TRANSCRIPT OF SEPTEMBER 26, 2024 DIGITAL EVENT:
AI-Driven Business Intelligence: Transforming Business Intelligence into Unparalleled Competitive Advantage with Daniel Hulme

INTRODUCTORY SUMMARY

We are at the dawn of a new era as artificial intelligence and machine learning redefine the business intelligence landscape. They offer unprecedented insights and decision-making capabilities—if you can capitalize on them.

At the same time, every industry faces challenges: the costs of operating legacy systems and outdated IT infrastructures, new skills versus old processes, new thinking versus old ways…

We spoke with **Dr. Daniel Hulme** to find out how we might shake off the old ways and become AI-driven organizations capable of capitalizing on business intelligence**.**

Daniel is a globally recognized Artificial Intelligence expert and an Emerging Technologies investor. He's the CEO of Satalia, an award-winning company that provides AI products and solutions for global companies. Satalia joined the world's largest marketing company in 2021, WPP, where Daniel is now the Chief AI Officer, helping define, identify, curate, and promote AI capability and new opportunities for the benefit of the wider group and society. And where he has been recognized as one of the top 10 Chief AI Officers globally.

Daniel started the session with a short presentation entitled: **Rethinking AI and its impact on marketing and humanity**.

TRANSCRIPT

[Presentation starts: Slides 1-20]

Daniel Hulme
I might be a bit bolder, actually, and say that this is rethinking AI and its impact on business and society. So when I do these kinds of talks, I usually cover three things. One is how we can practically bring AI to your organizations in a safe, responsible, and sustainable way. The second is how we can bring people and technology together to do amazing things. And then I also talk about the macro impact these technologies all have on society.

We can talk about superintelligence. We can talk about the impact on jobs. We can talk about post-truth worlds. What I'm going to do is I'm going to park the last two sections, and if people are interested in that, I'm very happy to lean into it. But I guess the first section is the more practical piece, which is the applications of AI.

I have been involved in AI for over 25 years. I am currently entrepreneur in residence for UCL (University College London), where I help spin out deep tech companies. I sold my AI company, which I've been developing for 17 years, to WPP, the biggest marketing company in the world. I sold that three years ago. I'm continuing to be the CEO of Satalia. So I continue to grow that organization within WPP and think of them as being WPP's Deep Mind. I'm also the Chief AI Officer at WPP, so I coordinate AI across about 120,000 people, which is good fun. And I invest in,

and I start companies. I have just started a company to try to solve machine consciousness. Very happy to lean into that as well if you'd like.

Organizations have been investing a lot of money over the past ten years with the hope that extracting insights leads to better decisions. Turns out that giving human beings better insights doesn't lead to better decisions. Decision-making is a completely different field in computer science. If you're old enough, it used to be called operations research. It's discrete mathematics and optimization. I would argue you need to start with the decision-making problem first and then work backward. Don't start with data.

Then, by the way, it turns out that human beings are rubbish at making decisions. I don't know if you read the book by Daniel Kahneman, *Thinking Fast and Slow*. Daniel Kahneman is a Nobel Prize-winning economist, and he argues that we have a fast brain and a slow brain, and there are parallels in terms of how these brains work with how AIs work.

Now, what I usually do is make these sessions very interactive, but I don't have time today to go through and explain to you why human beings are rubbish at decision-making. Other than maybe using this example, which is: If I have a delivery van to deliver packages around these 24 points around London, then humans are quite good at solving these spatial problems. After a few minutes, we'll draw a nice path around those points. Now, how long will it take you to get to the shortest path? The one that's going to save the most amount of time and petrol. Maybe you're thinking seconds, and you'd be right. It's somewhere between seconds and 20 billion years.

Put 24 times 23 times 22, etc., into a calculator, and you get the total number of routes. That's how many routes there are around all 24 points. If you had a computer that could look at a million routes a second, it would take 20 billion years to go through every single possible one and say this one I looked at X number of years ago was the shortest one. If I add another point on the map, it will be 25 times 20 billion years or 500 billion years.

