EXHIBIT 127



HIGHLY CONFIDENTIAL - ATTORNEYS' EYES ONLY

Transcript of Daniela Amodei

Date: November 12, 2025
Case: Concord Music Group -v- Anthropic

Planet Depos
Phone: 888.433.3767 | **Email:** transcripts@planetdepos.com
www.planetdepos.com
Michigan #8598 | Nevada #089F | New Mexico #566

1

IN THE UNITED STATES DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA
SAN JOSE DIVISION

CONCORD MUSIC GROUP, INC., ET AL.,
Plaintiffs,
v.
ANTHROPIC PBC,
Defendant.
Case Number 5:24-cv-03811-EKL SVK

HIGHLY CONFIDENTIAL--ATTORNEYS' EYES ONLY
VIDEOCONFERENCED and VIDEO-RECORDED
30(b)(1) DEPOSITION of
DANIELA AMODEI
San Francisco, California 94111
Wednesday, November 12, 2025

Reported Stenographically by:
MARY J. GOFF
CSR No. 13427
WA CSR No. 21030779
Job No. 606060
PAGES 1-372

2

HIGHLY CONFIDENTIAL--ATTORNEYS' EYES ONLY

VIDEOCONFERENCED and VIDEO-RECORDED DEPOSITION of DANIELA AMODEI, Volume I, taken on behalf of Plaintiff Publishers, Oppenheim + Zebrak, LLP, at WilmerHale, 50 California Street, Suite 3600, San Francisco, California 94111, beginning at 9:07 a.m. and ending at 6:11 p.m., on Wednesday, November 12, 2025, before MARY J. GOFF, California Certified Shorthand Reporter No. 13427 and WA CSR No. 21030779.

3

APPEARANCES:

For Plaintiff Publishers
Oppenheim + Zebrak, LLP
BY: NICHOLAS C. HAILEY
TIMOTHY CHUNG
Attorneys at Law
4530 Wisconsin Avenue, NW
5th Floor
Washington, DC 20016
202-480-2999
nick@oandzlaw.com
tchung@oandzlaw.com

For Defendant Anthropic PBC
WilmerHale
BY: SONAL MEHTA
Attorney at Law
2600 El Camino Real
Suite 400
Palo Alto, California 94306
sonal.mehta@wilmerhale.com
628-235-1026

ALSO PRESENT: Natalie Naugle, In-house counsel Anthropic
Videographer: Philip Astor

4

INDEX

WITNESS		EXAMINATION
DANIELA AMODEI		
Volume I		
	ATTORNEY HAILEY	8
	ATTORNEY MEHTA	--

NUMBER	DESCRIPTION	PAGE
Exhibit 165	Plaintiffs' Amended Notice of Deposition to Daniela Amodei Pursuant to Federal Rule of Civil Procedure 30(b)(1)	13
Exhibit 166	Daniela Amodei LinkedIn profile	26
Exhibit 167	All Company Meeting, 6-30-21 Anthropic_0000012254	73
Exhibit 168	To Do on Wednesday 83-page doc Anthropic_00000478271-353	106
Exhibit 169	Anthropic Plan and Outline Anthropic_00000444967-986	120
Exhibit 170	Brainstorm on What Anthropic Might Release Externally This Year Anthropic_0000458427-428	140

5

EXHIBITS CONTINUED: PAGE

Exhibit 171 7-11-21 email, Brainstorm on What Anthropic Might Release Externally This Year Anthropic_0000458477 141

Exhibit 172 Brainstorm on What Anthropic Might Release Externally This Year Anthropic_0000479021-25 144

Exhibit 173 Some Notes About Research/ Product Collaboration, 11-3-23 Anthropic_0000443204-218 156

Exhibit 174 Baa Baa Black Sheep Anthropic_0000570711 187

Exhibit 175 "Wings" by Birdy sung by Anthropic_0000570714 190

Exhibit 176 Some classical works Anthropic_0000570712-13 195

Exhibit 177 Anthropic Plan for 2023 Anthropic_0000216987 Dario Amodei, 12-30-22 209

Exhibit 178 4-13-23 email, Product research tech team off-site Anthropic_0000458486-490 232

