that information best phrase that question?

KAREEM SALEH: [1:45:57]: It is a great question.

As Sacha said a moment ago, you can always certainly ask.

Whether or not the person -- of whom that request is made will oblige, you know, depends.

I think increasingly we ought to be in the practice of at least asking what

variables the algorithms are taking into account and to what extent. I think that's a fine way to ask a question, you know, how are you -- what are the variables that you're using to make decisions about me, how are you using them? It is probably time for a policy debate on whether or not lenders or landlords, other folks who are using the algorithms to make high stake decisions on people's lives shouldn't be obliged to share at least their variable lists. I would be interested to hear if Michael, Sacha has anything to add to that.

MICHAEL AKINWUMI: [1:47:00]: What you said, Kareem Saleh, I would like to clarify that in -- just like we mentioned in the presentation, so usually the variables that actually go into the machine learning model may not be the variables that I'm actually making the decision, because the variables, there is a lot of complex interactions that actually go on at the core of the model itself. Why is it a good practice for consumers to request for the variables, and there are clear cases of disparate treatment where we know that protected classes, well protected classes are not allowed to go into the model or even some exist for the protected classes. We should probably pay attention to the fact that there may be some complex interactions which ultimately leads to what the data scientists call feature engineering that are actually being used to make the final decision. That's just the comment I would like to add.

JACY GAIGE: [1:48:13]: One thing I would add to, to both of those, sometimes there is more information out there about outcomes than you would think and you can use that information to leverage the conversations about what's driving the outcomes, because it may make the respondent interested in solving the problem themselves.

As Sasha was saying, a lot of the companies speak very differently to one side of the users than the other, the housing provider, the advertiser, whoever they're selling their model to, they may speak really candidly about what it is optimizing for.

They also can provide those entities, a lot of data about the results. This is big data. There is a lot of information there.

So those companies may be getting, you know, more information about who is being screened out works is screened in, who has seen ads that kind of thing. We can take a look at the outcomes and start to talk about what's driving those and how they can fix it.

CASHAUNA HILL: [1:49:14]: Thank you. That's very helpful.

The next question is, given that algorithms have targets, goals, objectives, is it possible then for one to assume that algorithms can be built or have been built with not just accidental or coincidental biases but that the systems have been built with -- intentionally built with biases built in? Does that question make

sense? I can ask again.

Given that algorithms have target goals and objectives, can one assume that algorithms can be built not just with accidental biases baked in, but that those biases can be intentionally discriminatory.

KAREEM SALEH: [1:50:11]: Yes. Certainly, one can build a machine that has, ma'am intent. I generally don't find that the folks -- for the most part, the most of the individuals I encounter with bias models tend not to be people of bad faith who have constructed them that way.

It hastened to be more people using new data, new techniques that are perhaps not well studied or not well validated resulting in unintentional bias.

Certainly one could make a terminator type machine that was built at the initial to be unfair.

CASHAUNA HILL: [1:50:51]: Thank you.

Do you have something to add, Jacy Gaige?

JACY GAIGE: [1:50:55]: No. I would just say that you can -- yeah, you can optimize algorithms and machine learning for what you want to. That can verge on intent again when talking about things like scams or profit margins where you look where you can exploit vulnerabilities or weaknesses of certain consumers. Like humans, it is complicated, but we should not rule out intent in all of its forms.

CASHAUNA HILL: [1:51:24]: Thank you. That's very helpful.

A question, I'm so sorry about the siren going directly past my house.

There is a question about what individual Fair Housing agencies can do. The question is, the information is fascinating, but what tools do individual Fair Housing agencies have at their disposal in this realm? For example, how can Fair Housing agencies create investigations to uncover AI discrimination and how would these groups go about implementing testing for AI and ML bias.

SACHA MARKANO-STARK: [1:52:11]: In terms of investigation, hopefully the first question from Cashauna that I answered is a bit helpful. I think one of the things -- you know, the referrals that we get, the cases that we see come to fruition are often based on one person's experience at the outset. If you as a Fair Housing agency are hearing from someone in your community that something has gone sideways or feels wrong or something like that, you can then sort of employ some of the public resources to try to figure out if they are there.

The one thing I didn't mention in the answer to the previous question was the

online reviews, a lot of different types of technologies and businesses, and those can be really fruitful and really leading.

Again, that's just open web, Google research that you can do to sort of try to figure out how things are feeling in your community and whether there is a potential issue that's fomenting with a particular type of technology.

From there, my theory of this, it is really not so different from investigating any other type of other case. I think that as lawyers, in particular, they can be stereo -- they can think too much of technology, how does it work, currently trying to teach myself blockchain and it is terrible.

What's going on, you're trying to figure out if there is a problem and the same techniques that you employ for any case, is the same I employ when it comes to algorithms and there you have to lay over the overlay of what Kareem Saleh and Michael Akinwumi have talked about, and potentially using the more defused forms of information that come out of big data that Jacy Gaige was talking about in order to build up the case.

CASHAUNA HILL: [1:54:25]: Thank you.

I was going to move into the last question if that's okay? Did you want to add something?

KAREEM SALEH: [1:54:31]: Just one very quick point. it is that those working in-housing have the benefit of the home mortgage disclosure and database and lenders are obliged to submit certain loan level information about their mortgage lending which includes these to the government every year and that data is all publicly available, individual Fair Housing groups I would imagine could go to that website and pull that data and do what we did in those maps and kind of construct adverse impact ratios for every mortgage originator in the country and we have done that work, if anyone wants to talk to us about it, feel free to reach out.

MICHAEL AKINWUMI: [1:55:14]: I was going to talk about the initiative offer, that's an initiative that we hope that there are educational materials that people in the Fair Housing, they'll actually understand the issues that they want to address or identify when it comes to the discrimination of bias.

More importantly, I think similar to what the leadership is actually showing, actually this is probably a good time to also embrace the technology, some of this, when it comes to using this, for example, we're limited in terms of the patterns you can see when it comes to testing some of the websites and the program review, you can actually develop or even write your own code to do the testing and to see if there is any discriminatory pattern That's what I would like to add.

CASHAUNA HILL: [1:56:12]: Thank you, Michael.

Thank you to all of the panelists.

This has been an absolutely enlightening conversation today. I certainly have lots of notes that I will also have to review again because I know very little about what we discussed today.

I'm thankful to you all for sharing the time and insight with us and judging from the amount of questions, I know that our attendees also found this conversation very enlightening.

I'm sorry that we can't get to more questions today, but we do have to wrap up.

I would like to thank everyone who attended for your participation in today's forum. We hope that you will join us for the next conversation.

Please check out the NFHTA website for a description and important information on registering for upcoming forums. Please also connect with the National Fair Housing Training Academy on LinkedIn for insights and information about upcoming events, including future forums and courses.

Thanks to everyone who made today's event possible, including our panelists and our interpreters. Thank you very much.

Finally, everyone should please be on the lookout for a survey that will pop-up when this training ends. The survey will allow you to provide feedback on today's conversation. The feedback is critical to improving these events and you should not take very long to complete.

Thank you again. We look forward to seeing you on the next NFHTA forum.

Take care.