Exhibit D



Planet Depos®
We Make It Happen™

HIGHLY CONFIDENTIAL - ATTORNEYS' EYES ONLY

Transcript of Russ Altman, Ph.D., M.D.

Date: February 27, 2026
Case: Concord Music Group -v- Anthropic

Planet Depos
Phone: 888.433.3767
Email: transcripts@planetdepos.com
www.planetdepos.com

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
DEPOSITION of
RUSS ALTMAN, Ph.D., M.D.
Palo Alto, California 94306
Friday, February 27, 2026

Reported Stenographically by:
MARY J. GOFF
CSR No. 13427
WA CSR No. 21030779
Job No. 620592
PAGES 1-309

2

HIGHLY CONFIDENTIAL--ATTORNEYS' EYES ONLY
VIDEOCONFERENCED and VIDEO-RECORDED DEPOSITION of RUSS ALTMAN, Ph.D., M.D., Volume I, taken on behalf of Plaintiff Music Publishers, Oppenheim + Zebrak, LLP, at WilmerHale, 2600 El Camino Real, Suite 400, Palo Alto, California 94306, beginning at 9:05 a.m. and ending at 6:37 p.m., on Friday, February 27, 2026, before MARY J. GOFF, California Certified Shorthand Reporter No. 13427 and WA CSR No. 21030779.

3

APPEARANCES:

For Plaintiff Publishers
Oppenheim + Zebrak, LLP
BY: COREY MILLER
Attorneys at Law
4530 Wisconsin Avenue
5th Floor
Washington, D.C. 20016
202-957-5043
corey@oandzlaw.com

For Defendant Anthropic PBC
WilmerHale
BY: ARI HOLTZBLATT
BECKY MAZ (remote)
Attorney at Law
2100 Pennsylvania Avenue NW
Washington, D.C. 20037
ari.holtzblatt@wilmerhale.com
202-663-6964

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

4

INDEX

WITNESS	EXAMINATION
RUSS ALTMAN, Ph.D., M.D.	
Volume I	
ATTORNEY MILLER	7, 302
ATTORNEY HOLTZBLATT	284, 304

NUMBER	DESCRIPTION	PAGE
Exhibit 291	Expert Report of Russ Altman 12-19-2025	11
Exhibit 292	Expert Rebuttal Report of Russ Altman, 1-23-2026	13
Exhibit 293	Expert Reply Report of Russ Altman, 2-13-2026	14
Exhibit 294	The Adolescence of Technology article, January 20226, Amodei Tab 2	64
Exhibit 295	Spirit of Asilomar EnTreaty 2025.3.1 Tab 3	71
Exhibit 296	Hacker Used Anthropic's Claude to Steal Sensitive Mexican Data 2-25-26, Mercury News, Tab 4	80
Exhibit 297	LLMs in Healthcare and Medical Applications: A Review Maity, Jyoti Saikia article Pages 1-25, Tab 5	97

Pages 13-20

13

A Yes.
Q Okay. I'm going to hand you what's been marked as Plaintiffs' 292.
(Exhibit 292 was marked for identification and is attached to the transcript.)
Q (BY ATTORNEY MILLER) You can set 291 aside. We'll come back to it.
A Thank you.
Q Do you recognize what has been marked as Plaintiffs 292?
A Yes. This is the "Expert Rebuttal Report" that I submitted. And it looks to be a copy of -- of -- and I think you want it to be a faithful copy of what I actually submitted --
Q Okay.
A -- and I'm --
ATTORNEY HOLTZBLATT: You can look through it, if you would like.
A -- would you like me to confirm that it's the reliable copy?
Q (BY ATTORNEY MILLER) Sure.
A So I'm going to do that. I'm looking at the pages. And once again, so far up to page 9 they look to be identical. I'm going to look through this. I'm going to randomly look at page 25 just in

14

a little bit more detail to make sure they match exactly, which they do. Good.
Following my signature, there's an "Appendix of Materials Considered" again. It's a total of 42 pages. So it appears to match my records and my copy, yes. Thank you.
Q Okay. And for the record, what has been marked as Plaintiffs' Exhibit 292 appears to be a copy of the expert report you submitted on January -- I'll get the date right -- January 23, 2026?
A That's correct.
Q Okay. I'm going to hand you what I have marked as Plaintiffs' Exhibit 293.
(Exhibit 293 was marked for identification and is attached to the transcript.)
Q (BY ATTORNEY MILLER) And do you recognize what has been marked as Plaintiffs' Exhibit 293?
A Yes, this appears to be a copy of my "Expert Reply Report" from February 13.
Q Okay.
A And if you'll allow me, I will quickly confirm it matches the copies that I have been given and are working with. Okay. And once again, after my signature is a "Table of Materials Considered,"

15

and I'm -- which does match my copy, so thank you.
Q And for the record, what I have marked as Exhibit 293, is that a copy of the report you submitted in this case on February 13, 2026?
A Yes.
Q Okay. You can put those aside for the moment. We'll come back to them.
THE DEPONENT: If I may say, I'm -- I'm sorry to interrupt. There's a distracting chat going on in the computer in front of me, which --
ATTORNEY HOLTZBLATT: Let's go off the record.
THE VIDEOGRAPHER: We are going off the record at 9:17 a.m.
(A break was taken from 9:17 a.m. to 9:19 a.m.)
THE VIDEOGRAPHER: We're back on the record at 9:19 a.m.
ATTORNEY HOLTZBLATT: And just as a housekeeping matter, I'm going to designate the transcript as "Highly Confidential, Attorneys' Eyes Only."
Q (BY ATTORNEY MILLER) Dr. Altman, you are a professor of bioengineering, genetics, medicine, and biomedical data science; is that correct?

16

A Yes.
Q And you have a courtesy appointment in the Department of Computer Science?
A Yes.
Q And those appointments are at Stanford University?
A Yes.
Q Have you published anything outside the fields of bioengineering, genetics, medicine, biomedical data science, or computer science?
A Do you mean in my academic work?
Q Yes.
A Those should mostly cover all of my publications. Although I must say that papers don't come labeled with respect to their field, and so that would be a judgment call about: Does this fall within genetics? Does this fall within medicine?
Q Have you taught any courses outside the fields of bioengineering, genetics, medicine biomedical data science, or computer science?
A Yes.
Q What courses?
A Biomedical data science at Stanford used to be called "biomedical informatics" and before that it was called "medical information science."

17

So I have courses are still running that used to have different names to the courses because of the changing names of departments and other organizations.
That's what I meant by "Yes."
Q Okay. Other than courses that used to have different names for the reason you just described, have you taught any courses outside the fields of bioengineering, genetics, medicine, biomedical data science, or computer science?
A Not to my recollection, no.
Q You are currently a board-certified physician in internal medicine and clinical informatics; is that correct?
A Yes.
Q When was the last time you treated a patient?
A Right before the pandemic, I had a clinic. Because of the exigencies of the pandemic, the clinic closed.
Q So that would have been 2020?
A In that time frame, yes.
Q Okay. So you have not treated a patient since 2020?
A That is correct.

