



# ON THE PLATFORM EPISODE: RESPONSIBLE AI AND THE CONTENT CHALLENGE

## VIDEO TRANSCRIPT

**Host: Martin Stoddart,  
Principal Director,  
Webscale Services,  
Accenture**

**Guest: Rumman  
Chowdhury, Responsible  
Artificial Intelligence  
Lead, Accenture**

**Martin:** Hello, and welcome to On the Platform. My name is Martin Stoddart, and today we're going to explore some of the ethical challenges and opportunities that companies face when introducing artificial intelligence into their content operations. I'm joined by Rumman Chowdhury, who is the responsible AI lead here at Accenture. Rumman, welcome to the program.

**Rumman:** Thank you for having me.

**Martin:** You're very welcome. Let's just start with a bit of context. So software and platform companies rely upon content as a key means for engaging their users, and content operations doing things like content moderation are essential, right, in terms of reviewing the content and labeling the content correctly before it appears on the platform.

And traditionally these operations have relied on humans to do the work, though increasingly AI is being introduced to automate away certain tasks or to guide the humans towards better judgments. For example, an increasingly broad array of objects or entities can be tagged automatically in user generated images. So AI can have a significant positive impact in terms of productivity for these operations, but implementation also comes with some challenges. For example, companies must take care to avoid human biases and prejudices creeping into the decisions that the algorithm are being trained to make.

So Rumman, let's, if we don't mind, let's just start off with a pretty simple question. When we talk about responsible AI, what does it mean, and why is it so important?

**Rumman:** So responsible AI is about creating a practice and a process of understanding what the consequences of the technology we are building will be in society and also, of course, on our consumers. Why is it important? Well, we're in an era right now where there's a lot of fear and misunderstanding about artificial intelligence. And the goal of this practice is not just to demystify, but also to provide solutions for some of the questions that we're all grappling with around the world right now.

**Martin:** And when you think about some of the



challenges, right, that companies are facing, what would you say are the main challenges in terms of successfully implementing this?

**Rumman:** At the moment I see a lot of companies that have AI proof of concepts, but those proof of concepts don't scale. And when we ask ourselves why, one of the answers they give is that, well, they're just worried that they're going to be violating some law or that it would have an unintended negative consequence. And that's pretty telling. It's quite difficult to take a POC and make it a full underprice product. And given the uncertainty of AI, you put that in the mix and people are quite worried about moving forward.

So in other words, for Accenture, to build AI and have a good business around AI, we have to address these problems as our clients face them.

**Martin:** And do you think the uncertainty that you're identifying there, is that just uncertainty on behalf of the companies themselves, or is it, are there multiple stakeholders where this uncertainty sort of plays a part?

**Rumman:** That's a really great question. And there are multiple people who are involved in all the conversations that I have. Yes, there are the technical folks who are interested in the explainable, understandable AI, as well as the new fairness tool that we've launched. But I get a lot of questions from, for example, from chief privacy officers, from chief strategy officers to say how can I understand this technology and how it will impact how people perceive my company to be, or how do I understand this technology, for example, in the privacy officer perspective, and make sure that we are still in adherence and compliance with regulations like GDPR. So there's quite a few people involved in these conversations. That's very reflected in the kinds of products that we're building.

**Martin:** So there's lots of people involved. There's obviously some trepidation, maybe, about the future. In terms of ensuring that future, then, you know, it seems that policy is a pretty big, pretty important factor there. How do you ensure that policy is well written for AI, and how do you do it in a way that it doesn't become unwieldy for moving the project forward?

**Rumman:** Absolutely. So there's two ways to think about policy, right? So there's the policy that states, governments, etc. will—and governing bodies will have, so political bodies. Then there's the kind of governance or policy that you pass within your organization. We have more organizations adopting things like ethical review boards. Accenture, we're building one ourselves, an ethical review board for artificial intelligence, which is different from the other types of review boards for technology that we might have. And of course we always run into the statement "regulation stifles innovation".

It is also worth appreciating that we are in a bit of an AI race right now with other countries around the world, whether it's China versus U.S. versus, you know, U.K. or France, we have everybody throwing their hat into the ring. And what the tension is right now is yes, we want safety, privacy, security, etc., but nobody wants to lose this battle. So our proposal is to develop a system of agile governance which adopts the agile methodologies that most technologists are quite familiar with and put them into the traditional waterfall methodologies that governance is created by.

**Martin:** You mentioned the race going on, and obviously there's different countries involved. The spread of use cases that are involved are quite significant as well. Essentially, does the scenario that you're implementing it in affect the level of risk or the pace at which you want to experiment along that agile process? I'm thinking, for example, you know, if you're moderating away, or if you're using content for autonomous vehicles, you know, it can literally be life or death circumstances of the content that's being reviewed. How does that impact the pace at which these projects can proceed.

**Rumman:** I think you're raising a really excellent example of the kind of tensions that can happen. So thinking about self-driving vehicles, and it's probably the best example of both the fear and trepidation, the potential for regulatory oversight, that in this desire to make sure regulation is a cycle innovation on one part, thinking of it at a user level, we pride ourselves on making frictionless experiences.



