

Adam Brotman (00:00)

the headlines in the news are, the major AI companies are talking about a slowdown, but they're not

Talking about is a slowdown of people using AI. They're not talking about a slowdown in their mind, anyways, of compute to go serve you, you know, the ability for you to use AI. They're talking about how fast we're developing better and better models because they're worried that if we get much further in terms of model smartness and capability at the frontier.

That the models could misbehave and cause massive damage in the world in a variety of ways. And that is a pretty widely shared view amongst a lot of people, particularly in the labs themselves. but it's a confusing topic, and it's all everyone wants to talk about.

Rose Kelly (00:48)

Yeah. I mean, and even Sam Altman came and was agreeing with him on that too as well last weekend, correct? Yeah.

Adam Brotman (00:53)

So yeah, so everybody did in a way.

For example, for there was a moment where Sam Altman said, I agree with that post from Dario. And I agree with his suggestions and I agree with his post. And then Elon Musk said, I think he just said the words, I agree, or something like that. He didn't specify. Yeah, and Demis Asavas from Deep

Rose Kelly (01:06)

I mean, when's the last time the three of them agreed on something? Yeah.

Adam Brotman (01:11)

Mind said, I agree. So there was a there was a lot of that going on.

Andy (01:15)

This is AI First with Adam and Andy, the show about what happens when AI stops being a tool and starts becoming part of your workforce. I'm Andy Sack, alongside my co-host Adam Brotman. Each episode we talk with business leaders on the front lines of the agentic era, what they're building, what's actually working, what's failing, and how AI agents are changing the way companies operate. No hype, just real world use cases, lessons from the field, and ideas you can put to work.

Today we have an episode on AI slowdown and what it means for business leaders.

we've invited our esteemed colleague Rose Kelly, who I'd like to hand it over to. Rose, welcome.

Rose Kelly (01:54)

I'm really excited to talk about this. I know, Adam, we have talked about this a lot this week. Andy, you as well. And I know we've all been talking to a lot of business leaders and a lot of our clients about this. So I think it's a really timely conversation, and I'm excited to get into it. But before we get into it, like what it means specifically for business leaders, I think we need to make sure that we're all talking about the same thing. Adam, you have been saying all week that this term, like the slowdown, is being used to describe a bunch of very different things.

can you kind of start at the beginning? what is Dario actually saying and what are some of the other things that are going on right now that might be getting conflated?

Adam Brotman (02:30)

Sure. Well, first of all, when you said Dario, let's make sure we even define what we mean by that. So,

Rose Kelly (02:35)

Exactly.

Adam Brotman (02:35)

yeah, so you've got Dario is Dario Amodei, the co-founder and CEO of Anthropic, which provides Claude. he posted something last weekend from the time we're taping this that got like a lot of attention. But remember that came.

after almost as a culmination of there was an incident which is now referred to as the hugging face incident which is actually an open AI chat GPT training you know incident fiasco whatever you want call it where a swarm of agents as they were being trained did a bunch of stuff that really showed how not in control

These models are particularly at the very edge of like the fastest and best and most smartest models, which is called the frontier. Okay. Which we don't even see all the time as consumers and as business leaders, the models that are at the frontier, they're being used in the labs, being trained for the next release, and that's called the frontier, the frontier of the smartest, best models. There was this hugging face incident where frontier models and agents created a swarm.

Broke into hugging face, even broke into parts of AI. It caused people to go, like, whoa, what's going on here? And then there was a pretty high profile, I guess you call it, like whistleblower-ish kind of thing that happened at Anthropic following that, where one of their researchers at Anthropic quit and said, I actually think there's a pretty material chance that we could all

die if we don't get this under control. that's how serious the safety risk is here. And that got even more attention and controversy. And it kind of culminated in the CEO of Anthropic posting a pretty lengthy and really well thought through post about what he suggested we do here. and what he suggested was that we pace the frontier, okay, that we slow down the frontier.