18

Q Okay. You're not a lawyer, correct?
A No.
Q You're not offering any legal opinions in this case?
A No.
Q You're not offering an opinion on whether Anthropic's conduct constitutes copyright infringement?
A That is correct. I'm not doing that.
Q And you're not offering an opinion on whether Anthropic's conduct is fair use under The Copyright Act?
A I'm not offering an opinion on that.
Q Do you have any expertise in linguistics?
A I would say to the degree that computational impacts some of my work, there's some expertise there, yes.
Q And how does computational linguistics impact your work?
A So for many years ago -- probably at least 20 -- I have been working in an area called natural language processing, NLP, specifically within biomedicine, scientific research, and healthcare.
And that requires some understanding of natural language, because we're processing it, and

19

so that is considered, I believe -- I know is considered a subdiscipline within linguistics.
Q You play the bass guitar, I believe?
A I do.
Q Do you have formal musical training?
A I don't know what "formal" means. Can you clarify "formal"?
Q Have you taken professional lessons?
A No, I have not taken professional lessons.
Q Have you -- have you ever had a recording contract?
A No.
Q Do you have any expertise in musicology?
A I took one music theory class as a senior in college. I don't think that gives me a real credential in musicology, so I'm going to say no.
Q Okay. Are you offering any expert witness in this case regarding the technical operation of Anthropic's AI models?
A Could you expand on what means "the technical operation" so I can answer more accurately?
Q How the testing models work.
A That is -- how models -- how LLM models work is not part of my remit as an expert in this

20

case.
Q You have not examined any of Anthropic's source code, right?
A I have not.
Q You have not examined any of the datasets that Anthropic used to train any of its AI models, correct?
A I have not.
Q Have you ever built a large language model?
A I have -- what do you mean by "built"?
Q Have you ever pretrained a large language model?
A Yes.
Q Which one?
A We are working on a project with the FDA where we are -- we have fine-tuned several open-source models in a -- in a -- in a head-to-head test that, I believe, included LLaMa and Qwen. I know it included Qwen, because Qwen was the most successful fine-tuning that we did.
Q Well, are you -- do you understand the difference between pretraining and fine-tuning?
A There is a legitimate debate about that, and it has become a fuzzier area recently. I do

Pages 29-32

29

HAI. And Google, you could consider, an AI company. I think they would definitely consider themselves an AI company.

And I -- so I'm very aware -- and I believe Amazon Web Services, which you may or may not consider as an AI company, has also given credits to HAI.

I may have benefited from those credits in either that seed project that I mentioned or the -- or the bigger project. And that would be something that I would -- I would ask my students, because they find out about these credits; and they eagerly use them, and they don't always tell me.

Q You're familiar with a company called Personalis?

A Yes.

Q What is Personalis?

A Personalis is a company that I found -- helped find -- found with five -- four cofounders. It -- currently it's listed on NASDAQ. It currently makes the best residual dis -- cancer disease detector assay that -- that I'm aware of. Just to give one or two more sentences. When you have cancer and you -- and you get a treatment and it becomes undetectable, the worry is always the return

30

of the cancer. And Personalis has built a very sensitive blood assay that is specific to each patient that allows the patient to take a blood test on regular intervals.

And with great sensitivity, we'll be able to tell the patient if their cancer seems to have returned because they detect the DNA of the cancer in the circulating blood.

Q Does -- do Personalis's products use AI?

A I think in the R& -- yes.

Q How so?

A I think in the early R&D, when they're doing basic data analysis, they have used machine-learning versions of AI to process their data. I do not think it is part of the product that they actually deliver. I think it was a research-and-development tool.

Q What is your current role at Personalis?

A I'm a scientific advisor to the company. It does -- it is not doing the things we thought it would do when we first started the company. This is a common phenomenon called pivoting, which you probably have heard of.

And we pivoted to an area that is not really my expertise, because I'm not a cancer

31

biologist. And so I give high-level business advice, and I'm an enthusiast and supporter of the executive team so that they feel supported by the founders -- original founders.

Q Have you ever had a day-to-day operational role at Personalis?

A When the company -- yes. When the company was brand new, I took a one-quarter, three-month sabbatical, to get a feeling for what it was like to be in a company.

I have had an entirely academic career, and so this was an opportunity to see what it would be like and to, of course, contribute to the mission at the time. So that was one quarter -- one quarter only, and then I went back to my regular job.

Q And when was that quarter?

A I would have to check. The company, I believe -- I'm going to start over, because I was just messing with my microphone.

Could you repeat the question?

Q Yeah.

When did you take the quarter-length sabbatical in order to participate in the operations of Personalis?

32

A It was within the first year of its founding, and I would have to check the exact dates.

Q Do you remember what year it was, approximately.

A I really don't. I want to say before 2010, but I would really have to check.

Q Other than the one quarter you spent at Personalis, have you ever worked in the day-to-day operations of a company outside academia?

A No.

Q Do you have a current financial interest in Personalis?

A Yes.

Q Are you familiar with an entity called Bayshore Global Management?

A Yes.

Q What is Bayshore Global Management?

A Bayshore Global Management is the family foundation -- philanthropic foundation of Sergey Brin, one of the cofounders of Google.

Q And Mr. Brin remains a co-controlling shareholder of Alphabet, right?

A I do not know that, but I -- I do not know that.

Q Okay. What is the nature of your

Pages 101-108

101

A It is context-dependent as -- as --
THE DEPONENT: Sorry.
THE VIDEOGRAPHER: Sorry. Can you just clip it on your right lapel?
THE DEPONENT: Right lapel?
THE VIDEOGRAPHER: Yes.
THE DEPONENT: I'm now putting my microphone on my right lapel. And now I'm going to have to ask you to repeat the question.
Q (BY ATTORNEY MILLER) What does human oversight verification and auditing require?
A So that is a context-specific question. The answer about appropriate oversight, verification, and auditing -- excuse me.
The appropriate human oversight, verification, and auditing depends on the situation. In that sentence I was referring to myself as a user of these models and tools, and so as I am a user in different roles, I apply different oversight, verification, and auditing approaches.
Q All right. Could you turn to paragraph 54 of the document you have in front of you, please?
A This is paragraph 54 of my -- of my expert reply report.
Q Yes.

102

A Yes, I have it.
Q And in paragraph 54, you state that, "The value of any individual class of training material, including song lyrics" -- oops. That's what it says, but that's not what I meant to be -- withdrawn. Ah. Turn to paragraph 53.
A Okay.
Q I'm sorry.
A I'm now looking at 53.
Q This. And do you see the third sentence in paragraph 53 states that:
"LLMs derive their capabilities from training on and exposure to vast amounts of heterogeneous text pulled from diverse sources, reflecting the full range of human language use, including creative, metaphorical, and nonstandard expressions"?
A Yes.
Q Do you consider song lyrics to be creative, metaphorical, and nonstandard expressions?
A Lyrics certainly can be that, yes.
Q Now I want you to turn to paragraph 54.
A Okay.

103

Q And in paragraph 54, you state that:
"The value of any individual class of training material, including song lyrics, cannot be evaluated in isolation, but must be viewed as part of a broader ecosystem of language data that collectively supports generalization and responsible use."
Do you see that?
A Yes.
Q Is it important that Anthropic include -- let me take -- withdrawn. Let me take a step back.
You understand that the copyrighted works that are at issue in this litigation are song lyrics?
A I -- yes, I have a rough understanding that song lyrics and copyright issues are at stake in this case, yes.
Q Is it important that Anthropic include the -- the song lyrics at issue in this case in Claude's training data?
ATTORNEY HOLTZBLATT: Objection, vague.
A Could you repeat the question?

104

Q (BY ATTORNEY MILLER) Is it important that Anthropic include the song lyrics at issue in this case in Claude's training data?
A So I did not -- as you know, I did not train Claude, and I really can't that question.
In general, they -- it constitutes text, but the strategies by which Anthropic decided to include or not to include things in its training are opaque to me, and I do not know them.
Q Would Claude be less effective if its training data did not include the song lyrics at issue in this case?
A Unless that's been specifically tested, that's an empirical question that I don't know the answer to.
Q Do you agree that lyrics, as text, do not constitute high-quality datasets for general-purpose LLM training, as they lack lexical diversity in factual information?
A That was a mouthful, so I'm going to ask you to repeat the question. I'll listen even more carefully.
Q Do you agree that lyrics, as text, do not constitute high-quality datasets for general-purpose LLM training, as they lack lexical diversity and

105

factual information.
A I don't have an opinion about that.
Q Do you agree that compared to other sources of text within Claude's training corpora, lyrics in general and including Plaintiffs' lyrics, do not constitute a high-quality dataset specifically for general-purpose LLM training?
A I really don't have an opinion on that as well.
Q Do you agree that the effect of Plaintiffs' lyrics on Claude's learned parameters or outputs is extremely likely to be near zero, if present at all?
A That one requires some thought. The corpus -- my understanding, without any expertise or inside knowledge about the training of Claude, is that it is trained on a huge -- the word could even "humungous" corpus, and I don't know the size of the lyric corpus that may or may not have been included in the training of Claude.
It seems to me from common sense only that it is likely to be a tiny, minuscule fraction of the text that -- that was provided to Claude, and therefore, it's quite unclear to me how critical or -- or it -- or what kind of difference it would

106

make. I cannot comment on that, but that's this -- the way I think about your question.
Q Do you agree that lyrics have minimal value to training a general-purpose LLM?
A I think that's an empirical question as well that I don't have access to the data or the experience to know the answer to.
Q Do you agree that with respect to lyrics, any informational utility they provide could be obtained from many alternative sources, so lyrics text specifically may not be imperative for the training of a general-purpose LLM?
A So you used the word "may" and "could," which -- which then changes the nature of my answer, which is: Yes, it may or could be the case.
And it would need to be empirically tested.
Q And a couple of times today you mentioned the CRISPR-GPT model that you have worked on.
Does that model use Claude in any way?
A No. It used mostly the products from OpenAI.
Q Do you know whether the OpenAI products that that CRISPR model used were trained on song lyrics?

