Exhibit C



Planet Depos®
We Make It Happen™

HIGHLY CONFIDENTIAL - ATTORNEYS' EYES ONLY

Transcript of Abhishek Nagaraj

Date: March 4, 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
ABHISHEK NAGARAJ, Ph.D.
San Francisco, California 94111
Wednesday, March 4, 2026

Reported Stenographically by:
MARY J. GOFF
CSR No. 13427
WA CSR No. 21030779
Job No. 620602
PAGES 1-291

2

HIGHLY CONFIDENTIAL--ATTORNEYS' EYES ONLY
VIDEOCONFERENCED and VIDEO-RECORDED DEPOSITION of ABHISHEK NAGARAJ, Ph.D., Volume I, taken on behalf of Plaintiff Music Publishers, Oppenheim + Zebrak, LLP, at WilmerHale, 50 California Street, Suite 3600, San Francisco, California 94111, beginning at 9:17 a.m. and ending at 6:52 p.m., on Wednesday, March 4, 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
BRET MATERA (remote DC)
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
CHRISTINA LUO (remote)
LYDIA BAILEY
Attorney at Law
2100 Pennsylvania Avenue NW
Washington, D.C. 20037
ari.holtzblatt@wilmerhale.com
202-663-6964

ALSO PRESENT: Liz Kim, in-house counsel Anthropic; Bram Isgur, Keystone; Shilpi Mukherjee, Keystone
Videographer: Philip Astor

4

INDEX

WITNESS		EXAMINATION
ABHISHEK NAGARAJ, Ph.D.		
Volume I		
	ATTORNEY MILLER	7, 287
	ATTORNEY HOLTZBLATT	284

NUMBER	DESCRIPTION	PAGE
Exhibit 307	12-19-2025 Expert Report of Abhishek Nagaraj, Ph.D. B/W	6
Exhibit 308	1-23-2026 Expert Rebuttal Report of Abhishek Nagaraj, Ph.D. B/W	6
Exhibit 309	2-13-2026 Expert Reply Report of Abhishek Nagaraj, Ph.D. B/W	6
Exhibit 310	12-19-2025 Expert Report of Abhishek Nagaraj, Ph.D. color	41
Exhibit 311	1-23-2026 Expert Rebuttal Report Abhishek Nagaraj, Ph.D. color	41
Exhibit 312	2-13-2026 Expert Reply Report of Abhishek Nagaraj, Ph.D. color	41
Exhibit 313	GPT3, Language Models are Few-Shot Learners Pages 1-75	42

Pages 17-20

17

language models, yes.
Q What model did -- what large language model did you build?
A In research use cases, we were working with other language models and adapting them to our purposes, so I have -- we have worked with a variety of different open-source models such as LLaMA and Qwen.
Q Did you -- you said you have "worked with other large language models and adapted them to your purposes." Have you ever developed a large language model from the ground up from the beginning?
ATTORNEY HOLTZBLATT: Object to the form.
A It's very hard to answer that question because both the word "developed" and "from the beginning" and "large language models" are quite complicated to interpret, and they mean different things to different people.
Q (BY ATTORNEY MILLER) What does "large language model" mean to you?
A In my -- just give me a second. I don't really have an opinion on a technical definition of a "large language model." The only thing I'll say is Claude and models like LLaMA are examples of large language models.

18

Q Okay. Would you say that Qwen, Q-W-E-N, is a large language model?
A Yes.
Q Have you been involved in training anything you would consider to be a large language model?
A Again, the word "training" is quite loaded, and so if your question is "have I worked with these models in a deep technical way," the answer is "Yes."
Q What do you understand model training to be?
A Again, that -- model training means many different things to different people. To me, working with the weights of the model in order to get it to perform certain behaviors is -- is one definition of "training."
Q Okay. And under that definition of "training" you just put forward, "working with the weights of the model in order to get it to perform certain behaviors," have you ever participated in that process with a large language model?
A Research in my lab. And I have written papers on that topic that have adopted -- that have used that technique, yes.

19

Q With which models?
A With Qwen and with LLaMA.
Q Okay. Do you consider Qwen to be a large language model?
ATTORNEY HOLTZBLATT: Object, asked and answered.
A I believe I already answered that question.
Q (BY ATTORNEY MILLER) Do you consider Qwen to be a general-purpose large language model?
A I don't really want to answer that question. I haven't spent time understanding the breadth of Qwen apart from using it in my own research.
Q So you don't know whether it's a general-purpose large language model?
A I do not have an opinion on whether it is a general-purpose large language model. It's not something I have thought about before.
Q How would you define a general-purpose large language model?
A As I said previously, I don't want to offer an opinion. And my job in this case doesn't involve offering an opinion on that question except to say that examples of general-purpose large language models include Claude and -- and LLaMA.

