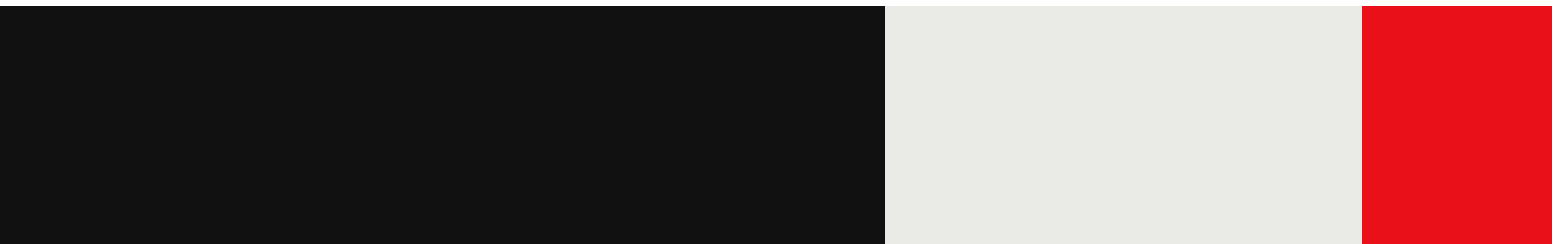Jacob Bourne, and Gadjo Sevilla

The Daily: What happens when AI learns from other AI content, and new features for the Apple Vision Pro

Audio

On today's episode, we discuss what lawmakers are most likely to tackle first when it comes to regulating AI, whether AI songs can win a Grammy, and what happens when AI eats up—and learns from—other AI-generated content. "In Other News," we talk about the newly announced features for Apple's Vision Pro AR headset and how this device could change the whole market. Tune in to the discussion with our analysts Jacob Bourne and Gadjo Sevilla.

Subscribe to the "Behind the Numbers" podcast on Apple Podcasts, Spotify, Pandora, Stitcher, Podbean or wherever you listen to podcasts. Follow us on Instagram

Verisk Marketing Solutions, a business unit of Verisk formed through the integration of Infutor and Jornaya, empowers marketers and platform partners with deep consumer insights to power precise and personalized omnichannel interactions. Verisk Marketing Solutions provides a unique combination of identity resolution, consumer attributes and behavioral data that integrates with marketers' existing technology and evolves with consumers' ever-changing behavior—all while maintaining the highest data security and privacy standards. To learn more about the consumer intelligence solutions available through Verisk Marketing Solutions, visit marketing.verisk.com.

Episode Transcript:

Marcus Johnson:

This episode is sponsored by Verisk Marketing Solutions. Verisk Marketing Solutions enables brands and platform partners to remove the guesswork around segmentation, timing, and

messaging to help you continuously maintain a complete picture of your customers and prospects identity, attributes and in-market behaviors as they change over time, confidently personalize interactions to the right person at the right time. Visit marketing.verisk.com for more information.

Jacob Bourne:

And it's kind of giving Gemini a new kind of problem solving skills based on interacting with the physical world as opposed to just being trained on internet data, and how it's doing that is actually it's fusing robotics and generative AI.

Marcus Johnson:

Hey gang, it's Thursday, July 6th. From my voice, you can tell that it's the morning again. Why do we keep doing this? I hope you had nice long weekends, Gadjo, Jacob and listeners. Welcome to the Behind the Numbers Daily, an eMarketer podcast made possible by Verisk Marketing Solutions. I'm Marcus, and today I'm joined by one of our senior analysts on the connectivity and tech briefings, usually based out of New York, it's Gadjo Sevilla.

Gadjo Sevilla:

Hi, Marcus. Hi, Jacob.

Marcus Johnson:

Hey fella. Hello. Hello. And also, we're joined by someone else on that connectivity and tech briefing team, one of our analysts based on the other coast, in California is Jacob Bourne.

Jacob Bourne:

Morning, Marcus. Morning, Gadjo.