The point here is that these are exponential problems. And there are many, many of these problems across your organization. I bet that you're either using human beings to solve these problems badly, or you're probably using algorithms that we developed ten or 20 years ago that are also solving them badly. Algorithms aren't a differentiator for your business. Now, if I build a system that I give data to and it makes a decision, tomorrow I give it the same data, and it makes the same decision; what I have is automation, and automation is amazing because you get computers to do things better than human beings. But the definition of stupidity is doing the same thing over again and expecting a different answer. So, I argue controversially that, by definition, automation is stupid. It's not AI.

I know that everybody who currently touches these words calls themselves an AI company. Fine. You get more clients, and you get more funding.

There are many definitions of AI. Unfortunately, the most popular definition, I think, is the weakest definition, which is getting computers to do things that humans can do. Over the past ten years, we've managed to get machines to recognize objects and images, to correspond in natural language. When we get machines to behave like humans, because humans are the most intelligent thing we know in the universe, we assume that that's intelligence. I would argue that "acting human" is not intelligent. That's a different conversation. There's a much better definition of AI, which comes from the definition of intelligence, which is "goal-directed adaptive behavior".

Goal-directed in the sense that you're trying to maximize your vehicles to maximize deliveries, you're trying to maximize your marketing spend to maximize reach, to maximize your workforce to maximize utilization. You have to have a goal. But the key word for me in this definition is the word "adaptive". What you want to do is build systems that make decisions, learn about whether those decisions are good or bad, and adapt themselves so next time you make better decisions.

And if I'm being honest with you all, you tend not to see adaptive systems in production. Most systems in production are automation. They're not necessarily intelligent. Now, if I held what we are currently doing in industry to this definition, well, I might controversially argue that nobody's doing AI. That's a ridiculous statement because I've built my career around AI. And the fact is that looking at AI through definitions and through technologies is a waste of time. And there's a much better way of thinking about AI.

I would argue that there are six applications and six flavors of friction that exist across any supply chain in any industry. So, all the frictions that you experience in your business can be mapped to one of these six applications.

The fact is task automation is using very simple algorithms. But the fact is, if you apply simple algorithms in the right way, you will drive a massive amount of value. You don't need to gravitate to shiny technologies like generative AI and machine learning.

The second category is content generation. Everybody now has the ability to generate generic content. WPP creates, I think, a quarter of the world's content in terms of ads. This is very important to us. The battleground is not creating generic content. The battleground is creating brand-specific, production-grade, differentiating content. I am very happy to go into detail about how we do that at WPP.

The third category is human representation. So, a few years ago, I would have been talking about how we can use AI to replace human beings with things that look and behave like human beings. I think what's interesting about large language models is that they enable us to not just create content but recreate how people perceive content. And that's a very powerful tool. We call those audience brains. I'm very happy to expound on those later.

The fourth category is what people called AI before generative AI: Machine learning and data science. You already know my opinion. Extracting insights from data doesn't lead to better decisions. The power of machine learning is not making predictions. The power of machine learning is explaining those predictions. So I'm a massive advocate of being able to build explainable algorithms, explanations that help you understand the world and make better decisions. I've already labored on the point about complex decision-making, but I just want to labor it even more because, I would argue, you get the biggest bang for your buck if you apply these types of algorithms to your organizations.

So if you have five staff members, for example, to allocate to five jobs, there are 120 possible solutions. Five times four times three times two times one. If I have 15 people to allocate to 15 jobs, there are a trillion solutions. To expect a human to solve this problem, we're wasting our time. 500 staff members, the number of possible combinations is a big number. And just to put this number in context, this is how many atoms there are in the universe. Optimization algorithms are a differentiator for your business. Not necessarily generative AI, not necessarily machine learning.

The final category is human augmentation. So this is where we can now build digital twins of people. So I can build a digital twin of my employee, trained on their email, their calendar, and their feedback. And I can ask that digital twin, if I put you on this project will you work well? If I put you on this team, will you thrive?

We're building digital twins of employees and our personal lives, and the marketing industry needs to learn to market to both people and AI.

