The Daily: Amazon's Anthropic AI deal, ChatGPT can now speak and see, and who's producing 10,000 humanoid robots?

Audio

On today's podcast episode, we discuss the significance of the Amazon-Anthropic deal and what's possible now that ChatGPT can talk to you and see things. "In Other News," we talk about whether Bard AI integrations can help Google catch up to the competition and why one company is producing 10,000 humanoid robots. 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

Episode Transcript:

Marcus Johnson:

This episode is made possible by Intuit Mailchimp. Ever heard of a clustomer? It's the result of marketers grouping customers with different behaviors into one big mess. But with

Mailchimp, you can use real-time behavior data to personalize emails for every customer based on their browsing and buying behavior, turning your clustomers into customers. Intuit Mailchimp, the number one email marketing and automations brand. Visit www.mailchimp.com/personalize for more information. Based on competitor brands' publicly-available data on worldwide numbers of customers in 2021, 2022. Availability of features and functionality vary by plan, which are subject to change.

Jacob Bourne:

The takeaway is that AI is becoming more human-like. Just with this simple upgrade, someone could have a vocal conversation with the chatbot, and people might want to turn to it more.

Marcus Johnson:

Hey, gang. It's Tuesday, October 3rd. Gadjo, Jacob, and listeners, welcome to the Behind the Numbers Daily an eMarketer podcast made possible by Intuit Mailchimp. I'm Marcus. Today, I'm joined by two of our folks who are on our connectivity and tech briefing. One of them is based out of New York, our senior analyst Gadjo Sevilla.

Gadjo Sevilla:

Hey, Marcus.

Marcus Johnson:

Hey, fella.

Gadjo Sevilla:

Nice to be here. Hi, Jacob.

Marcus Johnson:

Yes, sir. Thank you for hanging out today. We're also joined by that very person Gadjo was talking to. One of our analysts on that same team based on the west coast in California, our analyst Jacob Bourne.

Jacob Bourne:

Hey, Marcus. Hey, Gadjo. Glad to be here today.

Marcus Johnson:

Hey, chap. So today's fact, do you guys know where the lines of green Matrix code come from?

Jacob Bourne:

Ooh, it's a tough one.

Marcus Johnson:

I'm annoyed with myself for learning this because I feel like it's shattered a bit of the magic that is The Matrix. So here we go.

Jacob Bourne:

So maybe We don't want to know.

Marcus Johnson:

You don't want to know, but I'm going to tell you anyways.

Jacob Bourne:

No, no, but go ahead. You already started. So.

Marcus Johnson:

Simon Whiteley, creator of the green Matrix Code, says that it was actually created from symbols in his wife's sushi cookbook. He says he likes to tell everybody that The Matrix's code is made out of Japanese sushi recipes.

Jacob Bourne:

Lesson, the inspiration can come from anywhere, I guess.

Marcus Johnson:

Yeah. This has systematically ruined The Matrix for me. There must've been some people sat there looking at it going, "Huh, might make that for dinner."

Jacob Bourne:

Hungry for some sushi now, for some mysterious reason.

Gadjo Sevilla:

There are a lot of kanji characters, so that really makes sense when you think about it.

Jacob Bourne:

Yeah, that's true. That's true. Mm-hmm.

Marcus Johnson:

That checks out. Wait, there a lot of what characters?

Gadjo Sevilla:

Kanji.

Marcus Johnson:

I have no idea. What is that? I'm not cool enough for that, Gadjo. What does that mean?

Gadjo Sevilla:

Japanese characters, I think.

Marcus Johnson:

Okay. Is it like manga?

Gadjo Sevilla:

Yeah, but these are the actual-

Marcus Johnson:

I'm reaching.

Gadjo Sevilla:

Yeah.

Marcus Johnson:

You can say, "No. No, Marcus. It's not like that at all. You're wrong."

Jacob Bourne:

We're learning all kinds of new things today.

Marcus Johnson:

Yeah, mainly me. Before I start, Stuart who runs the team has told me to tell you about our November 3rd virtual event called Attention: Trends and Predictions for 2024. Our leading analyst at the company and executives from brands like Pepsi, Colgate, Palm Olive, Kendra Scott will be exploring trends like gen AI, which we'll be talking about today, retail media, and more to help professionals to plan for the year ahead. You can go to insiderintelligence.com/events/summit, insiderintelligence.com/events/summit to register today.