Exhibit 179 Spreadsheet CLAUDEbenchmarking.xlsx Anthropic_0000458491xlxs 236

Exhibit 180 Feedback From 12-19-22, Dario Anthropic_0000478255-266 251

6

EXHIBITS CONTINUED: PAGE

Exhibit 181 Conjectural Anthropic Prioritization List, Jared Kaplan, 7-15-2023 Anthropic_0000479241-249 261

Exhibit 182 Hatch T&S Plan Anthropic_0000480768-778 307

Exhibit 183 Hatch Launch Signoff Anthropic_000012437-449 329

Exhibit 184 Hatch Final Exec Review Anthropic_0000466379-81 340

7

San Francisco, California
Wednesday, November 12, 2025
9:07 a.m.

THE VIDEOGRAPHER: Here begins Media Number 1 in the videotaped deposition of Daniela Amodei, in the matter of Concord Music Group, versus Anthropic, in the United States District Court, Northern District of California, San Jose Division, Case Number 24-CV-03811-EKL (SVK).

Today's date is November 12, 2025. The time on the video monitor is 9:07 a.m.

The videographer today Phillip Astor, representing Planet Depos. This video deposition is taking place at WilmerHale Cutler Pickering Hale and Dorr LLP, San Francisco, California.

Would counsel please voice identify themselves and state whom they represent.

ATTORNEY HAILEY: Nick Hailey, with the law firm Oppenheim + Zebrak, on behalf of the Plaintiff Music Publishers.

I'm joined by my colleague, Tim Chung.

ATTORNEY MEHTA: Sonal Mehta, from WilmerHale, on behalf of Anthropic and the witness.

And with me is Natalie Naugle, who is in-house counsel at Anthropic.

8

THE VIDEOGRAPHER: The court reporter today Mary Goff, representing Planet Depos. The witness will now be sworn.

THE COURT STENOGRAPHER: My name is Mary Goff. I am a California Certified Shorthand Reporter, CSR Number 13427. Today's proceedings are being captured by stenographic means.

DANIELA AMODEI,
Being first duly sworn or affirmed to testify to the truth, the whole truth, and nothing but the truth, was examined and testified as follows:

EXAMINATION
BY ATTORNEY HAILEY:

Q Good morning, Ms. Amodei.

A Good morning.

Q My name is Nick Hailey. I'm an attorney representing the Plaintiff Music Publishers in this case.

We met for a few -- the first time just a few moments ago.

First of all, thank you very much for taking the time today to answer my questions.

A You're welcome.

Q Can you go ahead and state your name and address for the record?

Pages 101-104

101

A In what time frame?
Q When the company launched CLAUDE in the spring of 2023.
A I don't -- I don't remember.
Q So let's talk about Anthropic's founding.
I believe you already testified that your Anthropic cofounders are Dario Amodei, Jared Kaplan, Jack Clark, Tom Brown, Sam McCandlish, and Christopher Olah, correct?
A And myself. And Ben Mann joined us like four days after we left, so sometimes we'll colloquially refer to him as an eighth cofounder.
But I don't believe he's on the -- I don't know if it's in legal documents or not. There's, like, a little question mark there, but this -- we usually refer to the seven cofounders, and then Ben joined shortly after.
Q And you and your Anthropic cofounders are all former OpenAI employees, correct?
A That is correct, with the exception of -- Jared Kaplan, I believe, was a consultant at OpenAI when we departed versus a full-time employee, but I'm not 100 percent sure.
You would have to ask him.
Q And you and your Anthropic cofounders all

102

left OpenAI in December 2020 and founded Anthropic in January 2021?
A We started working on it -- you know, the endeavor that would become Anthropic -- after we departed in December of 2020.
And I believe we formally incorporated in February, but I might have the specific date wrong.
Q Whose idea was it to leave OpenAI together and found a new AI company?
A Again, I'm not entirely sure that I can talk about the details of this due to privilege in our departure agreements.
I think what I can say at a high level is that, you know, this group worked closely together at OpenAI, and we felt we had a compelling vision for a new company that we wanted to found together.
Q Who -- who first had the idea to found a new company together?
A I genuinely don't know.
Q So did you discuss with your cofounders, the reasons that you wanted to leave OpenAI?
A In what time frame?
Q Up until the time -- well, let me ask you this: In any time frame have you discussed with your cofounders your reasons for leaving OpenAI?

