EXHIBIT 137



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

20

language models include Claude and -- and LLaMA.
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] **scale that is required to train a general-purpose LLM.**

Q Can you train a general-purpose LLM on [redacted]

35

[redacted] **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 [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 129-136

129

economics of AI. We conduct research on the applications of AI in different chatbots. I also teach a very popular class called "The Business of AI" where we talk about many, many different applications of AI, including -- including projects and -- and -- and student reports, et cetera.

And never in any of those experiences has using an LLM to retrieve music lyrics been a topic of conversation. So it's based on my experience working with and teaching these models and -- with AI more generally over the last few years.

Q The last full sentence on this page states:

"The Clio report confirms that none of the Top 10 use cases for Claude.AI, which account for more than 60 percent of use cases analyzed, involve requesting song lyrics."

Do you see that?

A I do see that.

Q And then there's a footnote, right?

A Um-hum.

Q And you cite to -- the footnote points to Anthropic's interrogatory responses, correct?

130

A Correct.

Q Have you, yourself, reviewed any of the Clio reports?

A I have, actually. The Clio reports are a very interesting data source for anyone like myself doing the economics of AI. They provide a nice breakdown of how AI is used in the real world.

Q Did you review any of the Clio reports in connection with developing your opinions in this case?

A We did look at the category. I believe this came up in the -- in the -- in Dr. Smith's rebuttal where they talked about the extent to which Clio was used in the music industry, and so we looked at specifically how that was defined and how that was operationalized.

That was not actually in the report. That was in the -- in the -- they also have a website that shares some of the results, and so we looked at the website in order to delve deeper into those -- those questions.

Q In this sentence here when you point to "the top use cases for Claude," you're discussing relative use rather than absolute use, right?

A This -- this -- basically what they do is

131

they look at all the uses cases and then rank them. So in that sense, yes, it's -- it's a relative -- a relative analysis.

Q Why does relative use of lyrics matter rather than absolute use?

ATTORNEY HOLTZBLATT: Object to form.

A Can you clarify? Matter for what?

Q (BY ATTORNEY MILLER) Sure.

Why is it -- if lots of people are using Claude to search for lyrics, but they're -- but even more people are using Claude for other purposes, why does the fact that even more people are also using Claude for other purposes have any relevance?

ATTORNEY HOLTZBLATT: Object to form.

A Can you finish that question? Like relevance for whom or for what?

Q (BY ATTORNEY MILLER) For purposes of your analysis.

ATTORNEY HOLTZBLATT: Object to form.

A So for the purposes of my analysis, I need to understand the value that a general-purpose LLM developer like Claude might place on lyrics data. And so understanding their downstream consumer base and what it is that they're using the model for is pretty central to my analysis even in a relative

132

sense.

So if -- whatever the use cases are front and center for the developers will be the -- will then affect the types of data and the types of testing that they're subjecting the model to.

Q (BY ATTORNEY MILLER) In your analysis of the market for lyrics as training data for a general-purpose LLM, are you looking at a market where lyrics are the only training data for the LLM or where lyrics are one type of data, among many, for the LLM?

A The market that I'm looking for is assuming that there is already a large and diverse dataset on which a general-purpose LLM has been trained. And then the question: In order to fill in the gaps for gap filling, what's the potential of a market for lyrics in that -- in that scenario to develop?

Q Okay. All right. I'm going to upload the next exhibit.

(Exhibit 315 was marked for identification and is attached to the transcript.)

ATTORNEY MILLER: Do you want to break for lunch here or do you want to keep going for a little while?

133

ATTORNEY HOLTZBLATT: How are you doing?
THE DEPONENT: I can keep going for a bit more.
ATTORNEY HOLTZBLATT: Okay.
Q (BY ATTORNEY MILLER) Where is it? All right. The next exhibit should be in your -- in the chat. All right. And for the record, I'll designate this exhibit as Plaintiffs' Exhibit 315.
Do you have Exhibit 315 open in front of you?
A I do.
Q Do you recognize Exhibit 315?
A Yes, I believe this is "NBER Working Paper 33598 titled "Cloze Encounters: The Impact of Pirated Data Access on LLM Performance," dated March 2025."
Q And Cloze Encounters is spelled C-L-O-Z-E, right?
A Correct. C-L-O-Z-E, yes.
Q And that's a pun?
A No. Actually, it is a -- as it turns out, it's a psychology test that has always been spelled that way. I don't know why they chose the "Z" rather than the "S."
Q But using the name Cloze Encounters is