All right, today's real topic, we're talking all about AI. Anthropic and Amazon enter into a partnership, and ChatGPT learns to speak. In today's episode, first in the lead we'll cover Amazon's big AI deal and ChatGPT being able to talk. Then for In Other News, we'll discuss whether Google's Bard AI integrations can help it catch up, and 10,000 humanoid robots being made in Oregon. That's terrifying. We'll talk about the end of the world later. We'll start with the lead.

We've got a smattering of AI-related stories that we want to cover, and we'll start with Amazon Anthropic. So we're talking about Amazon investing 1 billion now, potentially three later, but up to 4 billion in San Francisco-based AI firm Anthropic. This reflects the earlier partnership between Microsoft and OpenAI, and this comes right after Amazon said it would use AI to boost its Alexa voice assistance conversational powers. Chris Vallance and Liv McMahon of the BBC explained that Anthropic has its own ChatGPT rival called Claude available in the US and the UK, which can handle tasks ranging from sophisticated dialogue and creative content generation to complex reasoning and detailed instruction. Anthropic will now be able to draw on Amazon's huge cloud computing power. In turn, Amazon developers can use Claude 2 to create new applications for its customers and enhance existing ones. As part of the deal, Amazon gets a minority ownership of the company and will incorporate Anthropic's technology in various products that it has across its business, including Amazon's Bedrock service for building AI applications. That's what's going on. Gents, your take on this Amazon-Anthropic deal.

Jacob Bourne:

I mean, I think the big takeaway here is that Amazon is trying to catch up with Microsoft and Google in the generative AI race. And it has technical investments going on in the background to do it, but on the foreground there is the image issue. It needs to project an image, and one

way to do that is to distinguish itself in some way. So that's behind the Anthropic partnership, or the Anthropic investment. Anthropic stands out among regenerative AI startups in that it's marketed itself as the responsible, ethical AI option. And what's behind that is Anthropic's AI product development is closely interwoven with its work on AI alignment, which is basically getting advanced AI in line with human interests and values. And it's an area of AI research that is far behind globally, and so it really stands out that way. And Amazon is hoping to leverage that image from Anthropic.

Marcus Johnson:

Mm-hmm.

Gadjo Sevilla:

I think it also helps Anthropic for its plans because they want to raise 5 billion over the next two years, and that's to build Claude-Next. So that's their language model that they say will be 10 times more capable than today's most powerful AI. So this is really a windfall for them. They can now invest all these billions of dollars to accelerate their research and development. So it helps them catch up with the likes of OpenAI, to some extent, and gives them a bit more of a ceiling to develop a competing product.

Jacob Bourne:

To a certain extent is a key point, too, because remember, Microsoft invested over 10 billion in OpenAI. Amazon not throwing quite so much cash at Anthropic could mean that it's looking also for other investment opportunities.

Marcus Johnson:

That was going to be my next question. How much should I read into that potentially 4 billion? Because it seems like a much smaller deal at the face of it. But to your point, there's more nuance to it.

Jacob Bourne:

Mm-hmm.

Marcus Johnson:

Who are you partnering with, who are you going to partner with next matters as well. Okay. Google has also invested in Anthropic, right? Over $300 million.

Gadjo Sevilla:

Yes. 300 million. Yeah.

Jacob Bourne:

That is true. And I think Anthropic having two major cloud providers that it's partnering with gives it access to a larger potential user base, also gives it access to both of those companies' chips for AI training, which as we know there's an industry shortage in AI training chips. So it's a wise move for Anthropic for sure.

Marcus Johnson:

Mm-hmm. So Google invested in Anthropic. They get a 10% stake. According to Reuters, Anthropic is continuing its partnership with Google despite the new Amazon deal, and is still planning to make its tech available via the Google Cloud.

So Anthropic was saying its model is safer, more reliable because it's guided by a set of principles, allowing it to revise responses itself instead of relying on human moderators. Is this pretty unique specifically to Anthropic?

Jacob Bourne:

As far as I've seen for the leading companies, yes. Anthropic has really invested in these safety features. Again, that work towards alignment, making sure that AI doesn't go rogue eventually, is baked into what it's currently doing in product development. In other companies, OpenAI also has an AI alignment team, but it's not as closely interwoven with the product development, I'd argue, as Anthropic.