107

A I do not know that.
Q For the CRISPR model to work, was it necessary that the training data for the LLMs on which the CRISPR model relies, included song lyrics?
A It is extremely hard to and, in fact, arguably impossible to understand how any particular piece of data that is fed in at the very beginning of the process into a network of billions or trillions of interacting variables, how that affects the output.
So I'm only saying that to say it is very difficult for me to opine on what the exact effect is. I can't say it's there, and I can't say it's not there, in the same way that every experience we all have contributes to our own personal neural network of the world called the brain.
And then connections that are made are not fully understood, but could be quite -- what's the word -- surprising and idiosyncratic.
Q In your report, you describe work being done on FDA -- excuse me -- on using LLMs in connection with FDA bioequivalence?
A Yes.
Q Is that the -- the fine-tuning you were referring to --

108

A Yes.
Q -- earlier?
Is -- does that project use Claude?
A No.
Q That uses Qwen, Q-W-E-N; is that correct?
A Let me -- if I can expand.
There are two pieces to that project. The first piece is the creation of a high-quality benchmark. As you may know, whenever you're building a system, it's actually good practice to predeclare what you expect it to be good at and, like, commit to that and then build the system and see how good you are.
So the first thing we did is we built a benchmark of questions and high-quality answers in collaboration with the FDA, and that's going to be a separate -- that's going to be a paper about this benchmark.
That benchmark can then be used for two things. It can be used for the fine-tuning, and -- but of course, you take some of that and put it aside, don't make it available to the LLM, and use it to test.
So when -- we built the benchmark, and in the benchmark we did use LLMs to help generate

Pages 129-136

129

page 3.
A Okay. Yes. Thank you. That's. That's right.
Q And if you would please turn to page 11.
A Yes.
Q In the "Discussion," I want to direct your attention to the last sentence of the last complete paragraph.
A On page 11?
Q Yes. And do you see that it says:
"We kept our grading scheme lenient because of the amount of hallucination and poor overall performance of the assessed models"?
Do you see that?
A Yes.
Q All right. Can we go back your report, please, to paragraph 57?
A Yes.
Q Beginning on page 23 and continuing on to page 24, you state that:
"Our results increased my confidence in Claude's potential medical or life science

130

applications."
Do you see that?
A Yeah.
Q So these results were all models performed poorly overall, and you had to keep your grading scheme lenient because of the amount of hallucination and poor over -- poor overall performance of the assessed model increased your confidence in Claude's potential for medical and life sciences applications?
ATTORNEY HOLTZBLATT: Objection, mischaracterizes the document.
A Could you just repeat the question?
Q (BY ATTORNEY MILLER) So these results where all models performed poorly overall and you had to keep your grading scheme lenient because of the amount of hallucination and poor overall performance of the assessed models increased your confidence in Claude's potential for medical and life sciences applications?
ATTORNEY HOLTZBLATT: Same objection.
A Yes, it did. And I can understand your question.
This had not been evaluated yet, so the very act of turning a wonder about whether an LLM

131

might be useful in pulling out this information, which I could describe to you, if you're interested, why recalling facts about biological reactions is important.
But we set up an experiment. I -- I discussed earlier the importance of benchmarks whenever you build a system. So a lot of my confidence is actually not about the performance of any individual LLM, which were all kind of uniformly not that great, but the fact that we now have a system for evaluating.
And in fact, the reviews -- it has occurred to me since you last asked me about the reviews, one of the reviews was: These models are changing all the time, and you're giving us this little snapshot, which will be of very limited use as soon as the next version of Claude or the next version of ChatGPT comes out, and so we -- we're working to figure out how to respond to that.
But that's why I didn't heavily rec -- lean on the actual performance of these. And I think the source of my confidence is that we did the experiment. We have numbers. And as we now try to do clever things to improve performance, I have a way to know when I'm getting better. And -- and

132

that comb -- I guess I'm being optimistic. I now have a system.
In engineering we say: You cannot intervene unless you can measure.
And now I can measure performance, which will allow me to intervene.
Q Starting at paragraph 58, you describe seven use cases for Claude in which Claude delivers what you call "measurable benefits in scientific discovery; is that right?
A Yes.
Q And these are the use cases you identified according to the -- in the methodology in Footnote 59 that we discussed earlier?
A That's right.
Q And the first use case beginning at paragraph 60 that you describe is "managing large multimodal datasets"; is that correct?
A Yes.
Q If Claude's training corpus did not include song lyrics, would that have any effect on Claude's ability to assist with managing large multimodal datasets?
A Could you repeat the question?
Q If Claude's training corpus did not

133

include song lyrics, would that have any effect on Claude's ability to assist with managing large multimodal -- multimodal datasets?

A So again, that -- that -- as I have said before, that's an empirical question. You would have to do an experiment.

Even the use of the word "lyrics" in your question could be interpreted multiple ways because there's written lyrics. But in a multimodal setting, it might be learning the lyrics from watching a performance or -- so "it's not tested" is my answer. And "lyrics" would have to be defined more clearly in order -- even to do that experiment.

Q But you're not offering any opinion that Claude's training corpus needed to include song lyrics in -- in either written or any other form in order for Claude to be able to assist with managing large multimodal datasets?

A I said in the report that "a diverse nonhomogeneous corpus is absolutely critical."

The degree to which it contributes to that, it's contributing. The magnitude needs to be tested separately.

Q In paragraph 61, you describe a potential use of LLMs as "interpreting health insurance

134

policies"?

A Yes.

Q Have you ever used an LLM for this purpose?

A As I said at the very bottom, I did have a small project with a large health insurance provider for this problem, but it was pre-LLM.

Q Okay. Are you aware of any published research showing that LLMs can be used to interpret health insurance policies?

A I'm not aware of any literature on that. I'm not sure the insurance policies would publish that, but I'm not aware -- to answer your question, I'm not aware.

Q In paragraphs 60 and 61, you don't cite to any papers from PubMed, BioRxiv, or ArXiv, do you?

A It looks like not.

Q And in fact, in paragraphs 60 and 61, you don't cite anything, do you?

A It appears no.

Q Okay.

A I -- these observations were made from my general impressions from attending conferences, reading the literature. It's -- it falls within -- in this case this falls well within my kind of

135

professional activities, and so this was a summary statement by me based on my expertise, not supported by citations in this case.

Q All right. The second use case or category of use cases you discuss beginning at paragraph 62 is "assisting with the development of research protocols."

Do you see that?

A Yes.

Q Have you used Claude to assist with the development of research protocols?

A The -- CRISPR-GPT, as I said before, was -- was an effort using the OpenAI tools.

Biomni is another tool generated by a colleague of mine. But I know that students in my lab have tested Biomni, and Biomni is at least partially based on Claude.

Q Have you ever used Biomni?

A I have not used Biomni. My knowledge of Biomni, as for many technical low-level tools, is usually based on deep conversations with my students about their experiences and their successes and failures.

Q In paragraph 64, you describe a study in which Claude enabled -- in paragraph 64, on -- let

136

me start over.

In paragraph 64, on -- I want to -- on the top of page 27, you describe a study using Biomni. And in middle of that paragraph, you state:

"Throughout these tests, Claude enabled Biomni, maintained all of the relevant work data, including code, figures, and logs so that it could be later audited and reproduced."

Do you see that?