134

referencing --
A Yes. The whole thing is a pun because we are -- we are testing the idea of -- Cloze, exactly.
Q And you're a coauthor of this paper, right?
A I am.
Q And do you stand behind what's written in this paper?
A I do.
Q Can you turn to page 2, please?
ATTORNEY HOLTZBLATT: And is it page 2 in the PDF or page 2 as numbered?
ATTORNEY MILLER: That's a good question. I just have page 2 in my notes. So let me figure out where it is.
Q (BY ATTORNEY MILLER) Ah, page 2 in the paper, not in the PDF.
At the bottom of the page, do you see the sentence that begins "From a theoretical point of view"?
A I do, yeah.
Q And the paper states:
"From a theoretical point of view, it is not clear how access to specific sources of copyrighted

135

content affects model performance."
Do you see that?
A Um-hum.
Q And then it continues:
"On one hand, it is reasonable to infer that any one source of content has very little bearing on performance because of the sheer volume and diversity of data used in training," correct?
A Correct.
Q And then I'm going to skip a sentence. The next sentence says:
"On the other hand, as evidenced by an active licensing market for such data, it is clear that training data and content are the lifeblood of generative AI foundation models and are, therefore, very valuable in determining model performance."
Do you see that?
A I see that sentence.
Q Do you agree that there's an active licensing market for training data for

136

general-purpose LLM models?
A Just give me a second to review the first two paragraphs, because I haven't read this in a while. Okay.
So just as context -- so this paper is not analyzing any -- any direct -- any questions directly as the results with this case, including it doesn't really analyze lyrics data, nor does it like at substitution or sort of cannibalization.
THE COURT STENOGRAPHER: Any sort of?
A Cannibalization. Sorry. An economics term.
But in the -- but it's trying to make sense of the emerging questions around copyright policy and intellectual property policy in this area.
And in paragraph 2 we talk about some of the ongoing lawsuits, as well as we give examples of publicly reported deals. I believe the whole paper has only examples of three deals. That -- those were all that we were aware of at that point, including one between Reddit and OpenAI. And -- and so throughout the paper we use this phrase of "licensing market" or "active licensing market" as a shorthand for these -- for these three deals that we

Pages 213-216

213

Q You understand that writing lyrics can be a separate job from composing music, right?
A I do know that there are people who specialize in writing lyrics for -- for -- as a -- as a profession, yes.
Q And if Claude is used to write lyrics, it's capable of completely replacing the individual who had previously written the lyrics, right?
ATTORNEY HOLTZBLATT: Object to form.
A That doesn't necessarily follow. Just because Claude can produce lyrics, it says nothing about to what extent it may or may not replace somebody whose job it is to write lyrics.
Q (BY ATTORNEY MILLER) Do you think it's likely that Claude could -- Claude -- do you think it's likely that Claude will be capable of replacing professional lyricists?
A Can you repeat that question?
Q Do you think it's likely that Claude will be capable of replacing professional lyricists?
A I don't want to offer an opinion on an analysis that I did not conduct. I didn't conduct an analysis empirically on Claude and its impacts on replacing professional lyricists, especially in the future as this technology and the profession more

214

generally will continue to evolve.
Q (BY ATTORNEY MILLER) Spotify doesn't currently label all AI-created works as such, does it?
A I don't remember the exact details now, but I believe Spotify has some efforts underway, but it does not -- it does not label AI-generated content on lyrics in a widespread way.
Q Okay. And Deezer does label AI-generated content, or at least it attempts to, right?
A I believe Deezer offers a product and a service, but it aims to label AI-generated content as such, yes.
Q Well, when Deezer's service labels an AI-generated content as such, it doesn't tell you which AI was used to create the -- the music and lyrics in the AI-generated recording, right?
A I did not review Deezer's and how this works. In fact, when I tried to dig deeper, there was no information on Deezer methodology or how it label s AI that I could find, so I can't comment on that -- on that point.
Q There's no way to go on Spotify or Deezer and identify all of the tracks available on those platforms that include lyrics written with Claude,

215

is there?
A The platforms that -- one of the things I know from my research on use of technology in -- in music creation, it's always been very hard to understand precisely what technologies are used in the production of any content.
In my own research I studied digital music production. It's very hard to know if a particular song was recorded in analog studio or using digital -- digital technology.
And AI is just another technology in the mix of -- that is used in combination potentially with -- with other technologies, so it's not possible to -- in a reliable way to know what technologies are used in the production of content in general, and that includes AI.
Q In paragraph 30 of your rebuttal report, you note that:
"Deezer determined that AI-generated songs account for .5 percent of streams."
Do you see that?
A Yes.
Q So that's a nonzero amount, right?
A Correct.