Marcus Johnson:

Morning fella. So today's fact, I'm tired. Apart from that, the can opener wasn't invented until almost 50 years after the can. What's going on? The first way to open cans was using a hammer and chisel. What the hell were we doing? Peter Durans, an English merchant, received the first patent for the idea of preserving food using tin cans back in 1810. Then for some ungodly reason, it took 50 years for the first modern day can opener that we know today to

be created by an American inventor patented by Ezra J. Warner. That's too long. I don't know how committed I would be to the can of alphabet spaghetti. Smash it open.

Jacob Bourne:

What happens when you don't have AI inventing things, I guess.

Marcus Johnson:

Exactly. So glad I was born later than this time, let's just say that. Today's real topic, AI, regulation rumblings, Grammy rules and chatbots eating themselves.

In today's episode, we'll cover how artificial intelligence might be governed and what happens when AI eats itself. Then for in other news, we'll discuss the impact of Apple's Vision Pro MR Headset. We start, of course with the lead, talking AI. Let's kick off with some regulations. US Senate majority leader Chuck Schumer, has launched an all hands on deck push to regulate AI, writes Christiano Vemer of the Washington Post. He explains that the high profile speech is expected to kick off a wave of legislative activity as lawmakers face pressure to set protections to prevent AI from being abused.

Mr. Schumer has dubbed the initiative, the Safe Innovation Framework with the name being an acronym for the goals of Security Accountability Foundations and Explain. One of its proposals is that a user should have the ability to ask how the AI chose its answer over another plausible option where that information came from. So that's one thing that maybe we'll start with in terms of getting regulated. But gents, what are lawmakers most likely in your opinion, to try and tackle first when it comes to AI regulation?

Gadjo Sevilla:

I think the first step would be to agree on the foundations and the parameters for regulation. I mean, they would need to understand basically how the technology works on a much granular level than perhaps their legislative rules would suggest. So it remains to be seen if they can get the buy-in from AI leaders, because definitely that's going to be an important component, and if the lawmakers can set the groundwork based on that. Either way, I personally expect there to be a lot of back and forth and it could take months to get a working framework with all the parties involved.

Marcus Johnson:

Jacob, there's a lot of places they could "start". I'm sure that there'll be lots of different rules coming out at the same time or within the same bill, but there's a few other US legislative efforts that are being kicked around. One centered around a bill requiring the US government to be transparent about using AI in interacting with the public, and it would let folks appeal decisions made by AI. That's one. There's a bill to establish in office to ensure the US remains a global AI front-runner. There's rules for testing AI systems before they're publicly available. There's protecting privacy rights. Where do you think they begin?

Jacob Bourne:

I actually think that they're going to start narrow. So things like algorithmic bias, misinformation, copyright infringement are going to be kind of the initial steps, especially with the bias. That's something that's kind of the age-old AI concern. We've already seen some tangible harm on that front. I think that's going to be the initial focus. I think you brought up a good point when you mentioned that bill for the agency to help the US stay an innovation front-runner, because one of the things that's happening around this call for AI regulation in the US is actually it seems like the government's kind of using it as an excuse to ensure that we see that it stays at the forefront of innovation.

I mean, originally at that congressional hearing with OpenAI and IBM, that agency idea that was floated was actually an agency to provide AI industry oversight and that it very quickly turned into ensure innovation. I think that's going to also be a running theme [inaudible 00:06:19]. The US, there's this call for regulation, but the US does not want to fall behind China, especially on the AI front. And so while we will see some efforts to regulate it, I think that the emphasis is actually going to be on innovation.

Marcus Johnson:

Yeah. Yeah, you were pointing out in an article you just wrote, that's a tough balancing act that they've got to try to figure out. Professor Yann LeCun, one of the three godfathers of AI, who is... Jacob, you said his title was?

Jacob Bourne:

He's a chief AI research scientist in meta.