20

Q I'm going to hand you what's been previously marked as Exhibit 307.
A Thank you.
Q Do you recognize what's been marked as Plaintiffs' Exhibit 307?
A I do.
Q What is Exhibit 307?
A Exhibit 307 is seems to be the expert reply that I wrote dated December 19, 2025, that was signed in San Francisco, California.
Q And is Exhibit 307 a true and accurate copy of that report?
A It does seem to me that it is a true and accurate copy of the report, although I would have to fully verify that.
Q I'm going to hand you what has been marked as Plaintiffs' Exhibit 308.
Do you recognize what has been marked as Plaintiffs' Exhibit 308?
A I believe this exhibit is my "Expert Rebuttal Report" assigned on January 23, 2026, in San Francisco, California.
Q And does Exhibit 308 appear to be a true and accurate copy of that report?

Pages 25-28

25

Q Well, in paragraph 21, you refer to "high-performing general-purpose LLMs," correct?
A The phrase that I use here is "when deployed in chat applications, high-performing general-purpose LLMs generate responses to user queries," yes.
Q You could also have a low-performing general-purpose LLM, right?
ATTORNEY HOLTZBLATT: Object to the form.
A Could you define how you see the -- I believe that general-purpose LLMs can vary in terms of their -- their -- their performance, both as a technical matter and -- and in other ways.
Q (BY ATTORNEY MILLER) You said you have used the model Qwen before?
A Give -- or deployed the model Qwen in our own research.
Q Is Qwen built to perform a wide variety of tasks such as understanding, summarizing, predicting, and generating new content?
A I believe so, yes. Yes.
Q So would you then consider Qwen a general-purpose LLM?
A I believe I already answered that question, but yes, I would think about Qwen as a

26

general-purpose LLM.
Q Okay. And how many general-purpose LLMs are there?
A It's very hard to say. I don't believe I have a number that I can give you.
Q In your report you also refer to "single-purpose LLMs," right?
A There is a -- there is a discussion between me and the expert from the Plaintiffs' side on -- on different types of -- different types of LLMs. If you can point me to the paragraph where I refer to "single-purpose LLMs," I can tell you what I mean in that context.
Q How do you distinguish a single-purpose LLM from a general-purpose LLM?
A Give me a second. I believe I address this in my report.
THE DEPONENT: So am I allowed to use these?
ATTORNEY HOLTZBLATT: Yes, you may.
A These are just very hard to flip through. So as I say in my reply report, Point Number 38, that "general-purpose LLMs are those which are generative AI assistants that can help with a broad range of tasks," and so this idea of "helping with a

27

broad range of tasks" is important to my definition of a "general-purpose LLM."
Q (BY ATTORNEY MILLER) How many different tasks does an AI need to be able to help with in order to be a general-purpose LLM?
A Again, I don't believe there is a -- a line in the sand that you can draw on that -- on that idea. But as -- but as I say in my -- in my opening report in Point Number 22 -- sorry -- in Point Number 21, that:
"There are applications that are spanning many industries such as healthcare, education, cybersecurity, software development, and numerous other industries."
So one can understand a "wide variety of applications" to include applicability across different sectors.
Q Is there a minimum number of parameters -- withdrawn.
Do you understand what the concept of "parameters" is in the context of LLMs?
ATTORNEY HOLTZBLATT: Object to form.
A I believe people in the industry use that

28

term in many different ways, so I don't -- I don't personally have a definition or understanding of -- of the word "parameters."
Q (BY ATTORNEY MILLER) You don't understand what parameters means?
ATTORNEY HOLTZBLATT: Object to form; mischaracterizes prior testimony.
A When training large language models, there's a variety of different technical decisions that have to be made. And there's lots of different considerations of, like, the technical architecture and algorithms required for these models. And parameters is one way to specify some of the features of that technical architecture.
Q (BY ATTORNEY MILLER) Is a model described as "having a certain number of parameters," do you understand what that means?
ATTORNEY HOLTZBLATT: Object to the form.
A Again, I would have to look at -- different model companies will publish things in different ways. There's no standardized definition of "parameters." And so I would have to look at the specific details of the technical architecture in order to actually make sense of what anything that is mentioned really means, including the

Pages 33-48

33

Q Are you offering any opinion about the effect of Anthropic's alleged use of lyrics on any other markets?