These six pathways allow you to navigate this complex world of AI – safety, security, ethics, and governance. There is a huge amount of misunderstanding and misinformation about these words. I know there are lots of people that have rebranded themselves as AI ethicists. I would argue there's no such thing as AI ethics. Ethics is a study of right and wrong. And the difference between AI and human beings is that AI doesn't have intent. Human beings have intent. And it's the intent that needs to be scrutinized from an ethics perspective. There are several other questions that you

need to ask yourself when inventing AI safely and responsibly. I'm happy to talk about that later. And I think that's probably where I'll close.

Other than that, we know that AI used to go out and solve problems across our supply chains. I would argue that, ultimately, what companies are trying to do is create a digital twin, a digital simulation of their entire supply chain. I think there are three digital twins you need to create. One is the flow of goods across your business model. The second is your back office processes – hiring, firing, onboarding, and offboarding. AI is not only making them more efficient, but they're also making them more effective. They're changing the way that we do back-office processes.

Finally, your workforce. Those people have hopes, dreams, desires, and skills. Companies organize themselves as hierarchies. Hierarchies are not conducive to adaptation. Remember, the more adaptive you are, the more intelligent you are. What we need to do is use AI to identify people and allocate them to work in a much more agile way.

I'm just going to close down by saying that there are three things that differentiate an organization. It's not technology. The three things are: data. It's data that makes the AI smart. If you have data that contains insights better than your competitors, you're going to win. The second is AI talent. If you want to differentiate yourself from your competitors, you need to use deep expertise in AI. And most organizations think they can attract AI talent, but they can't. You really have to be so pragmatic about your ability to attract, retain, and motivate deep AI talent.

And finally, the third differentiator is you – whoever's listening – it's leadership. It's the people who know and ultimately understand the transformational power of these technologies. If you apply the wrong technologies to solve the wrong problems, you're going to make a massive amount of misinvestment over the next few years, create a lot of tech debt, and essentially lose out against your competitors. On that note, I'm going to stop. I can talk on and on and on. I guess we can have a good discussion.

[Presentation ends]

Adam Burns
Where is AI already driving substantive ROI and competitive edge?

Daniel Hulme
Well, I guess it depends on what you call AI. So if you're applying optimization algorithms and, you know, 15 years ago, we built optimization solutions for Tesco, for PwC, where we can now allocate their resources 10-15% more effectively. You know, every 1% saving for those organizations and they're saving like tens if not hundreds of millions of dollars. So, applying optimization algorithms that have happened over the past two decades has driven a massive amount of value.

I would argue that machine learning hasn't driven the value that people think it should have driven because, you know, giving insights to human beings doesn't make better decisions.

Generative AI allows you to know things about the world and allows you to access that knowledge in nice, useful ways. So industries where creating content is important, like marketing, those industries are being disrupted. So we can now create ads, create content essentially thousands of times faster.

You won't be able to use AI to route your vehicles thousands of times more efficiently. It's not the right technology. So be pragmatic about making sure you apply the right types of AI to solving the right problems.

Sometimes, it literally is a very simple algorithm that you need to implement, and it will drive a massive amount of value. If you're talking about generative AI, then, you know, what we're seeing is organizations trying to use generative AI to write faster emails, to create PowerPoint slides better. That will probably improve productivity by 20-30%, but it won't differentiate your business.

Where you're in the business again of content creation of trying to understand audience perception, that's where generative AI is extremely powerful.

Adam Burns
How do we get from deciding better to acting smarter?

Daniel Hulme
So I think what you need to do is you need to be able to empower AIs to make decisions. And essentially remove, not the human entirely, it's very, very important to have a human in the loop, but, you know, humans are very good at finding patterns in four dimensions. We're very good at solving problems upto seven. But what we have a lot of is human beings often making decisions that they shouldn't be making. There's enough literature out there to show you that human beings are not good at making certain decisions. So I would automate that decision-making process. And where the AI struggles, that's where you bring the human to bring in those additional insights. So, I would be looking at places where human beings are currently making decisions and replacing them with an automation. By an AI.

Adam Burns
What are some of the most common pitfalls of implementing new technologies into existing processes?