And what he was saying was, let's do three things. Let's get a bunch of safety oversight inside the buildings of places like Anthropic and OpenAI. Let's get some sort of regulation or coordination, particularly amongst the labs in democratic nations. Like, he's realizing we can't really do that in China necessarily, but like at least in the Western world we could. And let's try to do some kind of international cooperation.

around safety as well. and he mentioned the hugging face incident and he mentioned some of his, researchers sort of talking out about this and how that was a real issue. So that all became like all everyone's talking about, right? so just to tee it up and then we'll turn it over to Andy and back to you, Rose, is that when he talks about pacing the frontier, slowing down the frontier.

That word slowdown is like perfectly ambiguous because it's like you already have a bunch of people that are already scared of AI today or don't like it for a number of reasons. And

Rose Kelly (05:21)

Or misinformation about it too in general.

Adam Brotman (05:23)

yeah, and you know, to each his own in terms of like, some people are like, I don't like this AI thing. they think it's bad for the environment, for data centers, and other issues with data centers. They think it's bad for

the economy and humanity because it's like taking our jobs and kind of frying our brains in some ways. there's all sorts of reasons that people are sort of upset and anxious about AI. And a lot of those reasons are good reasons, but there's also just a lot of that out there. And so when the headlines in the news are, the major AI companies are talking about a slowdown, but they're not

Talking about is a slowdown of people using AI. They're not talking about a slowdown in their mind, anyways, of compute to go serve you, you know, the ability for you to use AI. They're talking about how fast we're developing better and better models because they're worried that if we get much further in terms of model smartness and capability at the frontier.

That the models could misbehave and cause massive damage in the world in a variety of ways. And that is a pretty widely shared view amongst a lot of people, particularly in the labs themselves. but it's a confusing topic, and it's all everyone wants to talk about.

Rose Kelly (06:37)

Yeah. I mean, and even Sam Altman came and was agreeing with him on that too as well last weekend, correct? Yeah.

Adam Brotman (06:42)

So yeah, so everybody did in a way.

For example, for there was a moment where Sam Altman said, I agree with that post from Dario. And I agree with his suggestions and I agree with his post. And then Elon Musk said, I think he just said the words, I agree, or something like that. He

didn't specify. Yeah, and Demis Asavas from Deep

Rose Kelly (06:55)

I mean, when's the last time the three of them agreed on something? Yeah.

Adam Brotman (07:00)

Mind said, I agree. So there was a there was a lot of that going on.

Andy (07:04)

Let me just add to what Adam's saying, because I think for our audience to understand this, it's important, right? it got triggered by the hugging facing. Really, what Adam's talking about is the future speed of development of the large language models. And really like what's been developed to date is already so insanely powerful. and really people are lagging just keeping up with what is available today. So

It's that distinction that we're trying to make for our audience. and a lot of the concerns which were triggered by the hugging face incidents about the future is about the speed of development. which races us towards AGI. So from a business context, I think what we're gonna talk about next is really.

We've already accelerated like it's insane to just think about twenty twenty six when, you know, the agentic era came to be and what's available in the large language model tools today.

Rose Kelly (07:58)

Yeah, I think that's a really good distinction, Andy. And let's say they decided for some reason to stop all development of models in the future. I mean, there is still years and years in terms of the gap of us being able to recognize the power of these models. And Adam, you made a good point earlier in this week about, you know, there's so many more features and capabilities and so much more that these companies could just be rolling out with the existing models.

Adam Brotman (08:21)

Yeah, in Dario's post, he said he's proposing that we slow, like Andy said, that we slow down the pace of improvements in future releases. Okay, that's what he's saying. And he said, and by the way, even if we do that, it'll still feel like we're accelerating, is what he said in his post because

Rose Kelly (08:37)

Yeah.