ATTORNEY HOLTZBLATT: Object to form.

A My stated position, as I just said, is that the three markets that I outlined are all of the relevant markets that are needed to understand any negative impact of -- of Anthropic's alleged use of lyrics on the market for the copyrighted works, and so that analysis is comprehensive as in regards to the -- the assignment that I had, which is to "look at the effect of the use upon the potential market for or the value of the copyrighted work."

Q (BY ATTORNEY MILLER) For the purpose of analyzing the market for all training data for a general-purpose LLM, how do you determine how much data needed is to train a general-purpose LLM?

A I believe that my opening report addresses these points in some detail. For example, in paragraph 35, I say that:

"The bottom line is that the scale and variety of data required to train a general-purpose LLM."

That -- that gives my final conclusion. But let me go to paragraph 6 -- 36, which is that:

34

"I understand that training advanced LLMs requires an enormous number of tokens, which, in turn, requires an enormous amount of data. Without training a dataset of sufficient size, a general-purpose LLM could not draw the statistical inferences it needs to provide answers to general questions."

And so I rely on, for example, Dr. Kaplan's testimony where he says that:

"An LLM like Claude must be shown a vast dataset consisting of hundreds of millions, if not billions, of pieces of content."

And as previously described, training an LLM requires trillions of tokens. For example, the [redacted] And this is the very -- this is the kind of scale that is required to train a general-purpose LLM.

Q Can you train a general-purpose LLM on fewer than 40 trillion tokens?

[redacted]

35

[redacted]**, so the exact number varies of how many tokens it might take to train a general-purpose LLM. But the bottom line is that the scale and size of the data has to be enormous and vast in order for that training to produce a general-purpose LLM.**

Q Is there a minimum quantity of data necessary to train a general-purpose LLM?

A It's hard to say. I don't believe there is any -- again, any specific number that I can point to that you might be looking for beyond which you can train a general-purpose LLM.

The core point is that that data has to be an order of magnitude -- it has -- it has to be vast. So we are talking trillions and trillions of tokens in order to produce something that is both a general-purpose LLM, as well as something that is -- that is useful across a wide variety of tasks.

Q Do you consider Claude 2 to be a general-purpose LLM?

A I have personally never used Claude 2. And I have not studied its release documentation, so I don't want to offer an opinion in that regard.

Q You understand that Claude 2 was one of the models in operation and publicly available at

36

the -- during the relevant time period in this case?

ATTORNEY HOLTZBLATT: Object to form.

A In my reports, I looked at the launch date of the Claude.AI website, and I know a number of different changes happened on the technical infrastructure on the back end in terms of precisely what models are serving outputs to users.

And so in my analysis I look at the commercial operation of Claude, not any specific model or -- or what precisely was going on in the back end.

Q (BY ATTORNEY MILLER) So you can't say whether Claude 2 was a general-purpose LLM or not?

ATTORNEY HOLTZBLATT: Object to form.

A Again, the -- I do not want to offer -- offer an opinion on that on that regard.

Q (BY ATTORNEY MILLER) Well, you're offering an opinion that the market for all -- for all of -- all of the data needed -- needed to train a general-purpose LLM cannot consist, so my question is: Does Claude 2 fall in that category of general-purpose LLMs that you're opining about?

A Claude, as a general class of models, does fall into that category of general-purpose LLMs.

Q Do all of the Claude models fall into that

37

category of general-purpose LLMs?
A Claude 3.5, the Sonnet model, which is the specific one that I -- that I took a look at when I thought about the training and the one that I understand the best, is certainly an example of a general-purpose LLM. There have been follow-on models that have come that I believe are even more capable that also to me seem to be examples of general-purpose LLMs.
Q So what about models of Claude released earlier than Claude 3.5?
ATTORNEY HOLTZBLATT: Objection, asked and answered.
A I believe I already answered that question.
Q (BY ATTORNEY MILLER) So you don't know whether or not models released earlier than Claude 3.5 are general-purpose LLMs or not?
ATTORNEY HOLTZBLATT: Objection, asked and answered.
A I already answered that question.
Q (BY ATTORNEY MILLER) Can you say one way or the other whether Claude 2 is a general-purpose LLM?
ATTORNEY HOLTZBLATT: I think you have