Marcus Johnson:

That's it. He thinks each application would need its own rules even. For example, different rules would govern AI systems in cars. Different rules would govern those that scan medical images. This could get real messy, real quick, really expansive real quick. That pays to regulation in the US, as we mentioned before, significantly trials the European Union. They've just advanced an AI bill that after years of talks called the AI Act, the first major law to regulate AI and a potential model for policymakers around the world, Adam Santarion of the New York Times was noting, [inaudible 00:07:25] AI Act include a few things.

So European bill takes a risk-based approach to regulating AI, focusing on applications with the greatest potential for human harm. Ai, for example, AI that operates critical infrastructure like water and energy, AI and legal systems, AI used when determining public services, government benefits. AI companies would also have to conduct risk assessments before putting the tech into everyday use. Think of drug approval processes. Then also this bill would severely curtail the use of facial recognition software. They're still trying to iron that part out.

And then also, requiring makers of AI systems, ChatGPT, to disclose more about the data used to create their program. Just a couple of parts of this bill, major parts of this bill. What do we make of this bill? What has it got incredibly right, maybe a little bit wrong? Do we think this could actually be a model for others?

Jacob Bourne:

It's certainly very robust, man. Outside of China, one of the first very meaningful attempts to regulate AI. It was first proposed back in April 2021 and we haven't seen it get passed yet, largely because there's been so much change in the AI industry, namely like the biggest change being the rise of generative AI.

Marcus Johnson:

I know they probably had a lot ironed out until then and they were like, "Bloody hell."

Jacob Bourne:

Yes. Absolutely. [inaudible 00:08:43] they had to almost start from the drawing board here. I think it's expected to pass by the end of the year and I expect that we'll see further changes made to this draft in terms of how they approach this risk-based regulatory framework, whether they include general purpose foundational models within that. One of the things that

they're asking tech companies to do is to publish summaries of copyrighted materials that's used in model training.

The people building these models are saying, "Hey, this isn't technically feasible to do." I think what they mean there is doing that would be so expensive. It would be too expensive in light of the other very high cost associated with training AI models. It's something that they're seeing, they're looking at and saying, "This is going to be a further barrier to these AI startups becoming profitable." And so I think that's really going to be a sticking point-

Marcus Johnson:

There's that balance again.

Jacob Bourne:

... in that legislative process

Marcus Johnson:

Yeah, balance between innovation and regulation. Let's move from regulation to classifying AAAI. Juliana [inaudible 00:09:46] with NPR notes that the Recording Academy, the organization behind the Grammy Awards, just outlined new rules for AI use ahead of next year's competition. First, only human creators are eligible for the music industry's highest honor. Secondly, songs that include elements generated by AI can still be nominated, but there must be proof that a real person meaningfully, meaningfully, whatever that means, contributed to the song.

Two, Harvey Mason Jr, the CEO of the Recording Academy, says "If there's an AI voice singing on the song or AI instrumentation, we'll consider it, but in a song writing-based category. It has to have been written mostly by a human." I really don't know what that means. It seems like pretty ambiguous rules, but Gadjo, is this not an example of how difficult it's going to be to classify what's using AI and to what degree?

Gadjo Sevilla:

Absolutely, especially in industries that are just supported by original content. A classification system will be critical for musicians and creators just to ensure attribution and copyright. I mean, there's already a lot of instances where there are copyright infringements just because songs sound alike. The argument there is they're using the same musical building blocks. With

AI, I think it becomes a more complicated situation because there's so many ways AI can play into the creation of a song, whether it's the lyrics or the music, and it could be, I think, very difficult to really trace the providence of the original content in that respect.

Marcus Johnson:

So the US Copyright Office has issued updated guidance on submitting AI-assisted creative work for copyright consideration. The guidelines, they ain't exactly crystal clear, so I went and looked at them. And so you have to "Describe the authorship that was contributed by a human." For example, an applicant who incorporates AI-generated text into a larger textual work should claim the portions of the textual work that is human-authored, and an applicant who creatively arranges the human and non-human content within a work should fit.