Daniel Hulme
I would actually argue there are really only three questions you need to ask yourself when implementing these technologies. So, first of all, you need to make sure that you're applying the right technology. So those six applications are a good framework to start thinking about the right approach.

The three questions… The first question we've already mentioned is ethical: Is the intent appropriate? There are already well-established processes and structures to scrutinize intent. You don't need a whole new ethics committee associated with AI.

The second two questions are safety problems. Really, the only difference between AI and software is that AIs are typically opaque in how they make their decisions. And if you can make your algorithms explainable, which, by the way, is very hard, which is why you need deep AI talent. If you can make them explainable, it solves a lot of those words. It makes them transparent, it makes them auditable, it makes them governable. So explainability, I think, is actually the key challenge to solve technologically when implementing AI in production. Not only to adhere to the AI EU Act but to be able to make sure that you're using that explainability to surface insights that will allow you to make decisions.

The final question, which is interesting, is not, what happens if my AI goes wrong? We already know how to engineer systems, identify failure points and then mitigate those failure points. You have to now ask yourself what happens if my AI goes very right? If you implement AI correctly, it moves the needle significantly. There are many examples of this, and it can massively overachieve the KPI you set. Now, what it will do then is potentially cause harm elsewhere across that supply chain.

So, you know, for one of my clients, we built a workforce allocation solution, so we built an AI that essentially allocates 5000 people jobs. Now, historically, they would have 40 planners in a room doing that as their full-time job. Their full-time job was trying to figure out how to allocate 5000 planners to jobs. And we built an algorithm that can solve that problem in two hours. Now, the KPI for success for that project was if you can increase utilization, so if your algorithm can unlock 2.5% more utilization, that's billable hours, then the project will be an overwhelming success. We built an algorithm that increased utilization by 12.5%. But, if you looked at the resulting plans, what happened was people were driving long distances to clients, clients weren't getting the continuity they wanted, and employees didn't have time to train. So what it did was overachieve on one KPI, but it caused harm across other parts of the supply chain.

So, when you implement AI, you need to optimize not just for one KPI but holistically across multiple KPIs. AI can positively affect all of them.

Adam Burns
AI might help free up more time for people within an organization to be creative. Are organizations set up correctly to maximize this? Do we have the right metrics in place for an AI-enabled staff?

Daniel Hulme
Yeah. Absolutely. So I would argue that there's applying AI to improving corporate activities. Like writing emails faster and PowerPoints and things like that. And there's also applying AI to make your supply chain more effective and more efficient. So we create ads, and now we can create ads in seconds instead of weeks. This allows our creatives to create much more content. We can test our content against synthetic audiences to see how they think and feel about those ads, to be able to create better content. So what would have taken literally months and months and months can now be done in a day.

So I think, you know, it depends on the type of friction, the type of ask that you're talking about. But what we've seen, not just in using generative AI but in using other types of AI, is that if you apply them correctly, you can drive a huge amount of value.

Now, as I said at the beginning, my worry is that we gravitate towards shiny technologies, and we think that generative AI is a panacea for solving all problems across our organizations. And when it doesn't work, people say, well, you know, it's a failure of the technology. There is an AI bubble. It's nonsense. These technologies are incredibly powerful. These technologies have been around for many, many decades. The issue, the challenge, is people not using them in the right way.

Adam Burns
How do you know if you are using them in the right way?

Daniel Hulme
Ultimately, success is people using it. We built software, and people are using it because they value those technologies and tools because they make their jobs better. So success is people valuing it, either through their time or paying for it.

Adam Burns
And which bit of moving toward becoming an AI-enabled organization is the hardest for a business to get right?

Daniel Hulme
I think there are two extremes. There are some organizations I know that haven't even touched AI. They've got no idea really of its capabilities or power. They're just scared of it. Our challenge at WPP was preventing people from using AI. All the creatives are going out; they're using these large language models and technologies to enable them to do their jobs better.