38

asked that now five times.
A I have answered that question.
Q (BY ATTORNEY MILLER) I haven't gotten a clear answer, so I'll ask it again. I can sit here all day until I get a clear answer.
Can you say one way or the other whether Claude 2 is a general-purpose LLM?
ATTORNEY HOLTZBLATT: You can choose to use your time however you would like.
A I have answered that question.
Q (BY ATTORNEY MILLER) When you say that "the market for all of the data needed to train a general-purpose LLM cannot exist," are you saying that the market for all of the data needed to train Claude 2 cannot exist?
ATTORNEY HOLTZBLATT: Objection to form.
A Can you repeat -- repeat the question?
Q (BY ATTORNEY MILLER) When you say that "the market for all of the data needed to train a general-purpose LLM cannot exist," are you saying that the market for all of the data needed to train Claude 2 cannot exist?
A What I'm saying is that the market to train any general-purpose LLM cannot exist. The way that I have defined a general-purpose LLM is an LLM

39

that can do a wide variety of tasks and then one that relies on a large number of tokens in order to train.
So in that sense I do know that Claude 2 required a very large number of tokens to train, and so from that perspective, yes.
I'm also saying that Claude 2 would suffer from the same limitations in terms of building -- building the viability and the existence of a market to gather all of the data to train it, because I believe it took [redacted]
So order of magnitude, it's a very large dataset, and so all of my critiques would apply to that as well.
Q It's your understanding that Claude 2 [redacted]
Q If you may turn to the Zoom, please.
THE DEPONENT: Can we take a break in a few minutes?
ATTORNEY HOLTZBLATT: Sure. Do you want to take a break before we introduce this exhibit?
ATTORNEY MILLER: Sure, that's fine.
ATTORNEY HOLTZBLATT: Okay. Let's go off the record.

40

THE VIDEOGRAPHER: We are going off the record at 10:08 a.m.
(A break was taken from 10:08 a.m. to 10:18 a.m.)
THE VIDEOGRAPHER: We're back on the record at 10:18 a.m.
ATTORNEY HOLTZBLATT: I'll state for the record that I was able to confirm that Anthropic did produce the three expert reports in color, not In Black and White when they were disclosed in the litigation, so I'll renew the request that we mark the color versions of the reports as exhibits.
ATTORNEY MILLER: Yeah, we can go ahead and mark the versions that Dr. Abhishek -- Dr. Nagaraj brought with him.
Q (BY ATTORNEY MILLER) Do you have the stickers handy? Do you want to mark your --
A Can I just mark them with a pen for now?
THE DEPONENT: Can I borrow your pen?
THE COURT STENOGRAPHER: I'll just put the sticker over. That's fine.
A The report is --
ATTORNEY HOLTZBLATT: Yeah, don't -- don't mark them with a pen.
THE DEPONENT: All right.