Marcus Johnson:

So final question on this deal. There's this deal. There's the deal between, as we mentioned, OpenAI and Microsoft. OpenAI has ChatGPT and Microsoft. Earlier in the year, Google has made an investment in Anthropic. You've got Google's Bard coming out. Meta's releasing a product. Helping folks who aren't following this as closely as you guys, by the start of next year, is there going to be a clear leader in this race?

Gadjo Sevilla:

It's hard to say at this point because everyone's currently picking sides, making investments. There's two sides to this story. It's the big tech companies buying up the capabilities that they don't have, and there's also the AI startups deciding where they fit in the scheme of things. Will they just provide their technology and services, or are they going to compete with a big tech?

In the case of OpenAI, we've covered recently that they're now opening up all these features: web search, voice, image search. So it seems that they want to have a slice of that pie, whereas these other companies like Amazon, like Microsoft, they're integrating that into their products. Same thing with Google. So they already have the customer base for which to push AI. So there'll be, I think, a clearer division between AI companies and big tech companies using AI as we head into next year.

Jacob Bourne:

Yeah, and I think OpenAI's really trying to maintain its lead, and I think that's what we're seeing with all these updates and new features. Some things to watch over the course of the fall and early next year, Google coming out with its Gemini model that's supposed to be very powerful, Meta coming out with another powerful open source model early next year. Those are going to be two major releases that could steer the direction of the race.

Marcus Johnson:

Okay, interesting. You guys both quickly touched on features for ChatGPT. Let's talk about some of those. So ChatGPT has already learned to talk. Users of ChatGPT can now ask it questions by speaking to it and also post images, Gadjo something you mentioned, which can be referred to in conversations. So you can upload an image and then refer to that image during the conversation. Folks can choose from five different options, including male and female voices. Everyone who's subscribed to ChatGPT Plus, which costs about 20 bucks a month, will get access to the new talking version over the next two weeks.

So folks who aren't familiar with what this looks like, just to set the stage a bit further, there was a video from the Wall Street Journal by Joanna Stern, but she was speaking with ChatGPT. And I'll give you a little taste of the conversation. She asks ChatGPT, which now has a voice, she says, "Pretend I'm your best friend. Ask me how my day was." And it replies by saying, "Of course, friend. How was your day? Did anything interesting or special happen? I'm

here to listen." She then says, "I shot a video at work," and it replies, "That's cool. Shooting a video can be interesting and challenging. What was this video about? Any memorable moments during the shoot?" She says, "The video is actually about you." And it says, "Really? That's fascinating. I'm curious. What angle or perspective did you take on the topic? Was it about how AI and large language models work, or perhaps a demo of our interaction?" And she finishes by saying, "It was about how we can talk to each other now."

Admittedly, there was about a three to six-second pause after she'd asked the question or said something before it replied. But quite remarkable, gents, particularly in some of the questions it was asking, saying, "What was the video about? Any memorable moments? Did anything special happen today or anything interesting?" ChatGPT can now talk. What was your takeaway? So what?

Jacob Bourne:

Yeah, I mean the takeaway is that AI is becoming more human-like. Just with this simple upgrade, someone could have a vocal conversation with the chatbot, and people might want to turn to it more for information, personal advice, or just conversation, something that might be more awkward with typing. You get to hear the voice. It's more personal. And the conversation flowed rather naturally. I think they could probably make it even more natural, and then we're going to see even more potential use of this feature.

But I think that beyond the commercial aspect of this feature, this is actually an incremental step on the part of OpenAI towards achieving artificial general intelligence. This using different data beyond internet training data, like voice, physical movement, sound, and touch, is going to be really pivotal in terms of training the next generation of AI models and furthering their advancement. So I think that's something to know, is that yes, this has definitely some commercial appeal to it, but there's an underlying motivation for OpenAI, and that's to get more data to train more powerful models.

Gadjo Sevilla:

Yeah, I think the same goes for what Amazon showed with the AI update for Alexa, their smart speaker. So it was a similar demo. Not super seamless. There were pauses in between and some repetitiveness, but it was much more conversational, more engaging. And they're saying that their smart speakers and their smart screens, they now have built-in cameras so they can

try to understand gestures, tone. So it's a more nuanced approach, I think, to this type of interaction with technology.