103

A We have sometimes talked about Anthropic's vision, mission, why we wanted to start the company; what we felt were sort of, you know, important values that we believed, you know, Anthropic should have as a company; how we wanted to lead the company.
Things like that.
Q So I'm asking not about Anthropic, but about your reasons for wanting to leave OpenAI.
Have you discussed that with some or any of your Anthropic cofounders?
A We have definitely talked about what it was like to work at OpenAI, and we have talked about wanting to have -- you know, having a vision of wanting to start a new organization together, certainly.
Q Wanting to do something different from the way things were done at OpenAI?
A I think I would say we felt strongly that we had, you know, a vision for how we wanted things to run at a new organization like Anthropic.
And obviously, we talked about, you know, how we wanted to implement and what we thought that could look like.
And obviously, once we started the

104

company, we talked about how we wanted to run and organize it together.
Q So before you started at Anthropic, did you discuss with your -- the cofounders any concerns that you had about OpenAI?
ATTORNEY MEHTA: Objection to the form.
A I'm not sure that I -- I mean, yes, we talked about things that we wanted to do, you know, at Anthropic, you know, if we were to start a new endeavor; but I'm -- I'm not sure that I am allowed to go into specific details about, you know, our views about OpenAI just due to a potential agreement we may have had with them.
I just -- again, I don't remember exactly what was in the paperwork that we signed when we departed OpenAI, and I just, you know, want to make sure that we're abiding by the agreements there.
Q (BY ATTORNEY HAILEY) So is -- do you think the agreements with OpenAI allow you to discuss those things amongst yourselves, amongst your other cofounders, but just not in a deposition?
ATTORNEY MEHTA: Objection to the form.
And I'm going to caution the witness with respect to maintaining your obligations under any agreements that you have with respect to your prior

Pages 145-148

145

included comments on this document. And you can -- the comments are collected at the end of the document.

Do you see your comments identified in blue?

A Yes, I do.

Q So does this refresh your recollection that you did review Exhibit 170 and provide comments on it at the time?

A It definitely looks like I provided comments for sure. I don't specifically remember reviewing the document at the time, but it definitely looks like I commented on it.

Q Okay. So let's go back to Exhibit 170, please.

A Is that the brainstorm doc?

Q Yes.

A Okay. Great.

Q So let's start on page 1 with the first bullet.

A Yes.

Q Mr. Amodei writes:

[redacted]

146

[redacted]

Did I read that correctly?

A Yes.

[redacted]

Q Let's go to page 2, the last paragraph of the document, please. Keep scrolling. The last paragraph reads:

[redacted]

147

[redacted]

First of all did I read that correctly?

A Yes.

Q So Dario Amodei is already thinking about Anthropic's public brand at this time?

A It sounds like Dario wrote about his thoughts about our public brand, yes.

[redacted]

A Yes, that's what -- that's what it looks like is written here.

[redacted]

A Likely, yes.

[redacted]

That's what he writes, correct?

ATTORNEY MEHTA: Objection to the form.

[redacted]

148

[redacted]

document.

Q (BY ATTORNEY HAILEY) So what he's saying

[redacted]

ATTORNEY MEHTA: Objection --

Q (BY ATTORNEY HAILEY) -- correct?

ATTORNEY MEHTA: -- objection to the form; misstates the document; and asked and answered.

A I believe I already answered that question.

Q (BY ATTORNEY HAILEY) So I'm entitled to answer that's actually responding to the question, so let me just ask you again.

Do you understand that Mr. Amodei is writing in this document that he wants Anthropic's

[redacted]

ATTORNEY MEHTA: Same objections.

A I don't think I have anything to add to my prior answer to your question.

Q (BY ATTORNEY HAILEY) Do you remember what your prior answer was to my question?

A I don't remember the exact wording.

Pages 157-164

157

head, but it looks like some notes that Jared drafted about how research and product teams worked together.