41

ATTORNEY HOLTZBLATT: So I think that will be Exhibit 310 for the opening report, 311 for the rebuttal; and 312 for the reply?
THE DEPONENT: That's correct. I can give them to you, if you would like.
(Exhibit 310 was marked for identification and is attached to the transcript.)
(Exhibit 311 was marked for identification and is attached to the transcript.)
(Exhibit 312 was marked for identification and is attached to the transcript.)
A To confirm, the expert report is 310; the rebuttal is 311; and the reply is 312.
Q (BY ATTORNEY MILLER) Professor Nagaraj, are you familiar with the model ChatGPT-3?
ATTORNEY HOLTZBLATT: Object to form.
A I believe the -- I'm actually not familiar with that very particular phrasing you have. Because, again, ChatGPT is a service that OpenAI provides that's powered by many different models underneath, including -- I have heard of GPT-3, GPT-3.5, Chat-4o, et cetera.
Q (BY ATTORNEY MILLER) Let me rephrase.
Are you familiar with the model GPT-3?
A I am, in the sense that I have heard of it.
Q Do you consider GPT-3 to be a general-purpose large language model?
A My understanding of GPT-3 is that it is able to perform a wide variety of tasks, and that would be applicable to different industries. In that sense, yes, I do think GPT-3 was a large language model.
Q If you go to the Zoom and look at the chat, you should see the next exhibit there which I have uploaded to you. Please open that document.
(Exhibit 313 was marked for identification and is attached to the transcript.)
Q (BY ATTORNEY MILLER) And for the record, we're not going to put a sticker on it now, but we'll designate this exhibit as Plaintiffs' 313.
A I do have the exhibit.
Q You have Exhibit 313 open in front of you?
A It says "GPT-3.PDF."
Q Yeah. You have that open?
A Yes.
Q Does that appear to be an article titled "Language Models are "Few-Shot-Learners"?
A It is.
Q And the first named author on this article is Tom D. Brown.
Do you see that?
A I do see that.
Q Do you know who Tom Brown is?
A I do not.
Q And the second named author on this article is Benjamin Mann.
Do you see that?
A I do.
Q Do you know what Benjamin Mann is?
A I believe he was one of the people involved in starting Anthropic.
Q Have you ever read this article before?
A I have referred to it before, but I wouldn't say -- I have -- I have referred to it before, yeah.
Q Okay. If you could turn to page 74 of your opening report in front of you.
A Um-hum.
Q And do you see this is the -- at the top of the page it says "Appendix B, List of Materials Considered."
Do you see that?
A I do see that.
Q And this is the list of materials relied upon in forming your -- the opinions expressed in this report, correct?
A Correct.
Q Towards the bottom of page 74, do you see that there's an entry for "Tom B. Brown, et al."?
A I do see that.
Q And for the paper "Language Models are Few-Shot Learners"?
A Correct.
Q And that's the Exhibit 313 we were just looking at?
A It is, yes.
Q Okay. Could you turn back to Exhibit 313, please? I want you to turn to page 8. Are you there?
A Just a second. Yes.
Q And do you see at the top of page 8, there's a "Table 2.1"?
A I do see that.
Q And Table 2.1 lists a series of versions of the GPT-3 model; is that correct?
A The footnote says "Sizes, architectures, and learning hyperparameters of the models which we trained," et cetera, yeah.
Q And then the rest of the caption for

45

Table 2.1 says "All models were trained for a total of 300 billion tokens."
Do you see that?
A I do see that.
Q Does this indicate to you that GPT-3 was trained on a dataset of 300 billion tokens?
A I need some time to review the context for that, so let me get back to you. Thank you for the time. I have had the chance to review the context.
Q Okay. I'll repeat my question.
A Yeah.
Q Does the statement in the caption to Table 2.1 that "All models were trained for a total of 300 billion tokens" indicate to you that GPT-3 was trained on a dataset of 300 billion tokens?
A It does, yes.
I will note, though, that a lot of the results in the -- in the rest of paper talk about "few-shot training," which often involves providing additional tokens in addition in the tokens that something is trained on.
So a lot of the performance results rest not just on -- on the pretraining, but also on the -- in-context learning. In fact, this whole paper is about how, when you add more tokens to the basic

46

model size, the performance can improve quite a bit.
Q And you earlier agreed that GPT-3 was a general-purpose LLM. So does this then indicate to you that it's possible to train a general-purpose LLM on a dataset of 300 billion tokens?
ATTORNEY HOLTZBLATT: Objection -- I object. Mischaracterize prior testimony.
A The main thing that I find interesting from this paper is, firstly, that this model was indeed -- like, if you look at the "Limitation" sections, clearly to me seems to be very much a research model in the sense that it is not economically valuable.
There's a -- there's -- it basically struggles in many, many different things that modern-day LLMs are quite good at, and so I would think -- and as the history has borne out, the economic value of GPT-3 as a model for practical applications was very low.
Q (BY ATTORNEY MILLER) So I wasn't asking what you found interesting about the paper.
I was asking: Does the fact that GPT-3 was trained on 300 billion tokens indicate to you that it's possible to train a general-purpose LLM on 300 billion tokens?

47

A What it tells me is that when you take 300 billion tokens and add additional information such as in-context learning, then the model can kind of do okay on a -- on a -- few of -- a different variety of tasks. That's the main conclusion that I draw from this paper.
Q Could you turn to Footnote 53 in your opening report, please, which is on page 16?
A Yes, I'm there.
Q Okay. And do you see in the last sentence of Footnote 53, you state:
"Common Corpus, an open-license pretraining dataset contains roughly two trillion tokens"?
Do you see that?
A Yeah, the -- all of Footnote 53 says:
"The question of whether there's enough public domain data to train a general-purpose LLM with a technical and empirical one. But I have not seen any sources that indicate that public domain data would be sufficient. Common Corpus, an open-license pretraining