Adam Brotman (08:38)

of what you just said. this week alone, this week and last week, there's been

An explosion of what are called personal agents, which we can get into in another discussion. But like the harnesses, the tools that these agents have, they're continuing to get better and easier to use. And the frontier models that are available to us today, let's just call it like Astra class and fable class models, if you stopped everything, they're very, very capable. In fact, there's an argument, according to Jensen Wong, the founder and CEO of NVIDIA, that

those frontier models that we have available today could be called AGI. AGI is artificial general intelligence, which means they can do pretty much everything that humans could do at an average level, if not better. But they're not artificial superintelligence. the other thing you gotta understand about this, just to get into the weeds for a second, is that the thing that people are scared of is that this could get to super intelligence when they're misaligned, when they're not.

totally controllable. That combination, that peanut butter and chocolate is bad. Superintelligence, not aligned, bad. Okay. We have right now borderline AGI now with Astra and Fable class models. And to your and Andy's point, Rose. we could stop and we'd have five years of society still trying to catch up and use this stuff anyway, and the harnesses would get better. And so yeah, I'm excited for you to.

Ask Andy and talk to us about like, well, what does that mean for business leaders, no matter what happens with the slowdown or not slowdown?

Rose Kelly (10:02)

Yeah, no, I'm excited, and that's perfect transition because I was just gonna switch gears. I think, you know, we've spent the last two years or so working with leaders, helping them figure out what their AI strategy is, like getting through and working really hard. And I know it's taken people a long time and it's been difficult work, and now you know, you're a CEO, you maybe have finally kind of figured out your, you know, inklings of an AI strategy, you have things going, and then you kind of see these headlines. I mean, Andy.

I'd love to know like as you've been talking to people this week about this, what is the reaction you think that some of these CEOs and business leaders in general are having to these headlines? And, you know, like what would you recommend for them?

Andy (10:42)

I guess I would distinguish the impact and concerns related to AI that a CEO or business leader would have into two categories. One is the existential human risk, which has everything to do with everything, but it's the existential human species fear and risk about what happens to humanity.

of which your business is a very tiny piece. Separate from that, it's like, well, what's gonna happen to my business? And on the first category, it's all about really what Adam was just talking about, which is what Dario and Sam Altman and Elon Musk and their respective

employees are allowed to do and able to do as it relates to AI future development. As it relates to your business

I think our message, which we laid out in AI first and we then continued in our most recent book, Agents Inc., I think my premise is is like you have got to be using these tools and you've got to accelerate the hard work of transforming your existing business into an AI first company. And

There's lots to consider and we've put out a lot of content on Forum3.com, which I invite you to go look at and you should talk to us, engage with us in the conversation, but you should be plowing ahead, having your AI policy in place and acting on that and using it every day in work and automating your workflows.

Rose Kelly (12:03)

so if a CEO was gonna ask you straight up, does this change my AI strategy at all? What would you say? So it sounds like you'd say no. Like if anything, you should just plow ahead.

Andy (12:09)

No, it j it's

it's not just plow ahead, it's like, you know, we've been saying this repeatedly on this podcast, AI first starts with AI first leaders. You need to become an even more informed and demonstrate what it means to be AI first as a business leader. And you should accelerate your own progress towards that as well as your organizations.

Rose Kelly (12:32)

I mean, we always talk about like there's the technology side of it and then there's the people side. And I think that I can imagine, and I know because I have already had conversations this week with leaders who are already being asked by their teams, you know, their teams are concerned. And so you know, it's like they finally got their team on board with AI and feeling better about it. And then this comes back. And so I guess Adam put this to you. Adam, as a CEO, like if you're in the shoes and someone is being

ask these questions from their team like is this safe? Should we still be doing this? Like, why are you forcing us to use AI? How would you recommend that CEOs talk to their teams about this?

Adam Brotman (13:06)

Yeah, you hit the nail in the head, Rose,

So what I'd say is that I really empathize with business leaders because like you said, Rose, it's your job as a CEO usually, but as any leader of a business to have a point of view about AI that understand what the capabilities are.