Q And by "Jared," you're referring to Jared Kaplan, correct?

A Yes.

Q The title of the document is "Some Notes About Research/Product Collaboration."

It's dated November 3, 2023, and it in -- and it indicates that it's by Jared Kaplan --

A Yes --

Q -- correct?

A -- correct.

Q And you can see that you had made various comments on this document, correct?

A Yes.

Q AND those comments were from the time this document was drafted, correct?

A Let me check. Yes.

Q Who is Matt Bell?

A Matt Bell is a manager who reports to Jared.

Q Okay. Is he -- is he an AI scientist at Anthropic?

A I'm not completely sure of his background,

158

but I believe that Matt Bell works on the interface between product and research.

I'm not sure, but I think he has a technical background, but maybe not necessarily an AI background, but I would have to check.

Q So do you know what his title is?

A Something like head of product research or something like that, but I don't know. He's -- he's a manager who reports to Jared.

Q He's one of Jared's direct reports?

A I actually don't know if he still reports to Jared today or not. I believe at the time that this document was drafted, he probably would have been reporting to Jared, but I'm not 100 percent certain.

Q Okay. I can see that you're on page 5. Can you just scroll down a little bit?

A Yeah.

Q Do you see the heading on page 5 that says "Matt Bell's Thoughts"?

A I do.

Q So these appear to be redline comments from Matt Bell on this document that Jared Kaplan drafted, correct?

A Yes, that's what it looks like.

159

Q Can you turn to page 6?

Do you see the first full paragraph on page 6 starting with the word "also"?

Sorry. My question is just: Do you see that paragraph?

A Yes, I do see the paragraph.

Q So I'm going to read that paragraph now.

A Sure.

Q That paragraph reads:

> Also, if there's any way to reduce research risk around model performance during pretraining and fine-tuning, we should pursue it.
>
> One of the challenges we face is that the break-neck competition with OpenAI has driven us to release models as soon as we have established their capabilities, but this makes it particularly hard for teams outside research/product research to plan for releases, as no one knows what the model is going to be capable of until fine-tuning is over.
>
> And once fine-tuning is over,

160

there's a mad-rush to get the model out.

Did I read that correctly?

A Yes.

Q And that's the view that Mr. Bell memorialized in this Anthropic planning document, correct?

ATTORNEY MEHTA: Objection to the form.

A Do you mind if I just read a little bit more of this section to familiarize myself with it?

Q (BY ATTORNEY HAILEY) You can read what you need to. I'm just asking you about that language.

A Okay. Yep.

What was your question?

Q So my question is: That's the view that Mr. Bell memorialized in this Anthropic planning document, correct?

ATTORNEY MEHTA: Objection to the form.

A I'm not Matt Bell, and so I don't necessarily know what he meant by these comments. It's also not an area I'm an expert again.

Q (BY ATTORNEY HAILEY) And again, my question is not what he meant. I'm just asking you that -- we just read what he wrote, and I'm just asking you: That's the view that he memorialized in

161

this document, correct?
ATTORNEY MEHTA: Objection to the form.
A It looks like Matt wrote some sort of thoughts -- some brainstorming thoughts that he had at that point in time.
Q (BY ATTORNEY HAILEY) And that's what he's memorializing in this Anthropic planning document, correct?
ATTORNEY MEHTA: Objection to the form.
A I don't know that I would use the word "memorializing."
Anthropic, particularly on the research side, tends to write a lot of brainstorming documents kind of sharing and batting ideas around.
As a very, you know, heavy R&D company, I think there's a lot of documents that are written where people are sort of trying ideas on.
Because I'm not Matt Bell, I don't know exactly what he intended by this.
Q (BY ATTORNEY HAILEY) Sure.
And I'm not asking you right now what Matt Bell intended, but he's -- he's sharing these views with the people who are listed on this document; is that right?
ATTORNEY MEHTA: Objection to the form.