48

dataset contains roughly two trillion tokens."
Q (BY ATTORNEY MILLER) Okay. So Common Corpus then contains more than six times as many tokens as was used to train GPT-3, correct?
ATTORNEY HOLTZBLATT: Object to form.
A Again, like I said, with GPT-3, it's hard to know the precise tokens important for the performance because a lot of the GPT-3 paper is about additional tokens can help improve performance.
So again, hard to make a comparison between those two in terms of exact proportions.
Q (BY ATTORNEY MILLER) You say that:
"The question of whether there's enough public domain data to train a general-purpose LLM is a technical and empirical one."
Do you see that?
A I do.
Q Have you done any technical or empirical analysis to determine whether there is enough public domain data to train a general-purpose LLM?
A I know that there are a variety of different efforts underway to use public domain data

Pages 61-68

61

A It depends on the particular tasks. There's some tasks for which LLaMA is okay; there are others where which it is not.
I remember experiments in the lab where LLaMA was not functional in some of the things -- in some of the basic things we wanted it to do.
Q (BY ATTORNEY MILLER) Is it functional across a range of activities?
A I have not tested that, so I cannot say anything to that regard.
Q What's the minimum amount -- amount of data I would need to train a general-purpose LLM, as you define it?
ATTORNEY HOLTZBLATT: Object to form.
A The -- I think what I have on the record is -- is kind of what I already said, which is that [redacted] It has to be an order of magnitude -- so many, many trillions -- in order for these things to start becoming intelligent and for them to be valuable across different types of tasks.
I will also note that the relationship between training data size and the performance of these models is not linear. So just because you have a model that is -- uses half the data, it

62

doesn't mean that it's half as good. Often a lot of these models are not performant, are not economically valuable at a certain level; and only at these very high levels of tokens do they become economically valuable.
So again, these examples of what -- Claude 3.5 Sonnet offers a touch point for ballpark we're talking about, but I don't have an opinion on any specific number, if you're looking for one.
Q (BY ATTORNEY MILLER) Could I train a general-purpose LLM, as you define it, on one trillion tokens?
ATTORNEY HOLTZBLATT: Objection, asked and answered.
A Again, I do not have any specific numbers that I can give you today.
Q (BY ATTORNEY MILLER) Could I train a general-purpose LLM, as you define it, on 500 billion tokens?
ATTORNEY HOLTZBLATT: Same objection.
A Again, I'm not here to offer any opinions about specifically how many tokens are needed to train a general-purpose LLM.
Q (BY ATTORNEY MILLER) If you train a model on a smaller amount of data -- withdrawn.

63

Withdrawn.
You have had said that one of the features you look at in determining whether a model is a general-purpose LLM is whether it's economically valuable, right?
A What I said is that for the -- for the economic questions at hand, in order to -- in order to have a market where somebody is assembling a large and diverse data, it's important that the general-purpose LLM that's trained is functional and economically valuable.
Q Okay. And whether a model is economically valuable can depend, in part, on what other models are available in the marketplace, right?
ATTORNEY HOLTZBLATT: Object to form.
A The way I would characterize it is that the first question that any model developer needs to answer is whether -- even absent any other -- any other models, whether consumers would pay for -- for all of the -- and pay to a degree that exceeds all of the cost, time, and energy required to train a model. So that's often the first consideration.
And then there are also considerations that are relative performance and relative value.
Q (BY ATTORNEY MILLER) A model might be

64

economically valuable at a certain performance level if there's minimal competition, but not economically valuable if other more powerful models become available in the marketplace, right?
ATTORNEY HOLTZBLATT: Object to form.
A Do you mind repeating that question?
Q (BY ATTORNEY MILLER) A model could be economically valuable at a certain performance level if there's no competing models, but not economically valuable if there are more powerful models available in the marketplace, right?
ATTORNEY HOLTZBLATT: Object to form.
A Usually the nature of -- like I said, there's two reasons why something might not be economically valuable. One is independent of any competition, consumers don't find it useful. It's actually a very common occurrence in the business model.
And then the second is these questions around relative performance. And they're usually -- it's not an absolute question about -- they might be less economically valuable when -- where there's competition, even though -- so say something is -- is slightly worse. There still might be consumers who end up using the model for -- for other reasons,

65

especially if its cost is low, for example.
So -- so in figuring out the economic value -- in light of -- competition is much more complicated question. It depends on the pricing strategy, the marketing strategy, the different array of tasks.
These models are often -- some models are good at some tasks, some at others. So it's very hard to say -- it's not a simple comparison where there's, like, two models and the better one gets the full market.
Q If you train a model on a smaller amount of data, the model will still -- the model will still function, right?
ATTORNEY HOLTZBLATT: Object to form.
A Again, this relationship is actually quite nonlinear according to a variety of -- according to my understanding of how these models work.
And so there are certain levels where, again, like, if the dataset is even a little bit smaller than a certain threshold, then the models are often completely useless at certain tasks, while at higher levels these properties become emergent.
So to say that if a model is slightly smaller, it might still be functional, firstly, it