I mean, if I haven't lost you yet, you really hang on for dear life here. And then finally, applicant should not list an AI technology or the company that provided it as an author or co-author simply because they have used it when creating their work. There are the incredibly clear rules.

Jacob Bourne:

Yeah, and it really adds a level of complication for artists as well to really understand what the rules are going to be. Also, just it's going to be complicated to create the full frameworks for these rules.

Marcus Johnson:

Why this all matters, why do we need to pass between human and AI created content? Well, the majority, nearly 90% of consumers want to know if they're interacting with something created by AI according to Qualtrics. So it's something people are going to care about. Where did this come from? What's the source here? That's what Chuck Schumer was saying, to go back to the lead. One of the most important things here is letting people know why AI chose one answer and where that information came from.

Let's move to the final part of our lead, and we're talking about AI systems eating themselves. What happens when AI runs out of human-made content to hoover up? Well, generative AI programs may eventually consume material that was created by other machines with disastrous consequences, writes Mateo, one of the Atlantic, pointing out that eventually these programs will have ingested almost every human made bit of digital material that they're

trained on. They're already being used to enlarge the web with their own machine-made content, which will only continue to proliferate.

Mr. Wong thinks that to develop ever more advanced AI products, big tech might have no choice but to feed its programs, AI-generated contents versus human contents, or just might not be able to sift human fodder from the synthetic, a potentially disastrous change in diet for both the models and the internet researchers. So basically AI is being trained on human stuff. When that runs out, it's going to have to start using stuff it's created itself. And what the hell happens then? Jacob, how big of a problem could AI eating up and learning from other AI content be?

Jacob Bourne:

Yeah, I think it could potentially be a very big problem. One of the things we'll see happen is this kind of proliferation of low quality redundant content, which is already kind of a problem for the internet and search engines, especially the quality piece of that. If the AI model is already contained biases, then they're going to be regurgitating those biases and amplifying the problem. Same with the so-called AI hallucination problem, this kind of accuracy issue. If it's already producing inaccurate content and then your AIs are then trained on that inaccurate data, then it makes the problem exponentially worse.

Now, one outcome I don't think is inevitable here is that this means that this flood of AI and regurgitating data is going to mean that humans are going to lose their influence on the internet. I think we might see actually the opposite happen, and that's because it's really human-generated content that's required to provide a fresh source of data for AI model training. And so it's really actually a good thing for human workers. Now, one caveat to all of this is actually what we're seeing is AI companies are looking to move away from just internet data as a source of training.

For example, Google's DeepMind is building a new AI model called Gemini that it's hoping will actually surpass ChatGPT capabilities. It's kind of giving Gemini a new kind of problem solving skills based on interacting with the physical world as opposed to just being trained on internet data. And how it's doing that is actually, it's fusing robotics and generative AI. So basically, AI can now go out into the world and learn in a similar way that a human does.

Marcus Johnson:

Yeah. So Mr. Wong, really interesting article. I thought that him saying the problem with using AI output to train future AI is that despite the often incredible output from generative tools like ImageMaking and MidJourney, they remain sometimes shockingly dysfunctional as the outputs are filled with biases, falsehoods and absurdities. Elia Shamalov, a machine learning researcher at Oxford University says those mistakes will migrate into future iterations of the programs.

If you imagine this happening repeatedly, you'll amplify errors over time, in a study on this phenomenon, not yet peer reviewed, Shamalov and his co-authors described the conclusion of those amplified errors as model collapse, a generative process whereby overtime models forget almost as if they're growing senile. This does not sound ideal at all. That's what we've got time for the lead. Time for the halftime report.

Gadjo, I'll start with you. What to you is most worth repeating from the first half?

Gadjo Sevilla:

I think the idea that US lawmakers and many other legislative and regulatory bodies are moving fast towards trying to at least begin to understand what legislating AI is going to look like. And now it's become a huge challenge in all parts. But again, there's also this feeling that they need to do it expediently and sort of set the pace because it's still a competition across various regions.