Frankly, the reason why we love engaging with a lot of technological products is because they've made the experience seamless and frictionless. So let's think about a self-driving vehicle. Number one, the question is what is it that we're so scared of? Now, there are probably—not there are probably—there are definitely more deaths by traditional vehicles on the road driven by human beings than there have been from the autonomous vehicles that are out there. We've had one death, right?

And yet that death just had so much scrutiny. So really there is—I think what we're tapping into here is this latent fear of what we were talking about in the beginning of robots taking over, or robots taking our jobs, and this idea of losing human agency. So while that is to say that, you know, of course we need good methods of identifying people, making sure the car stops and fail safes, at the same time I wonder what is at the root of what we are so afraid of when this technology is actually so far better than the average person.

The second thought here is when we do create the system of, you know, of flagging, or stopping, or raising awareness of, you know, some potential issue, with the sort of famous case at this point of Uber, really what it turned out to be was that the algorithm was tuned so that it didn't register the woman as a person.

So in other words, you set a threshold—it's picking up things all the time. So you think about if you crank it completely in one direction where it's overly safe, it'll stop every two feet and you'll just basically have to drive the car yourself, right? So we want to avoid that. On the other end we get something like the automobile just runs over everything in its way, including people—

**Martin:** Right.

**Rumman:** —animals, and—

**Martin:** Definitely don't want that.

**Rumman:** No, we don't want that. So finding that middle ground that's, there is no science to it yet. You know, it is algorithmic value, but it's not optimized. We don't know what that is. And what happened in the case of Uber

was that the car flagged—it saw the woman but didn't flag her. It thought, you know, it probably put her in the category of like a paper bag or something like that—

**Martin:** Right, right.

**Rumman:** —so it's gonna be fine to hit this thing. So that's a really good example that brings up a lot of the moving parts when we think about responsible AI, what it means and how we build it.

**Martin:** And I suppose one thing that I think about there is that in that like sort of circumstance, ultimately, you know, there is a black and white, there is a right and wrong decision to be made, and we just need to make sure that the algorithm gets on the right side of it. In many cases, in content moderation, for example, the judgments that are being made on content are essentially judgments that would have been made previously by humans. And I suppose, you know, they could be carrying human prejudice, human biases into that decision-making process. I assume that's part of the thinking here. How do you stop those human prejudices just carrying forward into the algorithms that are going to be making the decisions on behalf of the humans going forwards.

**Rumman:** Absolutely. So one of the focuses that we have in our responsible AI practice is enabling human-centric AI.

So how do we decide what algorithms, what decisions algorithms should be making and what they should not be making. So we can think about how we use AI into basically two buckets. One is the AI enables the human to make a better decision. That's the enhanced judgment that we talk about a lot at Accenture. And the second half would be AI is making decisions because a human being can't.

An example of that, let's say, in content moderation, actually, is there is just a sheer volume of videos. It is difficult, if not impossible, for human beings, no matter how many bodies you throw at it, to watch every single video out there. One might even make a case that human beings shouldn't be watching some of the



videos and content that content moderation is supposed to flag. If we can create an AI that can identify that, all the better. But the other end, of course, is when AI should be enabling human decision-making, when AI is being used in very critical decisions.

For example, whether or not someone should be granted a loan, whether or not someone should receive bail or get parole. And these are decisions that we may decide should actually ultimately have human approval, even if it is guided by artificial intelligence. So there are very specific design decisions in how the human and the AI interact with each other, and how do we create this AI so it's not dictating an answer, but instead it is helping inform an answer, the person is able to poke and prod and ask, essentially, why a decision was made.

**Martin:** One thing that that brings to mind, actually, is obviously many of the companies that we're talking about here, the decisions that we're talking about, we can think of them as starting in, you know, in the U.S., for example, localized, but then very quickly those same decisions, that same kind of content review activity is going to have to roll out round the world. And in many cases those judgments, you know, you may need to be cognizant of local nuances to the judgments that are being made.

How does that play a part into designing a responsible AI?

**Rumman:** Absolutely. There's so much contextual awareness that has to happen. So a lot of this conversation really can and should happen at the local level, so engaging local teams, local individuals to say what is the cultural lens by which we view responsible AI or good AI. And I'll give you a perfect example. The GDPR comes out of a very European sensibility or priority placed on data privacy and security. You don't have that same level of privacy and security emphasis in other countries. China would be a good example.

But then we think about, let's say, in the U.S., what are our priorities here. Data privacy and security also don't rank as highly as something like, let's say, algorithmic bias or

media addiction. There's a lot of conversation here about addictive media that we don't have elsewhere.

We move to the Asian market, in Japan and in South Korea, because there is so much interest in robotics, which actually isn't as much of an interest everywhere else, there's a lot of talk about good AI, ethical AI, responsible AI being built around engaging with a robot. And that's a conversation we wouldn't necessarily have in the U.S. There's a lot of thoughts about robots for elder care, robots for childcare, for example, that those are very culturally nuanced conversations, absolutely.

**Martin:** Well, unfortunately, that's all we've got time for today. But

**Rumman,** thank you so much for taking part in today's episode.

**Rumman:** It was my pleasure. Thank you for having me.

**Martin:** Thank you. You're welcome. So we hope you all enjoyed today's conversation, and we hope that you will tune in again for the next episode of *On The Platform*.

Copyright © 2019 Accenture  
All rights reserved.

Accenture, its logo, and High  
Performance Delivered are  
trademarks of Accenture.