66

conflates this idea that measuring functionality of this model is extremely complicated. There is many, many different types of tasks and ways to do it.
And second, the relationship between data size and performance is nonlinear, and so it's very hard to say where that cutoff will be when functionality degrades beyond economic value.
Q (BY ATTORNEY MILLER) There's a lot there. Let me go through a few things.
So first of all, if the model is trained on less data, it is less powerful, it will still engage in next-token prediction and produce an output; the output just may not be very valuable, right?
ATTORNEY HOLTZBLATT: Object to form.
A It's possible it doesn't even produce any output. There are many reasons why models fail.
Q (BY ATTORNEY MILLER) Well, it's not just going to spit sparks and go up in smoke, right? It's going to do what it's been trained to do; it may just not do it very well?
A Well, there are lots of reasons why models fail, including, like, running out of context or technical issues, why they don't produce outputs.
So exactly how models fail is -- is --

67

it's a complicated question.
Q You said there are certain thresholds below which the models become completely useless.
What are those thresholds?
ATTORNEY HOLTZBLATT: Object to form; mischaracterizes prior testimony.
A Yeah, I do not believe I said that there are certain thresholds. In fact, I have been quite clear that I'm not here to offer any opinions on thresholds or -- or minimum sizes.
All I said is that when you change token size upwards or downwards, how exactly the model will respond, either positively or negatively, is very hard to say.
Q (BY ATTORNEY MILLER) Well -- but thresholds matter, don't they? If you're trying to figure out whether you can acquire, via public domain or licensing, a certain amount of data to build a model, don't you need to know how much data you need to do that?
ATTORNEY HOLTZBLATT: Object to form; and argumentative.
A My main assignment here, especially in that first market, was to understand the order of magnitude of the number of tokens it takes to train

68

one of these models and then really to understand the economics behind doing that, including the transaction costs involved.
And so to understand the core economic issues at hand, I don't need to know any specific thresholds. I need to know the order of magnitude task that we are talking about.
So for example, in -- in paragraph 37, I give you one example of kind of what I'm working with. It's like -- so I just tried to understand books, like, in order to understand [redacted]
So in order to do the economic analysis, I needed an order of magnitude of the types of

Pages 73-76

73

Q Okay. Could you please turn to paragraph -- withdrawn.

The market for all training data needed to build an LLM that you analyze in your report, that's not limited to just the market for the works-in-suit specifically, right?

A Can you repeat that question? Sorry.

Q Sure.

When you discuss the market for assembling all of the data needed to train a general-purpose LLM, you're not specifically focusing on the market for the works-in-suit in particular, are you?

A What I'm focusing on is all of the different types of diverse data required to train a general-purpose LLM. That goes beyond the -- the works-in-suite, which is 499 lyrics of songs.

THE COURT STENOGRAPHER: 499 lyrics?

A Correct.

Q (BY ATTORNEY MILLER) Is it your view that a general-purpose LLM developer would not license any data for training unless it already had access to some minimum amount of data necessary to make the model -- the model functional?

A Give me a second. I think I state in my reply report in paragraph Number 70 that deals,

74

which I call:

"gap filling" or last-mile deals would occur -- only occur in a context where general-purpose LLM developers had already been able to amass the data needed to train the models. In other words, in order for licensing deals to be feasible, there needs to be already a sufficiently large and diverse training corpus for a general-purpose LLM.

Q Assume for a moment that general-purpose LLM developers could not freely access large sources of potential training data like Common Crawl or scraping the Internet. And for purposes of this hypothetical it doesn't matter why they can't, but assume that they can't. Is it your opinion that LLM developers -- excuse me.

Is it your opinion that general-purpose LLM developers would just shut up shop and close down?

ATTORNEY HOLTZBLATT: Object to form.

A In your question there's an assumption that general-purpose developers exist and then they

75

are looking for -- for data and they -- what they can and cannot do. The history of the industry shows that the very existence of this technology and the fact that we know that models constructed this way can be quite economically valuable relied on amassing a large enough sample or a large enough dataset.