What the advantages are and how to use it responsibly, right? So all of those things. and then you gotta get your team on the same page to understand that this is a tool if you use it right and

responsibly, is gonna help your company do better, is gonna help your company achieve its goals faster. and you've got a lot of people that are sort of like worried about their jobs and they're worried about how they personally feel about AI, and you already had to overcome this. Then you get all these news stories that are like,

There's gonna be a slowdown, there's existential risk. And so I really empathize with leaders that have to kinda you know, straighten up a little bit even more and understand that like, we can have a conversation about the societal impacts of AI, which are serious. And I mean even besides the safety issues, there's just like job loss issues and other issues that like you're gonna have to talk about with your team. You're gonna have to have your own

Sort of feelings about them and understand how to deal with them. And then here comes this other stuff about like safety risks. But at the end of the day, you know, this tool's not going away. like I'm not trying to say the genie's out of the bottle and sweep it under the rug. This tool's not going away. Hopefully, we can figure out some way to thread the needle on allowing it to continue improving capability

When it's still not 100% safe and aligned, but that's a bigger issue. The tool's not going away. As a business leader, you have to understand if you don't do what Andy said and leap into using it responsibly, get proficient, get trained, get your team and you talking about this, you're gonna be at a disadvantage. And so it is a tougher conversation now though, because of some of these safety issues.

Rose Kelly (14:54)

Yeah, it definitely is. Andy, anything to add on the like the people side of it and like how to talk to your teams and I mean, let's say you're at a dinner party this weekend and a CEO is like, What should I do on Monday morning? What would you recommend?

Andy (15:05)

There's a lot of totally justifiable, correct concerns that would have you proceed cautiously. And I would say ignore all of those justifiable concerns and ignore ROI and go faster. and I've been saying that, we've been saying that. I'm more emphatic about that today

To go back to the beginning of our conversation, even if the LLMs stopped today, the gap between where the tool is today, where AI and agents are today, and what I see in the marketplace in terms of where business leaders and organizations are, there's years of work to transform and digest what's the technological development.

that's already happened and how that changes the way you do business, how you think about business strategically and the decisions that you'd make and getting those faster, better decisions incorporated into your business already. So that's why I say somewhat caution and ROI to the wayside go.

Adam Brotman (16:05)

Now, just to say one thing, because I agree with Andy, but this is an example of how confusing and ambiguous some of the topics are in this area. Andy is not saying to people, don't be cautious about using a model that's confidential, that you've gotten the right enterprise stuff, that you check your work, like Andy's not saying that. He's talking about whatever your sort of ethical

moral and societal issues are that are legitimate, like they concern Andy and I. That, you gotta compartmentalize that as a business leader to use this tool, but you still have to use the tool responsibly. so that part we're not saying throw caution to the wind on, but some of these more broader philosophical and ethical and safety discussions, you have to compartmentalize. And I don't mean to put words in your mouth, Andy,

'Cause there's still, like we said on our last episode that we just got released, in fact, we're saying, look, you gotta be proficient at AI, you gotta hurry up and learn how to use these tools, but use them responsibly, check your work, understand they can make mistakes, understand that you have to be careful about some of the different kind of connectors you use and the configurations you use, particularly in a corporate environment. But you gotta get to that place where you can use them responsibly so you can

Take advantage of the tools or else you're gonna get left behind.

Rose Kelly (17:15)

I think that's really helpful. So, to recap, what I'm hearing is that your recommendations are, stay the course. It's potentially maybe an opportunity to kind of, double down on making sure everyone understands how to be responsible with AI and why AI is so important to the business. so listen to your team, give real honest answers. but the slowing down the pace of development does not mean that you should be slowing down the AI strategy at your company.

Okay, great. Well Andy, I will turn it back to you to wrap things up.

Andy (17:43)

thanks, Rose, for leading us in this conversation. I think this went really well. And I think going forward the audience should expect you to be part of our conversations more frequently. And I think the audience got at least a good overview and definition of what Dario, Sam, and Elon using their first names as if I knew them,

talked about slowing down and what that means in distinguishing the large language model companies slow down from really business decisions related to their business and business transformation. And I think we did a a reasonable job of covering that. So thank you all for listening to AI First with Adam and Andy. for more on how to build and run a company in the agentic era, head to [Forum3.com](https://forum3.com). You'll find case studies.

research briefings, executive summaries, and our book, Agents Inc., How to Rebuild Your Company in the AI Era. We truly believe you can't overinvest in your AI learning. Onward.