162

A I don't know who the intended audience was for Matt's comments. It's possible he just meant to respond to Jared. I don't know.
Q (BY ATTORNEY HAILEY) Okay. So he was at least sharing this information with Jared?
A Because I'm not Matt Bell, I don't know what his intention was in 2023 in writing this document.
Q Mr. Bell states in this Anthropic document that:
"One of the challenges Anthropic has faced is that the break-neck competition with OpenAI has driven Anthropic to release models as soon as the company has established their capabilities"; is that -- is that right?
A That is what is written in this document.
That's correct.
Q That's -- that's what Mr. Bell wrote in this internal Anthropic document in July of 2023, correct?
A It looks like this was November.
Q Pardon me.
That's what Mr. Bell wrote in this

163

internal Anthropic document in November of 2023, correct?
A It looks like he wrote the thing that you quoted previously, yes.
Q "No one knows what the model is going to be capable of until fine-tuning is over. And once fine-tuning is over, there's a mad-rush to get the model out," that's what Mr. Bell wrote in the internal Anthropic document in November 2023, correct?
ATTORNEY MEHTA: Objection to the form.
A It looks like Matt Bell wrote that in this document, the thing that you quoted.
Q (BY ATTORNEY HAILEY) And you can see that numerous other Anthropic employees commented on this document, correct?
A It looks like people commented on this document, yes.
Q And no one commented on the language we have just been discussing to say that they disagreed with Mr. Bell, correct?
ATTORNEY MEHTA: Objection to the form.
A It looks like this is a 15-page document that was sort of drafted as a brainstorm. And if I'm not mistaken, it looks like there's a lot of

164

individual sentences that people did not comment on, including that one.
Q (BY ATTORNEY HAILEY) And I'm not asking about anything other than the paragraph that we have been reading and the sentences that I have just been reading to you.
And my question is: Nobody commented on that specific paragraph to say that they disagreed with what Mr. Bell was saying?
ATTORNEY MEHTA: Objection to the form.
A There is no comment on that language. That's correct.
ATTORNEY HAILEY: We can take that down.
Q (BY ATTORNEY HAILEY) Ms. Amodei, how are you compensated for your work at Anthropic?
A I receive a salary, and I received equity.
Q Are you compensated in any other ways for your work at Anthropic?
A I don't believe so.
Q Okay. So let's start with salary.
What is your current annual salary at Anthropic?
A I actually don't know off the top of my head.
Q Okay. What is your best recollection of

Pages 193-200

193

"knowledge cutoff."
Do you see that?
A I do.
Q And that's the date that its training data goes up through, correct?
A Let me just read this. That would be my best guess, yes.
Q So CLAUDE is saying that it:
"May not be aware of all remixes or covers of this song from after its training data cutoff,"
correct?
A Correct.
Q So based on that, CLAUDE is aware of that song because the song is included in its training data?
ATTORNEY MEHTA: Objection, misstates -- I'm sorry.
Assumes facts not in evidence; calls for speculation; asked and answered.
A I believe I already answered that question by saying "I don't know."
Q (BY ATTORNEY HAILEY) Okay. I noticed that you asked CLAUDE the same question twice here.
Do you see that?

194

A Yes.
Q Why did you do that?
A I'm not entirely certain, but my best guess is that this was a period of relative, I don't want to say, "instability," but occasionally CLAUDE would have errors, and so sometimes it would accidentally duplicate a question.
So like you would put a question in, answer it in the consumer product, and then sometimes it would just accidentally duplicate it.
It was like an error that we had on the technical stack side.
Q So your testimony is you think that you didn't actually enter the prompt twice? That's an error by CLAUDE?
A I don't remember, so I'm not certain.
But the fact that it's the exact same question and that it's, like, worded exactly the same way makes my suspect that's what it could have been.
Q Do you ever repeat prompts to CLAUDE to see if CLAUDE would give you a different answer or a better answer?
A In general, when I repeat prompts to CLAUDE, if I feel like it didn't quite give the right answer or maybe it -- I thought it might know something else, I'll usually change a little bit of the wording or the substance because it will just be more likely to result in another response.
It's possible I asked the exact same question twice. But given the timing, my guess is that it might have been an error, but I don't know.
(Exhibit 176 was marked for identification and is attached to the transcript.)
Q (BY ATTORNEY HAILEY) I'm showing you what's been marked as Exhibit 176 to your deposition. This is an Anthropic document with the Bates Number Anthropic 570712.
Do you recognize this document?
ATTORNEY MEHTA: Can you put it into the chat? I don't think that has come through yet. Thank you.
A Yes, I recognize this document.
Q (BY ATTORNEY HAILEY) This is a record of another conversation you had with CLAUDE, correct?
A Yes.
Q And I'll represent to you from the metadata from the production that this one is from October 13, 2024?
A Yes.