And so therefore, if it wasn't possible, I can't expect to be -- to have been an active area of research or development or product development. So in that sense, I wouldn't use "shut up shop," as the shop would never exist to begin with.

Q So it's your opinion that if LLM developers could not freely access large sources of potential training data like Common Crawl or scraping the Internet, then general-purpose LLMs would not exist?

A The --

ATTORNEY HOLTZBLATT: Object to form.

A -- it's my opinion that the history of the field shows that having access to a vast and diverse data was important for the discovery of key technologies that enable this -- this -- that enable the current general-purpose LLMs to exist, yes.

Q (BY ATTORNEY MILLER) In paragraph 32 of

76

your initial report -- or excuse me. I think it's paragraph 33. In the middle of the paragraph 33, you say:

"No market for assembling all such data exists or is likely to emerge because the transaction costs of conducting these transactions at the scale necessary to successfully train a general-purpose LLM would be prohibitively high."

Do you see that?

A Sorry. I'm having trouble finding it. It's a long paragraph, so let me just review the whole thing.

I believe you were starting at that -- after the common, right, "No market for assembling"? So I think the whole sentence says:

"Therefore, even though occasional deals for training data might materialize, no market for assembling all such data exists or is likely to emerge because the transaction costs of conducting these transactions at the scale

Pages 185-188

185

anyway?
ATTORNEY HOLTZBLATT: Same objections.
A I don't know what you mean by "a typical market." And I -- my opinions are grounded to this particular market that I'm analyzing.
And one of the things that I highlight in my report is that markets for intellectual property such as -- such as lyrics have unique characteristics that are important to take into account. In -- in particular, the nonrivalry of -- of those goods. Right. So if one party consumes it, that doesn't leave any less for the second party that wants to consume it.
As a practical matter, I know that in this case there are costs involved for Anthropic in identifying what constitutes lyrics on the open web and dealing with that separately, so I imagine those costs also factor into the decision of what precisely is or is not used in the training data. Yeah.
Q (BY ATTORNEY MILLER) How can any market function if the buyer always has the alternative of just taking the product for free?
ATTORNEY HOLTZBLATT: Objection, argumentative.

186

A Dr. Smith and I agree on this fact that the ability to exclude is an important factor of a well-functioning market. But what is very important to know is when considering markets for intellectual property, we often -- economists, we often think about them as a necessary evil. Right. So exclusion can create as much harm as it can create good.
And so the goal of intellectual property is really to balance those restrictions against the fact that access to knowledge is important for innovation and creativity in downstream markets. So the whole job of intellectual property rights is actually to strike that balance between what one can use and access freely versus cases where that should not be permitted, which is what makes intellectual property so unique relative to other markets.
So our logic for regular markets doesn't often apply in markets for IP. In market for IP, the legal frameworks are very important in helping figure out when it is okay to use something for free and when -- and when it is not.
Q (BY ATTORNEY MILLER) As we have discussed, you have offered the opinion that lyrics have minimal value for training a general-purpose LLM,

187

right?
ATTORNEY HOLTZBLATT: Object to form; misstates prior testimony.
A I believe the phrase I used is "low value."
Q (BY ATTORNEY MILLER) Can you turn to paragraph 64, please --
A 64.
Q -- which is on page 30?
A I made the same mistake again. Yes, I'm there.
Q And do you see the first sentence of paragraph 64 begins:
"In some, for all the reasons stated, lyrics have minimal value to training a general-purpose LLM"?
Do you see that?
A I do see that.
Q So you would agree -- so it's your opinion that lyrics have minimal value to training a general-purpose LLM, correct?
A Yes. So another phrase I used before in the testimony was "low value."
Q Is it your opinion that Anthropic could remove all of the works-in-suit from Claude's

188

training data with no effect on Claude's performance?
A I don't have anything to say to that. I don't have an opinion on that regard.
Q Is it your opinion that Anthropic could remove all copyrighted song lyrics from Claude's training data with no effect on Claude's performance?
A Can you repeat the question?
Q Is it your opinion that Anthropic could remove all copyrighted lyrics from Claude's training data with no effect on Claude's performance?
A Insofar as the lyrics have some value to training a general-purpose LLM, it's very hard to answer that question because, as I mentioned before, "performance" is very difficult word to -- to talk about, so I don't want to offer an opinion on what Claude's performance would be under different conditions of data. That's a very loaded question and requires lots of deep thought; one that I have not really given.
Q Don't you think it's an important question?
ATTORNEY HOLTZBLATT: Object to form, argumentative.