Jacob Bourne:

Yeah. And I think it also does, in addition to potentially bringing in more users, it also increases risk, too, because you're going to have people that might be attributing things to AI that are not there, like sentience, for example, because of this conversational aspect. And it's a risk that I think companies like OpenAI are showing that they're willing to take.

Marcus Johnson:

Mm-hmm. So it now has a voice, but it also now has eyes to a certain extent. You can upload images and then refer back to those during the conversation. A few examples, and two from the New York Times, one is uploading a video of the inside of your fridge and then the chatbot can give you a list of dishes that you could cook with the ingredients you have. Another one was students uploading an image of a high school math problem that includes words, numbers, and diagrams, and the bot can instantly read the problem and solve it. Good way to learn, a better way to cheat, the article was noting. And then another one from the Wall Street Journal showing a shot of a leaking hose with the prompt, "How do I fix this?" And it returns, Miss Stern was saying, seven steps of how to fix the problem. So it's about AI taking in information in different ways, providing context to the AI, and also with the voice aspect, giving it a more conversational element too.

Jacob Bourne:

Yeah. And I think with the image feature, it's important to know, too, that that represents a very massive technical breakthrough to get that to work. AI has long not been able to understand the contents of images, and so this is definitely, again, a step in that AGI direction.

Marcus Johnson:

Yeah. One distinguishing factor we should point out for folks is people listening might think, "Well, I have an Amazon Echo. I've been talking to Alexa for a while, trying to get her to do things." The difference here is before, the early Alexa, you had to find the right command. Whereas now, when you talk to it, something like this that is ChatGPT generative AI, after you've asked it, it's going to look for the right response in an ocean of information, and then

Page 12

finding it and pulling it to the front of its mind and then giving you that as opposed to you asked the wrong command and it's like, "I don't have that pre-programmed into my brain."

Gadjo Sevilla:

Right, and it's going to be customized to your need from what it knows about you. So it's constantly learning about you, what you buy, what you need, what your interests are. That's basically the whole difference here. So it's a learning algorithm.

Marcus Johnson:

All right, gents. That's all we've got time for, for the lead. Time, of course, for the halftime report. So Gadjo, I'll start with you for a quick takeaway on our first story regarding Amazon and Anthropic's recent deal.

Gadjo Sevilla:

So yeah, a quick takeaway is that investment from Amazon will allow Anthropic to accelerate its investment on its way to creating Claude next generation LLM, which they say it's going to be 10 times more capable than today's most powerful AI. So having money in the bank gives them the ceiling to catch up with OpenAI at this point.

Marcus Johnson:

LLM, large language model. Jacob, you're going to give us a takeaway for ChatGPT now being able to speak.

Jacob Bourne:

Yeah, it's a signal of two important trends going on in the generative AI sector right now. One is this push towards multimodal AI models, which are going to have more commercial appeal as well as be that incremental step towards AGI. The other thing to watch is companies making these tools able to operate more independently. So basically, more autonomous features like prop generation, for example. Those are the two things that are really going to influence the generative AI race going forward.

Marcus Johnson:

Well, that's it for the lead. Time for the second half of the show. Today in other news, can Bard's AI integrations help Google catch up? And why is one company getting ready to

produce 10,000 humanoid robots annually?

Story one. Gadjo, in a recent piece, you question whether Bard AI integrations can help Google catch up in the generative AI race. You explain that Google is releasing the latest version of its Bard AI, a ChatGPT competitor that it is infusing into popular services like Gmail, Drive, Maps, YouTube, Google Flights and Hotels. But Gadjo, the most important sentence in your article is what and why?

Gadjo Sevilla:

So the most important sentence would be Google has the opportunity to make generative AI truly useful to users by merging Bard with apps and services used by millions of customers. And what I mean by this is they already have a captive market of people who use these services without AI. Adding AI helps them in two ways. They gain adoption and at the same time, they can use what they learn from people using Bard on a day-to-day basis to accelerate their development. So as with a lot of Google things, we become the beta testers for the product and they improve it.

Marcus Johnson:

The line that jumped out to me was companies like Google, Salesforce, Roblox, and Accenture are integrating generative AI features into their services. This could increase the potential for adoption while reducing the dependency on standalone services like OpenAI's ChatGPT. I thought this was fascinating because it presents these two futures, one where you use one service for everything, one where you use specific services linked to specific companies. I imagine it's going to be a mashup of those two universes, but which one do you think is going to be the most prominent?