Marcus Johnson:

Jacob?

Jacob Bourne:

Yeah, I mean, I think we're starting to see an initial fate outline for what AI regulation will eventually look like. I think it's going to take some time. But parallel to that is all these new questions that continuously surface, these hard to answer questions about the implications of AI such as what happened, what's going to happen when it feasts on its own data, for example.

Marcus Johnson:

That's what we've got time for, for the first half. Time for the second. Today, in other news, just one story. New features for the Apple Vision Pro Mixed Reality Headset emerge.

Story one, the only story, we're talking about new features for the Apple Vision Pro. Apple recently announced its mixed reality headset, the Apple Vision Pro, with two hour battery life for $3,500 coming to the market early 2024 in the US. No controllers, the device is controlled using your hands, eyes, and voice. During its debut, Apple showed off features like projecting a huge movie screen in front of you or interacting with life-size personal photos or videos digitally popup in front of you.

Now, a few weeks removed from this debut, Apple has said it's working on some new features. The headset, centered around fitness and wellness apps think features that track breathing during yoga class. The device might also get full body 3D digital avatars and co-presence capabilities. What that means is using body trackers to help users feel like they're in the same room as other folks also using the mixed reality headset. This new Vision Pro device though does what for the AR headset market, Jacob?

Jacob Bourne:

The Vision Pro is still very much a work in progress. We've seen a preview at WWDC, however it's not launching until 2024. There will probably be some changes made before then. I think what is really important to know is Apple is [inaudible 00:19:34] a lot of resources to a headset at a time when everybody else is focusing on AI. I think that's a great thing for the virtual technology sub-sector. It also shows that the Metaverse isn't dead, but I think one overarching question that looms large here is, when somebody spends $3,500 on a headset, how often are they going to be putting it on?

Marcus Johnson:

I would live in this thing.

Jacob Bourne:

Well, I mean it remains to be seen how long often do they put it on? How long do they wear it? And it could be like what you're saying, Marcus, it's it'll be great for Apple if it's all the time.

Marcus Johnson:

It won't be.

Jacob Bourne:

The results of that are really going to be an important metric in terms of looking at the sales outlook as well as the outlook for a potential next generation model.

Marcus Johnson:

Yeah, time spent. A fascinating point. Yeah, people are not going to wear this thing all the time. I'm just saying that because it costs so much damn money. I couldn't see myself ever wanting to take it off. Gadjo, what do you think?

Gadjo Sevilla:

Yeah, so while there are mixed reality headsets out there, few are being marketed as a spatial computer. Apple actually sees this the same way they saw the first iPhone, which didn't have apps, is quite limited, but it had potential. The thing is, being on a spatial computer it already has an attached developer content and app ecosystem. In Apple's case, it is establishing mixed reality as the platform for the next decade. I do see it as a starting point, and there will be multiple generations of Apple headsets, and we're hoping they would become incrementally more affordable so they could get easily adopted into the mainstream as an alternative to PCs for very specialized applications and the entertainment as well.

Marcus Johnson:

Yeah. Yeah. Hundreds of thousands of apps that already exist on iOS being available through Vision OS. They've got that ecosystem that people can plug into that they're familiar with. If you see someone at dinner with somebody else with one of these on their head, it's me.

That's what we've got time for this episode. Thank you so much, gents, for hanging out. Thank you to Jacob.

Jacob Bourne:

Thank you, Marcus. Thanks Gadjo.

Marcus Johnson:

Thank you to Gadjo.

Gadjo Sevilla:

Likewise, Jacob. It's been a pleasure. Thank you very much.

Marcus Johnson:

Yes, indeed. And thank you to Victoria who edits the show, James, who copy edits it, Stuart, who runs the team. And thanks to everyone for listening in. We'll see you tomorrow hopefully with the Behind the Numbers Weekly Listen, an eMarketer podcast made possible by Verisk Marketing Solutions.