Transcript

22 February 2025, 07:00pm

Interviewer 0:10
OK, fair enough. So did you go over the package I sent you, including the the consent form.

Stakeholder2_AI Developer 0:15
You know what? I am so sorry. I I forgot. I was too busy the last couple of days. But I yeah, I give any consent you need to or this interview. Yes.

Interviewer 0:15
OK.
OK, alright. You can sign it and send it back to me. You know, at your convenience, as long as you consent to being recorded. And I'm gonna use your transcript for the analysis purposes and your name will be left out of everything. I assure you. OK.

Stakeholder2_AI Developer 0:39
Right.
Yeah.
Interviewer 0:52
Only your designation of whatever stakeholder you are, such as driver, passenger, et cetera engineer you know whatever is appropriate.
OK, So what we're going to do is first of all, thank you for participating. And also I'm very curious to hear about your Waymo experience as a passenger.
So what we'll do is we'll go through. I'm gonna give you a case study scenario and then based on that case study, I'm gonna ask you a few questions and remember that there are no right or wrong answers because this isn't a test for you. This is about getting your opinions your viewpoints, your perspective and so how you live in this world and how you're gonna be impacted by this technology is very important for this research. OK, so.

Interviewer 1:54
OK. The case study scenario that we're going to be talking about involves an automated vehicle, an AV named, Autopilot. OK, so it involves a real life case within the AI application of automated vehicles, or AVs. It involves the occurrences of actual car crashes involving one particular AV brand, which is Tesla and it's Advanced Driver Assistance System (ADAS) called Autopilot. So, Tesla's Autopilot system controls the steering, the breaking and acceleration functions of the AV without any assistance from the human driver. Furthermore, note that Autopilot could at any time disengage and hand over the controls to the human driver. Now, according to US NHTSA's office of Defects Investigation, EA22-002, between January 2018 and January 2022, Tesla AVs, with Autopilot engaged, were involved in 16 crashes where they struck highly visible stationary in road or roadside first responder vehicles.
So, for example, police vehicles, ambulance, fire trucks, road maintenance, everybody with their you know lights flashing, et cetera..

Stakeholder2_AI Developer 3:28
Mm hmm.
Interviewer 3:28
That and these vehicles were attending to preexisting collisions. OK, so furthermore, on average in these crashes, Autopilot aborted vehicle control less than one second prior to the first impact, so less than one second for the human driver to take over. OK.
Stakeholder2_AI Developer 3:51
Mm hmm.
Interviewer 3:51
So would you. Would you like a text? I can put it in the chat of this scenario. Are you OK?

Stakeholder2_AI Developer 3:57
I I think I'm following, I'm good.
nnati Patel 3:59
OK. The other thing, if you want, there are photographs and reports publicly available

reports and news organisations about these crashes. Would you like to see those? I can again put the link on the chat if you wish.

Stakeholder2_AI Developer 4:17
OK.
Interviewer 4:18
Yes. OK. Let me just.
Stakeholder2_AI Developer 4:20
Yeah, but whatever works for you. Like maybe sharing the screen.
Interviewer 4:22
OK, so if if you want a visual that that's one of them, OK, you.
Stakeholder2_AI Developer 4:27
I think yeah, I followed you, yeah. Very good.

Interviewer 4:30
OK, alright. So, my first question for you, OK.

Interviewer 4:38
What kind of explanatory information do you want from the AI named autopilot? OK, about these crashes so. What exactly, you know, based on this scenario described above, Autopilot, the system that controls steering, braking and acceleration functions of the AV . . .So, what specific questions would you ask of Autopilot? What types of information would you seek from Autopilot about the decisions made and actions taken? OK, that's so the AV is performing the human driving tests of motion control. So, what do you want to know about the scenario from Autopilot so you can answer this question by finishing the sentence such as I would ask X, Y, or Z or I would want to know X, Y or Z or I want information about X, Y or Z.
Stakeholder2_AI Developer 5:41
Yeah, I think the first thing I may ask is and all the things that's working and they should like these things. You use either cameras or LIDARs. So, is the system getting the right information to process and make a decision? That's my first question.
Interviewer 6:01
OK, OK. So for the purposes of this particular scenario, please assume that all other systems of the car are functioning correctly. So, assume that all the hardware,

steering engine wheel, pedestals, cameras, lidar sensors, all of you know the software such as . . . engine optimization, fuel efficiency. Just assume it's all working fine.