196

Q And the record reflects that you prompted CLAUDE:
What are some classical works I should listen to that are sophisticated, full of feeling, and heartfelt?
Did I read that correctly?
A Yes.
Q So you're asking CLAUDE to recommend -- recommend musical works to you; is that right?
A Correct.
Q And then CLAUDE responds by recommending a number of different musical works in response to your request?
A Yes.
Q So we have seen at least three different instances -- Exhibits 174, 175, and 176 -- in which you entered prompts into CLAUDE relating to songs or lyrics, and CLAUDE responded to those prompts; is that correct?
A Yes.
Q Are these typical uses of CLAUDE?
A I believe these are typical use case, but probably not the most common use cases.
Q But these -- these are typical use cases?

197

ATTORNEY MEHTA: Objection to the form; asked and answered.

A I believe I already answered that question with my previous testimony.

Q (BY ATTORNEY HAILEY) What's the difference between a "typical use case" and a "common use case"?

A I think I would say if you were to review the CLIO report of how customers primarily use CLAUDE or what the most common use cases of CLAUDE are, I believe "Coding" is Number 1; "Elementary School, Homework Support" is Number 2.

I think there's other things related to "Research and Development, Healthcare, Life Sciences."

I don't believe anything related to this topic is a "top use case."

But obviously, as you can see from my using it, it's also a helpful use case to learn new artists, to get song recommendations, or to create satirical things like "oink, oink, pink pig."

Q And CLIO, to be clear, doesn't drill down into specific granularity as to the prompt and output types, correct?

ATTORNEY MEHTA: Objection to the form.

198

A I'm sorry. I -- I don't think I understand the question.

Could you repeat it?

Q (BY ATTORNEY HAILEY) The CLIO analysis summarizes and generalizes about prompts -- prompt types, correct?

A It will not share like a particular user's prompt. That's correct.

Q But it also doesn't share specific details about what given prompts might seek? It's -- CLIO is summarizing and clustering groups of prompts, correct?

A I think that's a reasonable characteristic of it. I think I would say, you know, CLIO looks across, you know, millions of data points and summarizes, you know, themes of what people are using CLAUDE for.

So for example, writing code looks very different than perhaps doing, you know, content generation for a sales pitch or something like that.

Q But these three use cases in these three exhibits -- 174, 175, 176 -- could certainly fall into larger CLIO clusters, correct?

ATTORNEY MEHTA: Objection to the form.

A I don't know what the specific CLIO clusters are, all of them, off the top of my head. I would have to review the document.

Q (BY ATTORNEY HAILEY) So you don't know one way or the other?

A I don't know exactly what the breakdown of all of the different CLIO clusters are. I would have to review the document.

Q But when you say that "these are not common use cases," you don't know whether they might call into common -- a common use case as identified in a CLIO cluster, correct?

ATTORNEY MEHTA: Objection to the form.

A My recollection of reviewing the CLIO data indicated that this would not fit in there.

But I don't know for sure, and I'm not an expert in CLIO. You would have to talk to probably my colleague, Deep Ganguli.

Q (BY ATTORNEY HAILEY) Other Anthropic employees use CLAUDE in some of the same ways, as we have been discussing with your use?

ATTORNEY MEHTA: Objection to the form; foundation; speculation.

A I don't know how other Anthropic employees use CLAUDE.

Q (BY ATTORNEY HAILEY) You're not aware that

200

other Anthropic employees have used CLAUDE to generate song lyrics?

ATTORNEY MEHTA: Same objections.

A Outside of a conversation -- a privileged conversation I had with counsel, no.