A Yes.

Q Do you know whether that data was ever, in fact, audited?

A There is a published paper describing this work. It's -- it's the subject of Footnote 72, and that is where I would look.

So the answer to your question is: No, I don't know. And where I would look would be -- I would reread a citation, 72. I have also heard a talk about this work, and I just don't recall if they used these logs to audit and reproduce.

Q And I take it then that you don't know what results, if any, an audit of that data would have revealed?

Pages 141-168

141

not clear to me that I would do a PubMed-type query to Claude. It's a conversational chatbot. It doesn't kind of represent itself as a source of literature.

So I could do that and it might be okay, but I would instruct my students to be a little bit more like clever in their us of the -- to take advantage of what it's good at and take advantage of what the other resources are good at.

Q In order for Claude to answer a question -- "Who is doing research in my field on this topic -- on a particular topic?" -- does Claude need to have access to the underlying research papers?

A It almost certainly has access to PubMed, as -- as I do. It also access to the websites of the investigators -- or most academics have a little website with summaries of their research projects and their -- and their lab members and what not.

So I -- I think that it would help do it better job if it had some knowledge of the publications that are associated with those reviewers.

But that's not the only data source. It's one of the valuable ones.

Q If Claude's training corpus did not

142

include song lyrics, would that have any effect on Claude's ability to assist with scientific literature reviews?

A I don't know that that's been tested. It does provide diverse and nonhomogeneous text that gives it a better understanding, however incremental, of -- of the use of language and how to interpret language -- sentences and text.

Q And similarly, with the previous use case we were discussing, assisting with the development of research protocols, if Claude's training corpus did not include song lyrics, would that have any effect on Claude's ability to assist with the development of research protocols?

A I don't think it's been tested, and so I don't know, except through the fact that it's a nonhomogeneous -- it is a nonhomogeneous and diverse dataset, which in general, to my understanding, improves the performance of LLMs.

Q In paragraph 66, you discuss FutureHouse; is that right?

A Yes. That's correct.

Q And you have in paragraph -- in the first sentence -- withdrawn.

Paragraph 66 has four footnotes with

143

citations? Do --

A Yes.

Q -- you see that?

And the first one is Footnote 75, correct?

A Yes.

Q And it points -- Footnote 75 cites three sources; is that right?

A Let me just read through it. Give me a moment, please.

I have read the footnote. And it may be that it refers to four, but it's at least three. There's one at the very bottom.

Q Okay. You're correct. I missed one, but it doesn't change -- so let's go through them.

The first one is an article from Anthropic's website; is that correct?

A Yes.

Q The second is FutureHouse's website, right?

A Yes.

Q The third is an article from TechCrunch?

A That's right.

Q And then the fourth is an article from FutureHouse's website, correct?

A Yeah. That's right.

144

Q And then Footnote 76 cites an article from Anthropic's website, right?

A Yeah, it's actually the same -- it's the first -- it's the first citation in 75. I think it's the same one.

Q And Footnote 77 cites that same article from Anthropic's website?

A Yeah. These are just different little quotes that we pulled out.

Q And Footnote 78, which is on the next page, cites the same article from FutureHouse's website that was cited in Footnote 75?

A That's right.

Q Okay. In the second sentence of paragraph 66, you say that:

"In reviewing FutureHouse's work, I found they developed four specialized agents to address different aspects of workflow," and then you identified the four specialized agents --

A Yes.

Q -- right?

Did you actually review any of FutureHouse's work?

145

A I did go to the website. I have gone to the FutureHouse website where they provide what you could call descriptive promotional material about those tools to confirm that they existed, and I did -- and so that's what I did.
Q Have you ever used FutureHouse's service?
A No.
Q Have you spoken to anyone at FutureHouse?
A It -- I have a vague recollection, which I have not reviewed before today, of having some inquiries from them, because what they do is not so different from what my lab has published on.
But I don't think those conversations went forward. I think it was an indication of interest.
Q In your -- you didn't talk to anyone at FutureHouse --
A Other than this potential -- I'm sorry. I --
Q Let me finish --
A -- yes.
Q -- the question.
You didn't speak to anyone at FutureHouse to learn more about how their service works in order to develop your opinions for this report, did you?
A No. Other than this potential interaction

146

by email, I have not done a deep technical dive in kind of a due diligence sense of those tools yet.
Q And you don't -- in your report, you don't report any specific instances of anyone actually attempting to use FutureHouse to do a literature review, do you?
A That's correct.
Q Okay.
ATTORNEY HOLTZBLATT: If you're going to move on, then maybe we should take the lunch break now.
ATTORNEY MILLER: Sure. We can go off the record.
THE VIDEOGRAPHER: All right. We are going off the record at 12:59 p.m.
(A break was taken from 12:59 p.m. to 1:52 p.m.)
THE VIDEOGRAPHER: We are back on the record 13:52 p.m.
Q (BY ATTORNEY MILLER) Good afternoon, Dr. Altman.
A Hi.
Q Do you understand that you're still under oath?
A Yes.

147

Q Excuse me. All right.
In -- before the break we were discussing the seven use cases you identify in your report as "beneficial applications of Claude."I would like to continue with that.
If you could turn to page 29 of your opening report, please.
A Okay.
Q So the fourth use case you identify is "streamlining regulatory processes"; is that right?
A Yeah. That's correct.
Q And in -- excuse me.
In paragraph 68, you discuss a company called Bluenote?
A That's correct.
Q And you state in the second sentence of that paragraph:
"Through a partnership with Anthropic, Bluenote has reported that Claude-powered agents reduced regulatory document preparation time by 50 to 75 percent and in turn improved analytical throughput in quality control and lab operations."

148

Do you see that?
A Yes.
Q Have you verified the result Bluenote reported that Claude-powered agents reduced regulatory document preparation time by 50 to 75 percent?
A No. I was relying on the reports that I cited in -- in the associated footnotes.
Q Okay. Let's look at Footnote 80 then, which is -- that's the footnote you -- Footnote 80 is the footnote to that sentence, right?
A Yes.
Q And there are two sources identified in Footnote 80, correct?
A I believe so, yes.
Q And the first source is an article from the Anthropic website; is that right?
A Yes, that's how it would have come to our attention.
Q And the second source is an article from Epium, E-P-I-U-M; is that correct?
A That's right.
Q What is Epium?
A I would have to check to confirm this, but my recollection, it's one -- kind of an industry

149

newsletter, but I would need to check that.
ATTORNEY MILLER: I got kicked out of the Zoom. One second. Apologies.
Q (BY ATTORNEY MILLER) And do you see that the Epium article that you cite in Footnote 80 has an Anthropic production number attached to it?
A Could you repeat? Could you repeat the question?
Q Yeah.
Do you see in Footnote 80 where you cite the Epium article --
A Yes.
Q -- that citation begins with an Anthropic production number?
A Yeah. That's right. Yes.
Q And the number is Anthropic 569567?
A Yes.
Q All right. Okay. If you'll to the Zoom, I have just uploaded a document there. Please open that document.
A This is Tab 8?
Q Yes.
(Exhibit 300 was marked for identification and is attached to the transcript.)
Q (BY ATTORNEY MILLER) And we'll mark this

150

document as Plaintiffs' Exhibiter 300, which for the record bears the Anthropic Bates Number Anthropic 569567.
A I have -- I have it.
Q Okay. Do you recognize what's been marked as Plaintiffs' 300?
A I -- I believe it's a "Bluenote Deploys Claude-powered Agents to Streamline Life Science Operations."
It's a -- it looks like it's a capture from a website.
Q Is this a capture from Epium?
A It appears to be. It says May 23, 2025, and then gives an "Epium.com" URL.
Q Is this the Epium article you cited in Footnote 80 of your report?
A I think so, yes.
Q Okay. If you would turn to page 2 --
A Yes.
Q -- do you see at the bottom of the page, the document says "Source"?
A Yes.
Q And under "Source," it indicates "Anthropic.com/customers.Bluenote."
Do you see that?