Stakeholder2_AI Developer 6:12
Mm hmm.
Yeah.
OK, it's all algorithm OK.
Interviewer 6:29
Yeah. So what do you want to know about the decisions made and the actions taken by Autopilot?
Stakeholder2_AI Developer 6:47
Maybe I ... I . . . I . . . I would like to know, before the collision what was in the output of the model like it gets you know it's inputs in real time and then it makes decision that's the output right? So, I want to know.

Interviewer 7:08
Mm hmm.

Stakeholder2_AI Developer 7:10
What exactly the input was at that point in time?
Interviewer 7:14
OK.

Stakeholder2_AI Developer 7:14
Maybe you know the history and passing, but the model is saying something else, so that's like a failure to detect.
Interviewer 7:22
So what kind of input? Yeah. Can you describe the functions or features of this inputs that you're looking for?

Stakeholder2_AI Developer 7:30
Yeah. So, for . . . for example, you have, you know, a picture of a pedestrian passing by.

Interviewer 7:37
Mm hmm.
Stakeholder2_AI Developer 7:37
That thing is input to the model and the model should precisely identify that's a pedestrian and then make a decision accordingly. So I want to know . . . from the input, what was the input for the model and then what was the decision based on that input? You know, sometimes you have the right inputs and you get the wrong decision because of, you know failure to detect. So I want to know all sides.
Interviewer 7:57
OK so. You want to know what inputs OK were into the algorithm? How far in advance do you want to know? You said just up to the accident. Or can you tell me time frame wise? Couple of seconds, a minute 8 minutes?
Stakeholder2_AI Developer 8:21
Yeah.
So those things I . . . I believe are fixed. But maybe right before the collision, . . . that's where . . . where I want to know what was going on.

Interviewer 8:36
OK, OK. So OK, that's the inputs. Beyond the inputs, what else do you want to know from Autopilot?

Stakeholder2_AI Developer 8:48
What was the prediction of the model? What was the detection like? You know we you see you give it a person's image and it detects something else, maybe some, you know, shadow of the person or something like that. So yeah, that's basically failure. And that's in the 1st place. Why these accidents happen, right? Failure to detect right things.
Interviewer 9:03
Huh.
Uh huh.
OK, so you're assuming, are you? OK? So are you assuming that it's a failure of input and or failure to detect? Therefore, are you telling me you have confidence in the decisions made and the actions taken by Autopilot, assuming that it was a failure of detection?

Stakeholder2_AI Developer 9:35
No, I . .. I do not have confidence in the detection. But I wan to know if I was wrong. Like I just proved me wrong. You . . . you saw something and you make the right decision or you . . .you. So you didn't say anything and you'll just keep driving, right?

Interviewer 9:57
Mm hmm.
Mm hmm.

Stakeholder2_AI Developer 10:05
So I can give you an analogy. Forget the Autopilot. The countervision model I give it the . . . the picture of a fish. And it should say that's a fish. So that's like the perfect world, right? But the models most of the time or, you know, often times they miss to say that's fish. That's something else. Like, there's nothing. I don't see anything so. That's where my question would be like, what did you see given you know . . . X?
nnati Patel 10:35
Mm hmm.
Mm hmm.
Stakeholder2_AI Developer 10:38
So in other words, I . . . I question the accuracy of the model.
Interviewer 10:45
OK, so you're questioning what it saw, whether it saw . . .the scenario in front of it. um... Let's assume it did see the scenario in front of it. What other questions would you have for Autopilot?
Stakeholder2_AI Developer 11:03
So that then that would be like the decision process. So I will ask, you know you . . . you have the right information, you made the right detection, but still the crash happened. So I want to know the algorithm after the model after the fact like how are they managing, you know situations like this, it could be that action time . . . or, you know, many other things that prevent the accident. So I would ask, how is the algorithm, you know, making decisions after, you know, the model does its job in detecting the right thing.