Gadjo Sevilla:

I think for regular users, having it with the tools they already have, that's just an easier sell than having to step out of that box, get a subscription, use it without the known systems that can make these work successfully.

Marcus Johnson:

Right, and the context. Yeah, yeah. Are you going to use a chatbot or ChatGPT by itself in your house for everything? Are you going to use a specific ChatGPT linked to a certain service to further enhance that experience? Fascinating to see.

Story two. Jacob, in a September 19th piece that I had to double take when I read, you write that Agility Robotics is opening a first of its kind humanoid robotics factory in Salem, Oregon, which it will produce 10,000 humanoid robots annually, according to CNBC. Jacob, you note that Agility designed Digit, which is the robot to work flexibly alongside humans in warehouses and factories. And the robotic coworker can use stairs, crouch in small spaces, and carry up to 35 pounds. The most important sentence in your article is what and why?

Jacob Bourne:

So Agility here is winning on the mass production timeline front for humanoid robotics, surpassing Tesla. But we can eventually expect that companies like Google and Startup Sanctuary and others who are working on equipping humanoid robots with advanced general purpose AI could surpass Agility.

And the thing to know why is Agility doing this, well there's actually demand. I mean, the labor shortage in the US has made it difficult for manufacturers to run at capacity. And so humanoid robots, and of course other types of robots as well, are gaining in popularity.

But from a broader marketplace perspective, the thing to know is that degenerative AI tools have applications for robotics too. And we're going to see down the road, companies like Google, especially, are going to achieve breakthroughs in the robotic sector that's going to eventually supercharge consumer robotics more generally beyond vacuum cleaners, for example. So while Agility is at the forefront right now of that manufacturing push, I think AI sector leaders are really going to pull past eventually, especially for consumer devices.

Marcus Johnson:

Mm-hmm. I couldn't help but think of the battle droids from Star Wars. I think they're in the episode-

Gadjo Sevilla:

Right.

Jacob Bourne:

[inaudible 00:20:23].

Marcus Johnson:

... where Jar Jar Binks makes an appearance. And just seeing thousands upon thousands of those just being rolled off the factory floor is chilling. I was going to say there must be a huge demand for it if they're willing to make 10,000 of these.

Jacob Bourne:

Well, I was going to ask you, too, what your double take was about. But yeah, it's that 10,000 number that definitely-

Marcus Johnson:

Yes.

Jacob Bourne:

... is a lot of humanoid robots in particular.

Marcus Johnson:

Yeah.

Jacob Bourne:

Certainly, the robotic factory arm is pretty commonplace at this point.

Marcus Johnson:

Right.

Jacob Bourne:

But we have not seen something like this before. So one thing will be interesting to see, is the price point, how much they're going to charge per robot. It was at 250,000. I think they're going to have to drop that considerably. But how many deliveries they make once it hits the market should be really interesting.

Marcus Johnson:

Gadjo, you had something?

Gadjo Sevilla:

No, I was just going to ask about the price. I was very curious on how much this would cost. So.

Marcus Johnson:

Because you note that you can't rent them, that you can only buy them. So folks are going to have to fork out this much. It's just striking that there are that many. Now, I'm picturing all of those robots in one place, like all in an Amazon warehouse, working together at the same time. If you spread them out over however many companies that are going to buy them, one or two per workplace seems less dystopian.

But to your point, other companies are working on this. Tesla's has got Optimus bot. Robotic dog maker Boston Dynamics has Atlas. Google, Sanctuary working on humanoid robots as well. So more companies looking to get some of these out into the marketplace.

That's all we've got time for, gents. Thank you so much to my guests. Thank you to Gadjo.

Gadjo Sevilla:

Thank you very much.

Marcus Johnson:

Thank you to Jacob.

Jacob Bourne:

Great to be here. Thanks, Marcus.

Marcus Johnson:

And thank you to Todd who's editing today's episode whilst Victoria is away, James who copy edits it, and Stuart who runs the team. Thanks, everyone, to listening into this episode of the Behind the Numbers Daily, an eMarketer podcast made possible by Intuit Mailchimp. You can Tune in tomorrow to hang out with Sarah Lebow on the Re-Imagining Retail show as she speaks with analysts Sky Canaves and Zach Stambaugh all about what makes a great retail store.