151

A Yes.
Q That does appear to indicate that the information in this article is sourced from the article on Anthropic's website about Bluenote?
A Yeah, I must say that that source, since it's not numbered or connected to any text on that page, is hard for me to parse in terms of its -- the significance of that source, so I can't -- I can't tell why that URL is there.
Q Do you know whether the -- whether this article in Exhibit 300 was written with AI?
A I do not know, no.
Q Have you ever actually used Bluenote?
A I have not used Bluenote. It looks like it's a company that is competitive with the company Weave that I mentioned earlier.
And I -- I am not the target audience for Bluenote, which I think is drug companies -- drug companies who need to create regulatory documents.
Q In developing your opinions for this case, did you talk to anyone at Bluenote?
A No. I relied on this document and the description, which I found plausible, of the work they were doing.
Q In developing your opinions in this case,

152

did you talk to anyone who has used Bluenote?
A No, not in developing my opinions.
Q Okay. Do you know whether anyone has specifically used Bluenote to streamline regulatory processes?
A I only know that in the document they quote presumed clients of Bluenote, including Guardant Health, Eledon -- Eledon Pharmaceuticals, which are companies that I'm familiar with.
That -- and that is on page -- the big -- bit number is 68 at the bottom left in a paragraph that begins:
"Bluenote has engineered three core categories, and these -- they deliver transformative values to clients such as Guardant Health and Eledon Pharmaceuticals."
Q Okay. Do you know what -- withdrawn.
You haven't talked to anyone at Guardant Health, have you?
A No.
Q And you haven't talked to anyone at Eledon Pharmaceuticals, have you?
A No.
Q So you don't know what Guardant Health

153

used Bluenote to accomplish, do you?
A I don't know that.
Q And you don't know what Eledon Pharmaceuticals used Bluenote to accomplish, do you?
A I do not know precisely what either of them did. I had a sense of it from the beginning of that paragraph, but I did not confirm it. I relied on the contents of the report.
Q Okay. You can -- if Claude's training corpus did not include song lyrics, would that have any effect on Claude's ability to streamline regulatory processes?
A I think that's a -- a matter of empiricism. You would have to check it empirically. I would only say that lyrics would represent diverse and nonhomogeneous data that is the type of data that is required for building a high-quality LLM.
Q And you say that "lyrics represent diverse and nonhomogeneous data," what do you mean?
A So diverse and nonhomogeneous text. You're training an LLM. And essentially an LLM is a very large statistical model of language. And in order to have a good statistical model, you need to have lots and lots -- literally millions, billions of examples of text -- from which it can estimate a

154

probability distribution.
And if you choose -- but it -- just having billions of sentences in the ridiculous extreme, it would be one sentence repeated billions of times -- would lead to terrible LLM.
And there's a sense -- and I think it's known empirically from building LLMs that the more different the text is using words in different contexts, saying different kinds of statements with sentences of different lengths and on topics of a wide variety leads to a more robust statistical model of what words mean and how they work together to create sentences. And -- and that's critical for the LLM.
And so when I say "diverse and nonhomogeneous data or text," I mean fulfilling that criteria of not being entirely repetitive and offering no kind of incremental advantage over just repeating the same sentence millions of times.
Q How are lyrics diverse and nonhomogeneous?
A So diversity obviously is something you don't evaluate on a single item. You look at a corpus. Oh, and when I say "corpus," I mean a collection of text. And that is how you determine diversity.

155

So you look at the corpus -- in the example that I gave, the -- the extreme example, but perhaps useful for our conversation, if I'm repeating the same sentence a million times, that sentence doesn't have diversity or lack of diversity in it. It is the corpus of a million mentions of that sentence that would lead me to the conclusion: This is not diverse.
On the other hand, if I have a collection of data that included lyrics -- and I have never had that -- but if I did, then I would not just say: Is this lyric diverse?
I would say: As I look at the other text, which might include, for example, all of Dr. Altman's published papers, is that diverse?
And then you might be able to come up even with quantitative metrics of diversity that would give you confidence that as you're trying to estimate this probability distribution in the LLM, you're do -- going to get a good probability dist -- distribution that is robust for many situations.
Q So when you say -- excuse me.
When you said that "lyrics are diverse and nonhomogeneous," in what sense do you mean that they are diverse and nonhomogeneous?

156

ATTORNEY HOLTZBLATT: Objection, mischaracterizes prior testimony.
A When I -- when I used the shorthand of "lyrics are diverse and nonhomogeneous," having given you this explanation, now I can say: In the context of lots of other text that is different from those lyrics, they may well represent some useful diversity to improve the quality of the statistical model that the LLM is being trained -- that is being trained.
So that's what I mean by "diverse and nonhomogeneous" compared to the other elements of the corpus.
Q (BY ATTORNEY MILLER) In paragraph 70 in your report, you state:
"When used as a writing assistant, LLMs including Claude, can offer similar benefits to creative and academic writing workflows."
Do you see that?
A Yes.
Q Have you ever used Claude for a creative writing workflow?
A Let me start out by saying that I'm

157

disappointed that I separated "creative" and "academic," because when I write my academic papers, I do consider myself being creative. So I regret that actual exact wording.

But yes, for example, I have written grant proposals, proposals to get research funding, and I sometimes ask Claude for its opinion of some of the things I have written; and it has been, at times, useful.

Q Do you consider grant proposals to be creative writing?

A Absolutely for me they're at the pinnacle, because it's a proposal for what I plan on spending somebody else's money to do for the next period of time.

Q And your citation here to the benefits that Claude can offer is to -- a description about Claude at the Google Play App Store; is that correct?

A Yes.

Q Have you ever used Claude to write fiction?

A Well, as I mentioned in my opening report, I used -- I played with Claude to generate some potential lyrics of a kind of erstwhile song I was

158

writing, so you could con -- you could consider that using Claude for a creative activity.

Q Okay. We'll come back to that.

A Okay.

Q Turning to the next page, which is page 31, excuse me, the fifth use case you identify is "facilitating drug development and discovery"; is that right?

A That's right.

Q And you identify -- starting in paragraph 71, you identify two pharmaceutical companies who are using Claude, correct?

A Yes.

Q And the first one is Pfizer; is that right?

A That's right.

Q And the second one is Novo Nordisk, right?

A Yes.

Q And in Footnote 83, do you see that the document you cite to support your claim about Pfizer comes from the GenAI Gazette?

A Yes.

Q Are you familiar with the GenAI Gazette?

A Only insofar as I used it to -- to document the example.

159

Q All right. If you would please go to your Zoom, I'm uploading the next exhibit there. Please open it.

(Exhibit 301 was marked for identification and is attached to the transcript.)

A I see Tab 9.

Q (BY ATTORNEY MILLER) Yes. I'll mark this as Plaintiffs' Exhibit 301, which for the record bears the Bates Number Anthropic 570067.

A Yes.

Q And do you have Plaintiffs' Exhibit 301 open in front of you?

A I think so.

Q And does it appear to be an article titled "Pfizer Leverages GenAI for Faster Drug Development"?

A Yes.

Q And is this the document -- the article from the GenAI Gazette that you cited in Footnote 83?

A Let me just check. One moment. Yeah, that is one of the ones. I believe the second one I quoted in that footnote.

Q Do you know whether this article was written using AI?

160

A I do not know that, no.

Q Okay. Do you see at the top of page 2, it says "Source, SCNP.com"?

A Yes.

Q Do you know what SCNP.com is?

A No, I am not sure.

Q And then if you turn to page 4 --

A Yes.

Q -- do you see that, again, in the middle of the page it says "Source," and there's what appears to be a link to "SCNP.com"?

A Yes.

Q Okay. All right. Are you familiar with the South China Morning Post?

A I am not familiar with the South China Morning Post.

Q Okay. If you could go to the Zoom, I have uploaded the next exhibit there. And please open it.

(Exhibit 302 was marked for identification and is attached to the transcript.)

A I see Tab 10.

Q (BY ATTORNEY MILLER) Yes. And we'll mark this exhibit as Plaintiffs' Exhibit 302.

A Okay. I see it.