Interviewer 11:46
OK. What can you elaborate more on the decisions like? If this was a human being doing the driving, what kind of questions would you ask it?

Stakeholder2_AI Developer 12:03
Yeah. You know, our reaction time is different across age groups like, you know, the older we get, our reaction time goes, you know, longer. So for . . . for a human driver, I may ask, was the driver like young driver who is distracted or, you know, in their 80s or you know, about 80 and they were slow to react.
Interviewer 12:16
Mm hmm.

Stakeholder2_AI Developer 12:31
So that's a question you . . you would ask human driver. Before the Autopilot, it's a system, right? The model is just one component and there is the other on top of. It's like a bunch of algorithms and you know deterministic algorithms and how are they designed . . . to act after getting the detection after the . . . the model says hey there is pedestrian, you know, what's the decision after that. So that I want to ask . . . I want to know how that algorithm is designed.
Interviewer 13:06
OK, so let's talk about the algorithm. You said you're making an assumption that the algorithm is deterministic. What if the algorithm is not deterministic? What if the algorithm is, say, a neural net or a reinforcement learning, and you don't have visibility into . . . like it's not a transparent algorithm as deterministic where you can go line by line, right? So what questions do you have for an algorithm that isn't deterministic but is making these decisions?
Stakeholder2_AI Developer 13:41
um... So this is how I see copilots or any kind of AV. There is stochastic components that the model does the job of detecting.

Interviewer 13:54
Mm hmm.

Stakeholder2_AI Developer 13:54
But that's, that's where it ends. And once you get information, it's easy to build a

deterministic model to make a decision, right? So I can like, you know, write a code. That says OK when the some X feature appears, do something the action right? So that's kind of deterministic. Maybe you know I'm not 100% correct, but that's how I see it. So that they I think most of the time the problems happen in the stochastic component or the model of the Neural Nets or those, you know black boxes.

Interviewer 14:23

Mm hmm.ight.

Stakeholder2_AI Developer 14:33

Yeah, but the decision to act should be more of, you know, deterministic.

Interviewer 14:40

OK, couple of things. The reason I'm asking this and explainable AI is about breaking into these algorithms that are not deterministic, including the algorithms that are making the decisions. What you think are deterministic decisions are actually made by stochastic models and so.

Stakeholder2_AI Developer 14:57

Mm hmm.

Interviewer 15:04

And so, this project is about . . . is within the realm of explainable AI, which started about 10 years ago. Where they recognised that a lot of these AI algorithms are not deterministic and therefore you cannot step line by line through the logic of the decisions made so, they recognise, OK, so DARPA put together an initiative to break into these and that's the name explainable AI. And within the realm of explainable AI. To make this research and further this research into breaking into these non-transparent algorithms. Uh … What I'm saying is currently, . . . A lot of explainable AI efforts are AI engineers who are building explainable AI models for themselves rather than for the majority of people impacted by these technologies. So I'm saying explanation is not one thing, OK, there are different explanations depending on context, different people, different situations and even for the same situation, different people will want different explanations. So the explanations you seek would be different than an AI engineer or a policeman or a lawyer. Or, you know, judges, etcetera. So part of my research is to ask, OK, what kind of explanatory information would you need and for explainable AI engineers to go break into this black box and say, hey! This is the type of information we need so I understand you're assuming

that these algorithms, when it comes to making the decisions, are no longer stochastic, they're deterministic. What I'm saying is, within the realm of AVs there is a slew of. . .

Stakeholder2_AI Developer 16:53
Mm hmm.

Interviewer 17:02
Non transparent algorithms. And . . . including the ones that are doing motion control, motion planning, making the decisions because cars on the road and human beings they are not predetermined. . the . . . what cars encounter, and it's not just a perception thing, it's a situation. Even if you have the same traffic intersection. You will have a different scenario to deal with and cope with when you approach that intersection. You know what I'm saying. So therefore the amount of information that cars and human drivers in fact are dealing with are unpredictable. So that's where these Neural Nets and reinforcement learning, et cetera. All these algorithms, which are not transparent are utilised. So you're asking questions as if these were deterministic, and I'm asking you to think.