Unfortunately, some of those models are trained on data that we can't determine the provenance of. So you might end up creating an ad that violates copyright and all that kind of stuff. There's a risk of uploading client data sometimes, potentially, to large language models and having client data be available elsewhere. So we had to rapidly get control of our 100,000 employees to make sure they were using the right technologies in the right way across our supply chain.

So on the one extreme, you have organizations that are running away from it that are really scared of these technologies, all the way through to other organizations where it's being embraced wholeheartedly. And yeah, I think ultimately, when you start playing around with these technologies, you realize how useful they are, and then the proof is in the pudding. People only use them if they feel like they're valuable. And I think if we look at the trends in terms of AI usage

over the past three years, we'll see a massive uplift as well as a continued uplift of the usage of these technologies.

Adam Burns
So if we want to move towards being an AI-enabled organization, we need to get our people using AI. Do we just need to give them access, or do we need to make it fun?

Daniel Hulme
I don't know. I mean, I'm actually not a huge power user of AI, if I'm honest. Because, you know, of my job in some respects. As I said, AI is very good at finding generic information. My job as Chief AI Officer is to understand the future. And the future is not currently encapsulated in these large language models. They can tell you the current and the past. They can't tell you the future. But, you know, I think that anybody who uses these technologies will probably suggest that they're great tools for bouncing ideas against. They're great tools for synthesizing and capturing information. They're great tools to create stories for your kids.

So, I think, you know, large language models are perhaps the most powerful technology that we've ever created. At the moment, they're a little bit like intoxicated graduates. You know, they they get things wrong a lot. I would argue that they're now graduating to a Master's level, and in another 18 months, they'll graduate to a PhD level, where you are going to be able to give it a complex task. It will break that task down, create a hypothesis, and do experiments. Another 18 months after that, it will be a post-doc. And another 18 months after that, I will be a professor.

So, my prediction is that by the end of this decade, we will have a professor in our pocket. Now, how quickly that will percolate into industry, stimulate new innovations, and be adopted and embraced by people, I don't know. But that's the trajectory, the power of these technologies.

Adam Burns
What are the essential skills that every business will need within the next two or three years?

Daniel Hulme
I think being able to ask good questions of the AI is super important. So we found that with our creatives, the ones that are able to ask questions about, you know, some obscure artist in history or some historical event or some, I don't know, lighting and aperture setting. Those people are able to actually get better answers from a creative perspective than certainly me. So, developing your kind of broad understanding of humanity, philosophy, psychology, and history enables you to get more out of AIs.

So then, obviously, there's AI talent. You know, many organizations think they can attract and retain AI talent. The reality is that they can't. They might hire some, you know, relatively mediocre people that they'll lose in three years' time and again end up having systems and tools that are not sustainable, not supported, and not innovative. And, you know, I would ask yourself, if you want to build a differentiated supply chain using these technologies, you're going to have to engage with differentiated talent. And there's only a handful of DeepMind-type companies around. I'd argue that Satalia is one of them. So figure out how not to build your own Satalia but how you would engage with that type of organization.

And by the way, my talent want to be doing things that are purposeful. So if you have a strong purpose, rather than just a strong business model, it's actually purpose that's going to allow you to attract talent one way or the other. But it's also going to allow you to attract customers. So I'm a massive believer in the power of purpose.

Adam Burns
What's coming down the line?

[Presentation starts: Slides 21-28]

Daniel Hulme
I think everybody has heard of the singularity. The word comes from physics. It is a point in time that we can't see beyond. The AI community adopted it to refer to the technological singularity, which is the point in time when we build a brain a million times smarter than us.

So you can see there are six singularities, six points in time, that we have to face as a species over the next decade or so. And I'm going to use this macro framework, PESTLE, to capture them.

[Note: PESTLE stands for Political, Environmental, Social, Technological, Legal, Economic.]

The political singularity is essentially a post-truth world of AI. Misinformation bots and deep fakes have challenged our political foundations and continue to do so. But they now enable people to make clones of each other, which are being used to challenge the fabric of reality. Now, I think we can solve this problem if we use AI to identify and authenticate problematic content.