216

Q Do you think that percentage is likely to increase in the future?
A I don't want to speculate on what's likely to happen in the future partially because I'm not exactly sure how Deezer's algorithm works and how it will continue to work in the future under different circumstances.
So this number depends not just on the use of AI in different parts of the market, but also on the algorithms used and people's responses to the use of those algorithms. So it's very hard for me to speculate on what will happen to that number.
Q If the use of Claude to write lyrics becomes widespread and pervasive, would that have any adverse impact on the potential market for Plaintiffs' lyrics?
ATTORNEY HOLTZBLATT: Object to form.
A Not necessarily. Because as I outline in my report, a lot of the use cases of Claude to light -- write lyrics involve noncommercial users, includes personal use, and broadly account as market expansion.
And so there's certainly a possibility that a majority of the original lyrics created with Claude would be completely parallel or independent

Pages 257-260

257

of the effect of Anthropic's conduct on the market for the online display of lyrics; is that right?
A That's correct.
Q You agree that a market for the online display of lyrics exists?
A I do.
Q And you agree that's a relevant market to consider for purposes of your analysis?
A I do.
Q And you attempt to do an empirical analysis to ascertain whether the -- there's an effect on Anthropic's conduct on that market, right?
A That's correct.
Q If use of Claude to search for lyrics became unrestricted and widespread, do you think that would result in a substantially adverse impact on the market for the online display of lyrics?
ATTORNEY HOLTZBLATT: Object to form; vague.
A It's hard for me to, again, speculate about future consumer behavior, how different market participants would behave, et cetera, so I -- I -- I reserve -- I have no opinion on that -- on that question.
Q (BY ATTORNEY MILLER) Okay. In order to do

258

your empirical analysis, you do a series of -- you run a series of regressions, right?
A Correct.
Q Is a regression analysis always an appropriate method to evaluate cause-and-effect relationships?
ATTORNEY HOLTZBLATT: Object to form.
A Regression analysis is a well-accepted method to look at the relationship between different variables, and it's very commonly used in the academic literature.
Q (BY ATTORNEY MILLER) Are there circumstances in which regression analysis is an inappropriate or unreliable method for evaluating cause-and-effect relationships?
A I imagine there are circumstances when that's certainly true. Including regression analysis is a quantitative methodology. And so when you're looking at potential markets or markets where no data exists, then a regression analysis would not be possible to do.
Q If you have noisy data, can that make regression analysis unreliable?
A In some ways noise and variability in data is fundamental to all of econometrics and

259

statistics. It's something that is understood or taken for granted when one is doing regression analysis.
And so no, I think things like noise in data -- there are techniques in -- in econometric analyses that are built to account for -- for -- for that kind of variability.
We assume always that there is -- there is a level of measurement or -- or noise in the data, and so that rarely poses a challenge to econometric analyses.
Q Is there a minimum quantity of data or sample size you need in order to run a reliable regression?
A I have not seen a minimum number or size in any of the econometric textbooks I have studied from.
Q You attempt to use a regression to compare Claude usage to traffic to lyric data websites, right?
A In the analysis in the opening report, including what I discuss in paragraph 102, I talk about how I use website traffic data to the top lyrics websites to understand the impact of the launch of Claude.AI traffic -- and traffic to

260

Claude.AI on traffic to these other lyric websites.
Q Are you familiar with a -- with what's sometimes called a Type II error?
A I always get Type I and Type II confused, so you'll have to define that for me.
Q I think it -- are you familiar with the possibility of a false-negative error?
A Again, it depends on the context, but "false negative" is sometimes when you think -- you say that something is false, when actually it's true.
Q Can noisy data cause a false negative?
A If econometric analyses are not performed carefully, then you can make the wrong conclusions for them, sure.
Q You used as your -- the main variable interest in your regression, web traffic to Claude.AI, right?
ATTORNEY HOLTZBLATT: Object to form.
A Can you repeat the question?
Q (BY ATTORNEY MILLER) You use, as the main variable of interest in your analysis, web traffic to Claude.AI, right?
ATTORNEY HOLTZBLATT: Same objection.
A Usually regressions are defined both by a