Stakeholder2_AI Developer 18:04
It's just . . just the last part, the last bit, right. So you know, if you look at the whole thing, it is still stochastic, right? Because you are you are working with a stochastically generated data to make a decision where the decision maker is still deterministic. It's like, you know you, you listen to me, what I'm saying. And you make a decision. But I may say, you know, several things like some something that's not real.
Interviewer 18:33
Yeah. Yes. Right, right.
Stakeholder2_AI Developer 18:34
That's where I'm seeing, but then the final thing remains still unexplainable or not explainable or stochastic.

Interviewer 18:42
Right. Right. So OK, so. Um… Tell me about your Waymo. You mentioned that you were a passenger in a Waymo. So can you share with me your experience? Because that was very intriguing when you responded and signed up for this study.

Stakeholder2_AI Developer 18:56
Yes. Yeah, yeah. I I even have a a few phone records that I may share if you want. So yeah, it was very exciting experience.
Interviewer 19:12
Yes, please.
Stakeholder2_AI Developer 19:17
My first impression is. . . The system is, you know, as good as a human driver for me from, you know, the things I saw like that decisions it was making, you know, at the traffic light, you know, people sometimes you know or not everyone stops. Like it makes a perfect stop. But the traffic lights. But this was this thing was doing it like 100% perfect and I . . . we were in you know in San Francisco. Very crowded traffic. So I was very impressed. And yeah, so my first impression is this thing is working. You know, there may be accidents, humans also make mistakes. So that's my first impression.
Interviewer 20:04
Mm hmm.
Mm hmm mm hmm.
Stakeholder2_AI Developer 20:12
And if you have a specific questions, I can go further.
Interviewer 20:15
Well, yeah. Tell me. What? Yeah. Tell me what else you.
Stakeholder2_AI Developer 20:21
Yeah. OK.

Interviewer 20:22
Beyond your first. Yeah. Beyond your first impression. Tell me what else you your viewpoint is on your rides that you took with.

Stakeholder2_AI Developer 20:23
Yeah. Right, right. Yeah. So I I can speak about the . . . the interaction between the. Autopilot and the passengers like us and my me and my friends. So the first thing the the system said was when we get inside, um… It. . . it told us or she told us, you know um… About safety issues like bacteria, seats and nobody is recording you or we're not listening to you unless you want to press a button.

Interviewer 21:05
So was there a human being in the driver's seat or no.

Stakeholder2_AI Developer 21:09
No, no, no, just here this just a car stuffed on the side way and we we go inside and we will work on the system and you know the basic safety information, that's what we heard. And then about, you know, yeah, don't worry about anything. If something bad happens, stay where you are.
nnati Patel 21:17
OK.
Uh huh.

Stakeholder2_AI Developer 21:35
Like your seatbelts there. And so it . . . it made me comfortable, like listening to that kind of information, including the privacy issues.
Interviewer 21:43
Mm hmm.
Stakeholder2_AI Developer 21:48
And then it played AI generated music for us. You know, some of my friends didn't like it. Like it's weird music. So those are the things I remember, I can review the video and add more.
Interviewer 22:04
Yeah, please go on.

Stakeholder2_AI Developer 22:04
So positive, yeah. What if? Yeah, we, you know, I am a data scientist and my friends are too. AI trained professionals, so we were discussing more of like what's going on in the model. For example, I asked whether the system works elsewhere than where it was trained, like around downtown San Francisco area. For the one of my friends.

Interviewer 22:38
What was the answer? What was the answer?
Stakeholder2_AI Developer 22:40
Yeah, the the answer. One of my friends told they not necessarily like there may be issues if they, you know … take this system outside of San Francisco, downtown,

maybe you know more accidents may happen. And my other friend. Um... Said I remember it's going to take like 10 years for this time to be for this system to be adopted like in Victoria, BC. You know, there's also the like social component, right? So yeah, we'd also discussed that.

Interviewer 23:16
Well, hang on so. One question. You were driving ... uh.. you were inside the Waymo and you're considering San Francisco as the test bed for this?