161

Q Does this appear -- does what I have marked as Exhibit 302 appear to be an article titled "Deriving Patient-Centric Innovation in Life Sciences Using Generative AI with Pfizer"?
A Yes.
Q Okay. I will represent -- withdrawn.
Do you know in the previous exhibit, 301, we were looking at on page 4, there's a link to SCNP.com?
A Yes.
Q Do you know where that link goes to?
A I don't know, no.
Q Okay. If you go back to Exhibit 302, I will represent to you that when you click on that link, Exhibit 302 is where the link takes you.
ATTORNEY HOLTZBLATT: Objection --
Q Have you ever --
ATTORNEY NAUGLE: -- failure to establish authenticity.
Q (BY ATTORNEY MILLER) Have you ever seen Exhibit 302 before?
A You are referring to what is also known as Tab 10 in our Zoom --
Q Yeah.
A -- exchange? I believe that I saw Tab --

162

Tab 9, and I don't believe that I clicked through to Tab 10.
Q Okay. Could you turn to page 2 of Exhibit 302, please?
A Which is also Tab 10?
Q Yeah.
A Yeah.
Q It's better if we don't use -- the tab numbers are for my internal tracking, so it's better if we --
A I'm just going to continue to do that because everything is labeled as "Tabs" in front of me, and there are no exhibit numbers in front of me.
Q Fair.
Do you see at the top of page 2, this article states:
"The content of this article has been produced by our advertising partner"?
A Yes.
Q And if you go back to page 1, do you see right under the headline, it says "Paid Post, Amazon Web Services"?
A Yes, I do see that.
Q Does this article appear to be a paid

163

advertisement?
ATTORNEY HOLTZBLATT: Objection, lack of foundation.
A It's very hard for me to tell what it means to be a "paid post from Amazon Web Services," which is not Anthropic and is not South China Morning Post or -- or MING Labs, which is mentioned at the bottom of a previous -- so I really am unable to know what it means to be a paid post, and I cannot dissect the significance of that in this example.
Q (BY ATTORNEY MILLER) Well, this is the -- withdrawn. Withdrawn.
In preparing your opinions in this case, did you speak with anyone at Pfizer?
A I have definitely spoken with people at Pfizer, and I have been a consultant to Pfizer in the past.
But in preparing for this testimony, I don't believe I spoke with Pfizer -- anyone from Pfizer, and I certainly didn't discuss the deposition, which I was under instructions not to discuss.
Q Do you know how Pfizer has implemented an LLM?

164

A I do not -- well, I do not, except insofar as it's described in this article.
Q In preparing your opinions in this case -- or in developing your opinions in this case, did you speak with anyone at Novo Nordisk?
A No.
Q Do you know how Novo -- do you know how Novo Nordisk has integrated Claude into his core research infrastructure?
A Only insofar as it is described in the exhibit that -- or in the footnote article that I just -- that I referred in my report.
I found these reports to be credible because they matched things that I had heard at some panel discussions from other companies, including Amgen, which is a biotech company in which I participated in a panel where my Amgen colleague described very similar -- very similar experiences, but I did not have a -- a citable even news article for that, so I used these instead.
Q Does Amgen use Claude? Do you know?
A I do not know. They made general references in the language at the workshop to their use of LLMs. And they, to my memory, did not specify which LLMs.

165

Q If Claude's training corpus did not include song lyrics, would that have any effect on Claude's ability to facilitate drug development and discovery?

A I think that's primarily empirical question about the impact, which you would have to test.

However, in the context of a larger corpus containing many other different documents, the introduction of lyrics could be considered the introduction of diversity and nonhomogeneity into the corpus, which in general, leads to a more robust LLM.

Q All right. Let's turn to -- the sixth use case you identify starting on page 33 is "augmenting healthcare delivery and patient care"?

A Yes.

Q And you begin by referring -- by -- excuse me. You begin by stating that:

"LLMs have been used for knowledge extraction such as protocol and literature synthesis";

is that right?

A Yes.

Q Is -- the "literature synthesis" you

166

reference here, is this different from the literature review we were discussing earlier?

A Yes.

Q What is literature synthesis, as you used it here?

A Literature synthesis here is -- give me a moment to refresh, although -- I'm -- I'm just going to take a moment.

Literature synthesis is more focused than literature review. This is where you have several papers on the same topic and you want to combine their findings for a kind of -- to generate -- in order to generate a summary of what they say as -- as a whole instead of individual studies.

So you would only synthesize things that were similar and comparable so that you could come up with kind of a consensus statement of what they showed.

Q And what is protocol synthesis?

A Protocol synthesis is a similar idea, except now you're looking at protocols that are meant to do some task.

Like: What is your protocol for treating a patient or making a diagnosis?

And you can imagine that there would be

167

several publications from specialists about how they do it.

And you might want to say: Okay, let's generate a summary of what is the current best practices among these several experts who -- who have published a protocol?

The protocols may not be identically the same. That's the challenge. But it may be that when you extract them out into a synthesis, they share lots of features that give you a good idea of the general best practices of people who are expert at these protocols.

Q If Claude's training corpus did not include song lyrics, would that have any effect on Claude's ability to augment healthcare delivery in patient care?

A I don't know. I think that's an empirical question that one would have to test. I do think that lyrics provide -- in the context of the entire training corpus used to create an LLM, potentially provide diverse and nonhomogeneous compared to the other elements of the training corpus; and therefore, they could contribute to a better LLM.

And when I say "a better LLM," I mean a better LLM that better approximates a statistical --

168

a statistical distribution of how language can and is used.

Q When you say "that's an empirical question that one would have to test," you haven't tested it?

A I have not.

Q And I think earlier you suggested that there may be sort of formal statistical measures of diversity within or among datasets?

A I could --

ATTORNEY HOLTZBLATT: Objection, mischaracterizes prior testimony.

A -- I can imagine statistical measures of diversity. I know that they exist in fields that I work in, very disparate fields, but sometimes those measures are quite general purpose.

Q (BY ATTORNEY MILLER) You have not attempted to calculate any statistical measures of diversity with regard to any of the training corpora used for Claude, have you?

A No. We did use these ideas in our fine-tuning of Qwen that I have discussed in response to other questions because we knew that the diversity was important. And we -- we introduce diversity, and it's possible that my student

Pages 173-176

173

So the seventh use case you identify starting on page 36 is "supporting effective time management" --

A Yes.

Q -- is that correct?

A Yes.

Q If Claude's training corpus did not include song lyrics, would that have any effect on Claude's ability to the support effective time management?

A I don't know. And I consider that an empirical question that would -- could be answered with experiments -- with controlled experiments.

I can only say that LLMs, to the extent that they require a diverse and nonhomogeneous corpus of documents, would benefit from lyrics marginally if those lyrics, compared to the rest of the corpus, were different -- that's what I mean by "diverse" -- and different in different ways, which is what I mean by "nonhomogeneous."

Q I would like to direct your attention to paragraph 84 where you discuss "supervision of patients who are unable to care for themselves such as elders or children"?

A Yes.

174

Q And in this paragraph you don't describe anything beyond prototypes, do you?

A Let me read the paragraph to make sure that that's true.

That's right. I am aware -- I am exposed to ambient intelligence academic work quite a bit. It's an area of interest at Stanford.

And I have -- and many of my colleagues are experts in this area, so I didn't refer to specific systems. I kind of integrated my knowledge -- my preexisting knowledge of the kinds of things that are happening in ambient intelligence and the sources for my enthusiasm about them.

Q Are you aware of any ambient intelligence systems that have actually progressed to the stage where they're available for use?

A Yes.

Q Which ones?

A The most impactful ambient intelligence systems so far in medicine have been the ones that listen to an exchange between a physician and a patient in the room and basically take notes.

This has been so far the killer app in medicine, and we hope there will be better killer apps. And "killer" is not a good word to use in the context of medicine. I want to acknowledge that. It's a phrase.

But this has been a remarkable early win for LLMs.

175

Q Do you consider yourself an expert on AI safety generally?

A In academics we like to be specific, so to say things are general is very difficult.

I think I am -- I have quite a bit of knowledge about AI safety in the areas within which I work, which is medicine, biotechnology. But -- and scientific research.

And I think it -- it trails off in terms of my familiarity with the full academic literature. It kind of trails off as -- as expected.