The environmental singularity is consumption. As far as I'm concerned, consumption gives people access to goods and services that enrich our lives. But the fact is that we are putting pressure on our planetary boundaries. And I know that if we apply these technologies in the right way, if we apply them by algorithms, we get at least half the amount of energy that we need to run this planet. And I know that because every project I've done, we typically reduce the amount of carbon by 25%. If you can optimize your supply chains, if you can create digital twins of organizations, you can reduce the amount of carbon even more.

This is not my expertise, but the social singularity is when we cure death. There are scientists who believe there are people alive today who don't have to die. AI is preventative medicine. It is able to monitor ourselves and clean ourselves out. A bit like a car, you stop any damage, and that car will never have a breakdown. And I don't know what the world will look like when we realize very soon that people amongst us won't have to die.

The technological singularity is when we build the brain a million times smarter than us. And we become the second most powerful species on this planet. We think that's now going to happen in the next ten years, maybe 20 years. And my advice to people is that when it comes, look busy, be nice to each other, and hopefully, it will bugger off to a different dimension.

I've actually just started a research company invested in by WPP to solve machine consciousness. I believe that machine consciousness is going to become very important over the next several years.

And by the way, whilst large language models are incredibly energy efficient, they take a long time to learn. There is a new, emerging technology called neuromorphic computing, modeled on how brains work. Our brains operate on the power of a lightbulb. We learn very quickly. And the fact is, these technologies are going to spawn the robotic revolution. And those robots are going to start to realize self and potentially become conscious.

The legal singularity is when surveillance becomes ubiquitous. I'm in the business of understanding perception and trying to influence or persuade a change in perception. And that's an incredibly powerful position to be in. What we want to do is make sure that we're using that power responsibly and safely and preventing bad actors from using it to accumulate more wealth and power.

Finally, economic singularity concerns the impact on jobs. It was coined by a very good friend of mine, Calum Chase. For the past 18 years and the next 10 years, AI has been freeing people from mundane tasks, enabling them to do more important, impactful, and purposeful things. I think there is going to be a Cambrian explosion of new innovations created by AI. It will displace and disrupt jobs, but it's going to allow humanity to grow.

So beyond ten years, nobody knows what they're talking about. Just to give you two extremes of the argument… The one extreme of the argument is that when you can fill whole jobs, we probably will. And if that happens very quickly, we create an economic imbalance. Our economies won't be balanced. It will lead to social unrest. We need a lot of mechanisms to mitigate this risk. We can think of the UBI [Universal Basic Income], four-day working weeks, etc.

The other extreme of the argument, which is controversial but bear with me, is that we should accelerate as fast as possible towards the singularity. We should automate food, health care, education, energy, and transport production and dissemination. We should remove the friction from the creation of those goods, bringing the cost down to zero.

So imagine being born into a world where paid work isn't available, but everything you need to survive and thrive as a human being is free.

Now people say to me, Daniel, what would I do without a paid job? A job defines who I am as a human being. I know lots of people who don't have paid jobs. They're not sitting at home bored and depressed. They're using their time and assets to contribute to humanity. And I usually ask my audience: What would you do if you didn't have to do paid work and everything is free? And the first answer usually is golf. The second answer usually is travel. I'm indulging in my hobbies, reading, and spending more time with my friends and family. If you keep pushing people, they say the same thing. They'll do all those things and then do something that tries to make the world a little bit better.

I think we all have an innate desire to make the world better, but we are all pretty much born into economic constraints that prevent us from doing this. AI has the promise, if implemented correctly, to free everybody from those economic constraints.

I'm just going to close by saying I joined WPP because I believed in this purpose. WPP's purpose is to use the power of creativity to make the world better. I think a better world, is a world where everybody is economically free to use their creativity to live beyond themselves. I think that if we apply AI in the right way, we can create this future.

Adam Burns
You said of the political singularity, "I think we can solve this problem if we use AI to identify and authenticate problematic content." But we're not using it…

Daniel Hulme
No, so what we're doing at WPP when we're creating content across our supply chain, we can show that content to synthetic audiences – brains that essentially represent a segment. But what you can also do is build brains that represent every corner of society. You can build a brain that will represent a minority group, a political party, a culture, a newspaper, or even an ad-compliance-claims framework or a sustainability framework. And when we are then creating content across our supply chain, you can essentially show it to these thousands of representatives to make sure that you're not triggering any communities, breaking any laws, or causing any harm.