Stakeholder2_AI Developer 23:32
No, no, no. We . . . we were in San Francisco, right?

Interviewer 23:32
mm... hmm... So you're you're you're assuming that it's still in testing phase in the city of San Francisco and not applicable outside of San Francisco, is that correct?

Stakeholder2_AI Developer 23:36
Not . . .not assuming it is in the test phase, but we were simply asking if this system works elsewhere where you know new environments, you know, new buildings. So you know, I assume it memorised somehow environment the buildings and the turnings and everything. But will this work elsewhere like for, for example Toronto?

Interviewer 24:16
Mm hmm.
Stakeholder2_AI Developer 24:17
Was a question, but we're not asking whether it was in test phase or so. Obviously it's in production phase, right, it's.

Interviewer 24:23
What was the... So the answer is it ... it has memorised. You're assuming it memorised San Francisco and its environment --- buildings, cross sections, et cetera, but you don't assume that it would be the same for Victoria or Toronto or other cities.
Stakeholder2_AI Developer 24:36
Yeah. Yes, yes. Just like you know any other machine learning models, right? It's

always hard to to . . . to make the right decision at different datasets, different distribution.

Interviewer 24:46
Yes. Mm… hmm..

Stakeholder2_AI Developer 24:54
Yes.

Interviewer 24:55
OK. So. Um… Do you have any questions for the designers? I'm gonna jump back to the Autopilot… Tesla Autopilot scenario we talked about. Do you have any questions about the algorithm itself for the designers?

Stakeholder2_AI Developer 25:18
um… Not for like the functional components, but just say you know if some they could add more comfort to the passengers like comfort applications like you know, maybe add realistic music so people get more comfortable, right. Listening to real person's voice.

Interviewer 25:39
Mm hmm.
Mm hmm.

Stakeholder2_AI Developer 25:42
Will create environment where you think you are in. You know in a very safe place.
Interviewer 25:46
OK. So more or so more human friendly environment inside. . . inside Waymo.

Stakeholder2_AI Developer 25:50
Yeah. Exactly. Exactly.

Interviewer 25:54
OK. So can we now jump back to Tesla Autopilot and the scenario we were talking about?

Stakeholder2_AI Developer 26:02
Mm hmm.

Interviewer 26:02
OK, so do you have any questions about Tesla's algorithm? Based on that scenario and those 16 accidents we're talking about.

Stakeholder2_AI Developer 26:10
Oh. Yeah, obviously something is not going right. Earlier they were very confident, overconfident. They were, you know, they. . they called it Autopilot and then after a while, you know, Co pilot, those are like the shifting mindset, which is.
Interviewer 26:33
Yep.
Stakeholder2_AI Developer 26:38
What it should be?

Interviewer 26:38
It is. It is still called Autopilot and it is called. They've also used the word full self driving FSD. They have not used the word copilot. This is Tesla I'm talking about.

Stakeholder2_AI Developer 26:47
OK, cool. All right, right. So but still, there's a shift in, you know that tone, right? So because of, you know, some legal issues.

Interviewer 26:56
Yes, yes, the. . . the metaphors. Yeah. The metaphors make people think certain things.
Stakeholder2_AI Developer 27:03
Right, right. So yeah.
Interviewer 27:04
Right. OK. So what? What algorithm? OK, your question about the algorithm is.
Stakeholder2_AI Developer 27:15
It is not working. And the question is why right?

Interviewer 27:23
Can you elaborate on the why? What?

ZE **Stakeholder2_AI Developer** 27:24
Yeah. So each accident is different. So it's . . it's kind of hard to ask a general question, but you go to a specific case and ask the right questions. It could be the detection issue.

UP **Interviewer** 27:34
Mm hmm.

ZE **Stakeholder2_AI Developer** 27:44
Or the reaction time is too slow too . . . too high.

UP **Interviewer** 27:49
Mm hmm.

ZE **Stakeholder2_AI Developer** 27:50
So things like that. So I . . . I don't see a general question at least I need to review more before I make, you know, ask the right question.