So with regard to -- when we're talking about safety and when we're talking about principles and especially if those principles are present and relevant to the areas that I do know, then I feel on solid ground even talking to a physicist or a chemist or some -- something -- or a social scientist outside of my expertise because we're talking at a high level about shared challenges.

But when those same experts go to detailed challenges of safety in their field, that's when I would have to say: No, that is not my expertise.

176

So there is an element of generalizability to my expertise, which is founded in my actual specific expertise in those specific areas that we -- that I mentioned.

Q Do you have any experience working on developing safeguards or guardrails for AI models?

A So I feel that the answer is yes. And let me explain that. When we were building our LLM for the FDA, in that case and in that specific context, the risks are giving a reviewer at the FDA bad information, and so the entire reason that we did a training of a special model that was domain specific and narrow was to try to mitigate.

And so in that -- I know that -- I re -- an FDA reviewer getting bad information might not consider -- be considered a safeguard, but I think -- I can make the case that they are reviewing data that might lead to the approval of a drug that might be taken by millions of Americans or people globally.

And so even in that small example I consider myself building safeguards, and that's why I say "Yes."

Q When you were building the model for use

Pages 205-208

205

responsible deployment of AI models as a general matter?

ATTORNEY HOLTZBLATT: Objection, compound question. Go ahead.

A I'm going to need to read it for -- briefly just to refresh.

Q (BY ATTORNEY MILLER) Okay.

A Thank you. I have re -- now reviewed the "deploying responsibly" section, and so I probably can -- if you could repeat your question.

Q Sure.

In the "Deploying Responsibly" section of your report, are you discussing the responsible deployment of AI models within the context of specific use cases or are you discussing the responsible deployment of AI models generally.

ATTORNEY HOLTZBLATT: Again, objection, compound question.

A I can say that my discussion -- starting in paragraph 46 of any testimony -- right away mentions "improving healthcare delivery and medical research," so to that extent, I was narrowing my considerations to the area that I know best.

And as I have said in other settings -- in other settings with you in other -- response to

206

other questions, that doesn't mean that the general principles used to address deployment in healthcare and medical research aren't similar to other areas.

But I -- I was speaking primarily from the area that I know best. And even in the ex --

Q And --

A -- I'm sorry. To continue --

Q -- go ahead. Sorry?

A -- and even in the examples that I give in 47 and to some extent 48, yes, I continue to refer to examples within healthcare as -- as the source of my comments.

Q Do you consider yourself an expert on the responsible deployment of AI models?

A I think I have substantial expertise in that deployment within healthcare and scientific venues and perhaps regulatory because of my work with the FDA.

And to the degree that that specific expertise is applicable to other areas, there is some general expertise. But one has to be very careful when claiming expertise, and so my strongest expertise would be in the areas within which I work directly. And then it would be high-level general expertise in other areas.

207

I should say that at a -- it's a little bit more than that. Because as I say at the end of chapter -- of paragraph 46, I am, as a mentioned before, very involved in the Stanford Institute for Human-Centered AI or HAI.

And there, all the applications are coming down the pike, not just healthcare. I tell them about healthcare, but they tell me about other areas.

And so I do test my understanding and expertise and to get some sense for when my expertise is very peculiar to medicine and other times when: Oh, you guys are dealing with this exact same issue.

So that's the kind of dynamic that happens at HAI, which is one of the reasons I love being part of it.

Q When you discuss the responsible deployment of AI models, are you considering that from the point of view of the -- of a person using those models or from the point of view of the company developing those models?

A I think it's fair to say when we think about responsible deployment -- deployment is a thing. And there are many more than the two that

208

you just mentioned. There is the company that built it.

In healthcare, there is the clinician who is providing care. There is the hospital or clinic that they work. There is the patient. There's the patient's family. There's the insurance provider who is paying for the interaction.

And I think all of them have safety interests. They're not always perfectly aligned, but all must be considered when you're looking at a deployment.

It's why deploying things in medicine is, I'll use the word, "fun" to study, because it's so complex that if you can do that well, you probably have some good general principles for exporting to other areas.

Q Can you turn to paragraph 122?

A Okay. I have 122.

Q No. This isn't right. I'm sorry. Turn to paragraph 126.

A Okay.

Q I have the wrong number written down.

A 126. That's the first paragraph of my "Conclusion."

Q Yes. And that paragraph begins:

Pages 241-244

241

with it.
Q In your report you also reference that Northeastern University has implemented the learning mode technology?
A Yes.
Q In preparing your opinions in this case, did you speak to anyone at Northeastern University about their use of Claude?
A No, I did not. I relied on the -- the Anthropic reports of that collaboration.
Q Okay. Are you an ecologist?
A Excuse me?
Q Are you an ecologist?
A Ecologist? No.
Q Are you a zoologist?
A No.
Q Are you a marine biologist?
A No.
Q Are you an environmental scientist?
A Nope.
Q Do you have any professional or academic expertise in public policy development?
A I would say a little bit, yes.
Q What is your experience with public policy development?

242

A So my public policy experience -- education, if you will, or experience is multifold and includes aspects of my service on the FDA advisory committee to the commissioner; aspects of my service to the NIH director as a member of his advisory committee; my -- my work at HAI which, as I have told you, is a broad 360-degree assessment of the public impacts of AI.
And there I think we already discussed briefly my work participating in the AI Index effort, which takes a broad look at the public impact of AI on an annual basis and the AI100 study, which takes a broad look at AI every five years.
In addition, I have some service things. Like, I testified to the U.S. Senate in October about AI in healthcare and -- and drug discovery. And as part of that, visited the individual -- several individual senatorial offices where I mostly met with their staff, explaining and discussing in very concrete terms, the ways that they're considering helping their boss, who is a senator, consider writing either legislation or policy pieces.
I have recently become a fellow -- it's a senior fellow by courtesy at the Hoover Institution, which is a think tank, which you may be familiar with, that thinks about public policy.
And I think it's -- I'm quite sure it's my expertise at the intersection of AI and biotech which is why they invited me there to participate not in basic science discussions, but bringing my basic science expertise to discuss -- to discuss policy issues of interest to them.
Let me just think if there's anything else I wanted to mention. So those, and perhaps a couple that I'm forgetting, has been my kind of on-the-job training for understanding policy issues.
Oh, we have an -- at HAI, we have a "Policy in Healthcare" subcommittee, and I serve on that subcommittee. It's the HAI Health Policy Committee. It's led by my colleague, Michelle Mello, who is a lawyer at Stanford Law School. And it's made up of a very multidisciplinary group, and I serve on that.
And so I think the combination of all these activities give me some expertise in public policy re -- especially related to AI.
Q Have you ever held a public office?
A No. The appointments to advisory committee to the FDA commissioner, and the advisory committee to the NIH were official appointments at those federal -- federal agencies unpaid and only for the term of my service.
Q Have you ever served in local government?
A No.
Q Have you done any research or empirical analysis on how governments and communities use AI?
ATTORNEY HOLTZBLATT: Objection, compound question.
A So you -- I think you asked two questions about my research on how governments use AI, and I haven't done research there.
And there was another part to your question.
Q (BY ATTORNEY MILLER) Sure. I'll break it out.
A Okay.
Q Have you done any research or empirical analysis on how governments use AI?
A No. I have attended many talks about it as part of the HAI, so I listen with great interest. I have not done the work.
Q Have you done any research or empirical analysis on how nongovernmental community organizations use AI?