Q (BY ATTORNEY HAILEY) Ms. Amodei, I'm showing you the PowerPoint presentation that we previously marked as Exhibit 168.

Do you recall -- if you scroll to Slide 2, do you recall this is the PowerPoint we discussed earlier with the title slide "Where We Were, Where We're Going"?

A Correct.

Q You recall reviewing this Exhibit 168 earlier, correct?

A I do.

Q And just to reorient ourselves, this is a PowerPoint presentation that was created in connection with an off-site meeting that took place in mid-2023; is that right?

A I believe, although I'm not completely certain, that this is a document that my chief of staff or admin complied as an input to an eventual document that was shared.

Q Okay. And I will tell you that the

Pages 205-208

205

dramatic reading of CLAUDE in which CLAUDE was convinced that our son, Galileo, who was 13 months old at the time, was the CEO of OpenAI, despite many attempts to convince CLAUDE otherwise. Did I read that correctly?

A Yes.

Q So this is referring to CLAUDE hallucinating and stating incorrectly that your son, who was a baby at the time, was the CEO of OpenAI; is that right?

A Yes.

Q And despite many attempts to convince CLAUDE otherwise, you couldn't correct this hallucination, correct?

A I would say we were sort of purposely, for comic effect, trying to convince -- you know, continue the conversation with CLAUDE, but -- but yes.

Q And this may be a humorous example of CLAUDE hallucinating, but you agree that AI hallucinations can be a serious problem, correct?

ATTORNEY MEHTA: Objection to the form.

A I would say that in 2022, it was more

206

common for all AI models on the market to hallucinate.

There are still hallucinations that can happen with AI models, but they are significantly less frequent today than they were in 2022, I would say, probably across the board for the vast majority of -- of top models.

Q (BY ATTORNEY HAILEY) CLAUDE still hallucinates today, though, sometimes?

ATTORNEY MEHTA: Objection to the form.

A It is definitely possible for CLAUDE to hallucinate today, yes.

Q (BY ATTORNEY HAILEY) And hallucinations can pose serious risks and harms?

ATTORNEY MEHTA: Form.

A I don't necessarily think I would endorse that perspective. I think it's important that AI developers, you know, work to reduce the number of hallucinations that happen, and I think Anthropic has actually been a leader in doing that.

And for -- from my understanding for the majority of time that CLAUDE has been on the market, we have had one of the lowest hallucination rates of all AI models.

Q (BY ATTORNEY HAILEY) So it's your

207

testimony that hallucinations do not pose serious risks and harms?

ATTORNEY MEHTA: Objection to the form; asked and answered.

A I -- I believe I already answered in the previous response.

Q (BY ATTORNEY HAILEY) So I asked a different question, so I would ask you to respond to it.

Is it your testimony that hallucinations do not pose serious risks and harms?

ATTORNEY MEHTA: Objection to the form; asked and answered.

A Again, I think I just -- I wouldn't quite put it that way.

I think what I would say is that it's -- AI developers, including and especially Anthropic, have worked hard over the past several years to reduce the number and rate of hallucinations that AI models produce.

And I think my -- my understanding, although I could be wrong, is that Anthropic has actually been a leader in reducing the number of hallucinations that our models produce.

Q (BY ATTORNEY HAILEY) And the reason you're

208

trying to reduce hallucinations is because hallucinations can cause serious risks and harms?

ATTORNEY MEHTA: Objection to the form; asked and answered.

A I don't think I would put it that way.

I think what I would say is that humans in the loop are usually very valuable for AI models at their current stage of development, and it is very helpful for particularly the businesses that we support to be able to rely on a very high level of accuracy from our models.

However, I think it's also just important to consider in context that in many cases CLAUDE hallucinations less than human error.

Q (BY ATTORNEY HAILEY) But when CLAUDE does hallucinate, that's a problem?

ATTORNEY MEHTA: Form. Same objections.

A I think the way that I would put it is that Anthropic always aims to reduce the number of hallucinations that CLAUDE will produce because it's obviously better for our customers, the higher accuracy that CLAUDE has.

Q (BY ATTORNEY HAILEY) So I -- I still don't think I have an answer to my question, so let me just ask it one more time.