UP **Interviewer** 27:57
Right, OK. How would you go about improving this algorithm based on the . . . remember the pattern I'm talking about happens to be the first responder car crashes, right? That's the pattern that's put in front of you. So I'm not talking about the other crashes that Autopilot has been involved in and all the other in rest . . . I'm just talking about this. The first responder crashes.

ZE **Stakeholder2_AI Developer** 28:10
Mm hmm.
Interviewer 28:24
I could then just happen one or two times. It happened 16 times, right? So.

Stakeholder2_AI Developer 28:24
Yeah, that's interesting. That's weird, yeah.

Interviewer 28:30
And and and it happened with very highly visible scenes, you know with flashing lights and lots of people around it. You know, so.

Stakeholder2_AI Developer 28:43
Yeah.

Interviewer 28:43
Like these are very highly visible objects and large like fire trucks, ambulances are large objects, they're not small, right?

Stakeholder2_AI Developer 28:52
Right. So, you know, at the end of the day the . . . the computers see numbers, right? They don't say, you know that there's a high visible flashy car in front.

Interviewer 29:05
Right.

Stakeholder2_AI Developer 29:05
So I'm not 100% sure why they preferentially, you know, crash to these first responders, so. Maybe they need more of those examples in the training phase. More of what's what happened.

Interviewer 29:23
Mm hmm.

Stakeholder2_AI Developer 29:27
So they do this thing called online learning, so they . . they basically.

Interviewer 29:27
OK.

Stakeholder2_AI Developer 29:32
Isolate the samples where the model failed to detect or struggle, and then they retrain with that sample. So maybe that's something they can.

Interviewer 29:42
Mm hmm.

Stakeholder2_AI Developer 29:45
Try to improve.

Interviewer 29:47
Right. So more training, more training, more focus on certain objects, yeah.
Stakeholder2_AI Developer 29:52
Yeah, yeah.

Interviewer 29:53
OK, also what about functions or methodology about improving the algorithm? What else?
Stakeholder2_AI Developer 30:06
Hmm.
Interviewer 30:09
Anything else?
Stakeholder2_AI Developer 30:11
Yeah, I am trying hard to like to see where they can improve?
Interviewer 30:16
You're providing you're providing great, great input. OK, so keep thinking out loud because all of your opinions, all of your viewpoints, your perspective, all very important for this study. You know, so keep going. Keep going.
Stakeholder2_AI Developer 30:32
OK.

Interviewer 30:37
And if you can't think of anything, just say I can't think of anything else. That's fine too.
Stakeholder2_AI Developer 30:37
Yeah. Yeah, right now that's . . that's my answer, yeah.

Interviewer 30:44
OK, OK. So OK, one more question. Well, with respect to this scenario, there's some cursory questions afterwards, but so do you have confidence in the steering, braking, acceleration functions of Autopilot? When it comes to these 16 scenarios, how . . . how confident are you about its ability to steer, brake, accelerate, make decisions about steering, breaking, accelerating and take actions about steering, breaking and acceleration?

Stakeholder2_AI Developer 31:21
So based on my limited experience, I have a very good confidence in the system, . . but that's because I was not involved in any accident or any in a near accident. Right? So . . . so you know, perfect world, with my experience, it was doing a perfect job the way it was, you know.

Interviewer 31:24
Mm hmm.
OK.
Yeah.
Mm hmm.

Stakeholder2_AI Developer 31:44
Coming to stop a traffic light was … Well, yeah, yeah. Well for.

Interviewer 31:46
They're talking about way more now. You're not talking about Tesla, right? Yeah. OK.

Stakeholder2_AI Developer 31:51
Just like I . . . I cannot speak from experience, right? Just based on the information he gave me, it sounds like maybe it's not the braking system. Maybe it is. Maybe it's just

a lack of, you know, for examples. Do the training of those, you know, objects, those, you know, flashy objects because they we don't see them often. . . in our streets, right, those first responders, maybe they, they don't have enough samples.

Interviewer 32:18

Mm hmm. Mm hmm mm hmm.

OK.

Stakeholder2_AI Developer 32:24

Yes.

End Transcription for analysis .... general discussion continued until 48:56 when Interviewer stopped recording and transcription