Actually, I'm an advocate for essentially trying to get policy or government to say if you are creating content of any form, you need to first determine the harmfulness of that content, number one, by potentially using this mechanism. And two, if it is deemed to be harmful, you need to then authenticate it. Does it come from, you know, a human being that has their own intent and whatever? Or does it come from some sort of misinformation bot designed to polarize people? We do have a mechanism that we could use to determine how harmful content is, whether it's a tweet, a newspaper article, or an ad. And then say, does this come from an authentic trusted source?

Adam Burns
Can we revisit the idea of AI consciousness? I think that's something that does worry people. What is your new company trying to achieve safely?

Daniel Hulme
The company's north star is trying to understand consciousness in pursuit of building safe and efficient AI. When we use the word "safe", we're not just talking about safe for human beings; we're also talking about safe for AI.

Over the next several years, I think people are going to think that AIs are conscious, and they're not, which I'm not too worried about. But there's also a concept called mind crime, which is we might end up building AIs that are conscious, that can suffer, that do suffer, and not realizing it. And we have a duty of care, not just as human beings and to animals, but we potentially now have the duty of care to AIs.

And I know that sounds crazy, but these are the questions we'll be facing over the next ten years.

I believe that we can actually build AIs, evolve them, and align them to a value system. At the moment, people are building big brains—large language models—and then essentially playing whack-a-mole, trying to prevent those LLMs from doing anything bad.

That's a very silly thing to do.

What we can do is we can actually evolve AIs in environments where the environment can actually enforce a value system. You could create an environment and evolve AIs in that environment where you have to fight and kill each other to survive. Or you could create an environment where you have to cooperate and be altruistic, sacrificial, and kind to survive.

I believe we can do this. We can build AIs with embedded value systems. I guess the question I ask people is if I gave you a choice between a superintelligence that's a zombie and a superintelligence that has consciousness, which one would you choose?

Let me just give you an example of a zombie superintelligence. Imagine you created a super-capable AI and gave it the task of solving cancer. The easiest way of solving cancer is by eradicating human beings. Now, that AI might be very, very good at achieving its goal, but it could be very harmful for itself and the rest of humanity.

So, I guess the point about potentially building consciousness into a machine is that, first of all, it might respect other things that are conscious, but it might also be able to be aware of the decisions that it's making—not just on itself but on other species.

We also realize that building conscious AIs might be a very silly thing to do. We need to pause it. In this case, I would be advocating as hard as possible to make sure that we don't build these technologies.

Adam Burns
What type of person is the hardest to convince that AI can help them decide better and act smarter?

Daniel Hulme [
There are usually two. One that has been doing their job for the past 30 years and thinks there's no way that some mathematicians or AI are going to come along and be able to do that job better than them.

The fact is that sometimes you can't. But we've learned to take people on the journey. We've learned, first of all, to validate their understanding and hypotheses about the world. So rather than telling them the stuff that they don't know, we've kind of told them stuff that they do know. We've empowered them with data and numbers. And then once you start to win their trust in the usefulness of these technologies, then you can start to show them things that they don't know. So that's the first challenging stakeholder.

The second challenging stakeholder, if I'm honest with you, is clients. So there are people out there who want to build a career out of AI, and they're being encouraged by consultants who have rebranded themselves as AI experts. They're saying, "You can go build these technologies yourself, we'll help you build them, we'll help you support them." And the reality is, if you've never built, scaled, and deployed software in your organization, you are not going to do the same thing for AI.

And my worry is, is that there are consultants out there that are maybe, you know, through no fault of their own, encouraging organizations, trying to get them to go build their own technologies. And again, unfortunately, this ends up being a massive mis-investment. People think it's the fault of technology, and really, it's the fault of human decision-making.

Adam Burns
Fantastic. Thank you very much indeed, Daniel.