Pages 249-252

249

Q Are you an artist?
A I am definitely not a professional artist. However, I consider my work playing music and my work taking photographs to be in the direction of art.
Q Are you a graphic designer?
A I have no formal training in graphic design. However, in my publication of scientific manuscripts, there is an unbelievably frequent need to render data and workflows and other things so that we must -- we -- we -- I do a lot of graphic design and -- and critiquing of graphic design of my students.
As part of my job, it is -- it is a -- kind of an accessory requirement of being a professor, but I do not have training in graphic design other than that.
Q When you refer here to "writers and content creators and developers," do you consider yourself a writer of the kind you're referencing here?
A Yeah, so I think there's a legitimate question about, like: What are the creative sectors?
And I take that very expansively. So I do

250

write a lot in my job, and to that -- and it is novel content. It's science. It's not what many people would consider. They're thinking about other things.
But my text is definitely novel to me and it's definitely a professional -- a proficient -- a professional requirement, so I do consider myself to be a content creator in my writing.
Q And then you say:
"Claude has helped content creators amplify the visibility of their content."
Do you see that?
A Yes.
Q And in the next paragraph, you state that:
"I gained exposure to this during an event with HAI," and then you describe a symposium with -- that HAI put on in 2023?
A That's right.
Q Do you have any basis for your opinion about how content creators can use Claude to amplify the visibility of their content, other than your experience at the HAI symposium in 2023?
A Well, the footnote to that Footnote 147 to

251

that claim was a -- a pretty interesting Anthropic collaboration with Pulpit AI where reverends or -- or a -- you know, clergy who have written sermons have used Claude using Pulpit AI. This is not an expected thing. It was unexpected.
Have used this AI to generate other materials that they can distribute to -- to their flock, if you will, to kind of amplify the effect of the written -- and my understanding is they do -- they write their own sermon, but then it is used -- the -- Claude is then used to create numerous kind of offsprings of that. And that's -- that was the example that I cited.
But I agree that there are many other examples, including some of the ones we talked about at this -- at this HAI symposium.
Q The HAI symposium you discussed, according to paragraph 106, occurred on May 24, 2023, right?
A Yes.
Q That was before Claude was -- the first Claude model was released to the general public, right?
A Yes. In those days most of the discussions used ChatGPT as the LLM.
Q So I take it that this symposium did not

252

focus on the use of Claude?
A It -- it referred to LLMs generically. We didn't know Claude would exist, so of course, we didn't refer to Claude.
But I think the idea was to think about not only the models that were extant at the time of the -- at the symposium, but the anticipated future proliferation of such models.
And since the models are in some ways similar -- they're not identical, but they are in some ways similar, we thought that this would be a good general-purpose discussion at the time.
Q Are you a member of the clergy of any religion?
A I -- I had a one-year license to perform a marriage, which may have been called "clergy" and -- at the website that I paid $12 or whatever, but no, I do not consider myself clergy.
Q So I take it then you are not in the target demographic for use of Pulpit AI?
A So I do not think they would target me. Although, I could imagine use -- I give speeches all the time. So depending on how specific to religion Pulpit AI is, I would not rule out that I could find uses for it. But let's be honest. I am not the

Pages 257-260

257

in helping it generate responses to your queries?
A I think if -- if song lyrics were included in Claude, then they contributed to Claude's performance. As I have said, when you train an LLM, you give a huge corpus of text. It builds an elaborate statistical model of language. And then when you query it, it uses that model to create output.
So -- so I do think that all of the training data in a -- in this -- in a distributed sense contributed to the output, and that is always the case modulo alignment.
Q In paragraph 115, you discuss Claude's ability to make recommendations of music in response to queries?
A Yes.
Q In Claude's training corpus did not include song lyrics, would that have any effect on Claude's ability to make music recommendations?
A I don't know. I don't -- because the LLM is a very complex statistical model and because there's multiple inputs, including lyrics, song reviews, album reviews, dis -- critic discussions of songs, indexes of songs that are available for -- for purchase, it is very hard to know the answer to

258

that question, so I don't know.
Q Are you a historian?
A No, I don't think so. No.
Q Do you have any academic training as a historian?
A Only the history classes that were taken as part of my general education for my B.A.
Q You have never taken any graduate-level history classes?
A I have not taken graduate-level history classes.
Q Have you ever published any hysterical research?
A No.
Q Prior to this -- your work as an expert in this case, had you ever conducted any research on the history of technological development?
A Could --
ATTORNEY HOLTZBLATT: Objection, assumes facts not in evidence.
Go ahead.
A -- could -- could you ask it again, please?
Q (BY ATTORNEY MILLER) Yeah.
Prior to your retention as an expert in

259

this case, had you ever conducted any research on the history of technological development?
A The only reason I'm pausing is my field of biomedical informatics is a relatively new field. We are aware of that. It's spun out of computer science.
And so I have written about that. And I -- but I have written it as a participant, and I don't think that would be bona fide historical analysis.
I wouldn't tell my -- I wouldn't tell my colleagues it was a bona fide historical analysis, but it is a history of the technological gen -- development of this subspecialty of computer science which is looking intensely at biological and medical data. So in that sense I have published a little bit.
Q Prior to your retention in this case, had you ever conducted any research on the history of the use of technology in education?
A No, I don't think I have done that.
Q In your rebuttal report you cite several historical examples of the introduction of new technologies that had significant risks associated with them?

260

A Yes. I believe I did, yes.
Q And the -- the first one is the "rollout of electricity"?
A Yes.
Q And the second is "the discovery of x-rays and radiation"?
A Can you give me the paragraph just so I can refresh what I said in those areas?
Q Sure. So beginning at paragraph 25 --
A Thank you.
Q -- and continuing through the next several paragraphs.
A Yes. The choice of those examples was made to make them easily accessible and understandable as examples.
I, of course, have examples of technological rollouts in my fields, but they're very specific. They're not of general interest. And they would not, if you would, hold a reader's attention. And they would be foreign.
And the reader would say: Well, I don't know. Is this true? Is this really how it went?
Relying on historical -- some historical documents and the fact that these are very accessible examples, I thought it would be useful as

Pages 265-268

265

about it in a way that supports this hypothesis that I had.

So that's the nature of the research that I did here, not in -- not a novel in a -- research done by my group, which doesn't do historical research.

Q Are you an economist?

A Definitely not.

Q In your opening report, beginning at paragraph 119, you discuss how "LLMs lower barriers to entry for new business."

Do you see that?

A Yes.

Q Have you done any empirical research on whether LLMs have lowered barriers to entry for new businesses?

A No. These -- these -- this paragraph and the examples that I cite are examples that I'm aware of either from my own consulting or from reading about them that establish some evidence that -- that we're seeing a change in the level of barrier for entry.

Q And empirical research on changes in barriers to entry for new businesses is the kind of thing that an economist would study, right?

266

A Economy -- economists and -- and allied fields, yes.

Q You are also not offering any opinion on the overall macroeconomic effects of generative AI --

A No --

Q -- right?

A -- I'm offering no opinion on that.

Q You haven't studied that?

A Offering no opinion.

Q And you don't have any relevant expertise on that, do you?

A No, I have no opinion to offer for -- on macroeconomics, regular economics, and those issues.

Q Okay. In paragraph 119, you say that you:

"Have personal experience with AI-related business ventures through my work consulting with a -- with a start-up in the pharmaceutical industry, InsightRX, that is using Claude to create agent systems to build models of drug use and toxicity, which is a critical component of drug development"?

267

A Yes.

Q What is the relationship -- nature of your relationship with InsightRX?

A As I said in that sentence, I'm a consulting -- I consult with them as a scientific advisor.

Q When did you begin that consulting relationship?

A Two to three years ago.

Q And what sort of advise do you provide to them?

A I -- one of -- as -- as I mentioned to you, one of my areas of expertise is drug response -- the -- the science of drug response. Why drugs work, why drugs don't work when people take them.

And a lot of my AI research is actually focused on that question. So InsightRX knew about that expertise. And because they were building tools to help the pharmaceutical industry and their -- and their customers, which tend to be hospitals that are using drugs, they initially asked me to join for that reason.

But then we have a shared interest in AI. And you know, they knew that I also had AI

268

credentials. And so in the more recent interactions with them, it's been more about the use of Claude and the creation of agents than it has been about my other expertise.

So over the course of the relationship it's moved from pharmacological-type discussions to AI discussions.

Q What sort of AI discussions do you have with InsightRX?

A The creation of agents. So as I said in this -- they're creating agent systems to build models that -- for how people use drugs and what toxicity they cause.

They're a small company. They're still in the very much at-risk phase of life. Maybe all companies feel that way. But certainly this is a start-up. It's not much older than my relationship with them.

And so they're not only thinking about their current customer base and serving them well; they have to come up with a strategy for growth in the future, and they -- Serge, who is the CEO and he's my main contact, when we meet, was telling me about the exciting work he was doing -- you -- creating agents to help him kind of create the next