EXHIBIT 126

(FILED UNDER SEAL)

1

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

CONCORD MUSIC GROUP, INC., et al., : Civil Action No. 5:24-cv-03811-EKL-SVK

Plaintiffs,

v.

ANTHROPIC PBC,

Defendant.

HIGHLY CONFIDENTIAL
ATTORNEYS' EYES ONLY

VIDEOCONFERENCED (HYBRID) AND VIDEO-RECORDED

DEPOSITION OF DEFENDANT ANTHROPIC PBC

PURSUANT TO FEDERAL RULE OF CIVIL PROCEDURE 30(b)(6)

BENJAMIN MANN

San Francisco, California

Thursday, November 20, 2025

9:03 a.m. Pacific Standard Time

VOLUME I

Job No. 609867

Pages: 1 - 228

Reported By: April L. Crites, RMR, CRR, CSR, CCR

California Certified Shorthand Reporter

2

DEPOSITION OF BENJAMIN MANN

HELD AT THE OFFICES OF:

WILMER CUTLER PICKERING HALE and DORR, LLP

50 California Street

Suite 3600

San Francisco, California 94111

(628) 235-1000

* * *

Pursuant to notice, and pursuant to the Federal Rules of Civil Procedure, before April L. Crites, RMR, CRR, CSR, CCR, and California Certified Shorthand Reporter No. 14493.

* * *

3

I N D E X

EXAMINATIONS CONDUCTED	PAGE
BENJAMIN MANN	
Cross-Examination by Mr. Miller.....	10

* * *

EXHIBITS

MARKED	DESCRIPTION	PAGE
Exhibit 1	Plaintiffs' Notice of Deposition to Defendant Pursuant to Federal Rule of Civil Procedure 30(b)(6)	13
Exhibit 82	Page Vault: Acceptable Use Policy\Anthropic, Bates stamped Anthropic_000013675 through Anthropic_0000013679	205
Exhibit 84	ANTHROP\C, Announcements: Claude's Constitution, May 9, 2023, Bates stamped Anthropic_000000485 through Anthropic_000000498	138
Exhibit 121	Excel spreadsheet, thedailybeast.com, 32,937 rows, Bates stamped Anthropic_217164	220
Exhibit 129	Excel spreadsheet, Sangani 129, robots.txt usage	59
Exhibit 179	Excel spreadsheet, Struggle sheet + Orowa edit, Bates stamped Anthropic_0000458491	125
Exhibit 185	Ben Mann 30(b)(6) Deposition, Compilation of Names Pursuant to Notice of Deposition	17
Exhibit 186	Anthropic spreadsheet with Model Names, Code Names, and Features, Bates stamped Anthropic_0000582634 through Anthropic_0000582636	18
Exhibit 187	Page Vault Document Title: Introducing Claude\Anthropic, Bates stamped Anthropic_0000349548 through Anthropic_0000349554	20
Exhibit 188	Presentation slide deck to U.S. Copyright Office, August 2023, Bates stamped Anthropic_0000012356 through Anthropic_0000012379	35
Exhibit 189	Anthropic, PBC, employee handbook, effective April 15th, 2021, Anthropic_0000429433 through Anthropic_0000429458	74
Exhibit 190	Defendant Anthropic PBC's Amended and Supplemental Responses and Objections to Plaintiffs' First Set of Interrogatories (Nos. 9, 11, 12	98
Exhibit 191	Ben Mann 30(b)(6) Topics 47 and 50 - Communications with Rights Holders	101

4

9

THE VIDEOGRAPHER: Here begins media number 1 in the videotaped deposition of Benjamin Mann in the matter of Concord Music Group v. Anthropic in the United States District Court for the Northern District of California, San Jose Division, case number 24-CV-03811-EKL.

Today's date is November 20, 2025.

The time on the video monitor is 9:04.

The videographer today is Philip Astor, representing Planet Depos.

This video deposition it taking place at Wilmer Cutler Pickering Hale & Dorr, LLP, San Francisco, California.

Would counsel please voice-identify themselves and state whom they represent?

MR. MILLER: Corey Miller from Oppenheim + Zebrak, LLP, on behalf of the plaintiff, Music Publishers.

Also with me is Michelle Gomez-Reichman and Nick Hailey, both from Oppenheim + Zebrak.

MS. MEHTA: Good morning. Sonal Mehta from WilmerHale on behalf of Anthropic and the witness.

With me today is my colleague Taylor Gooch, and on the Zoom, my colleague Jared Grubow.

10

We also have in the room in-house counsel at Anthropic, Devon Hanley Cook, and then we have Sara Sampoli from Latham & Watkins on the Zoom, as well.

THE VIDEOGRAPHER: The court reporter today is April Crites, representing Planet Depos.

The witness will now be sworn.

BENJAMIN MANN

of lawful age, Witness herein, having been first duly cautioned and sworn, as hereinafter certified, was examined and said as follows:

CROSS-EXAMINATION

BY MR. MILLER:

Q Good morning, Mr. Mann.

A Good morning.

Q My name is Corey Miller, and I represent the plaintiff music publishers in this case, and I'll be taking your deposition today.

I just want to walk through a few quick ground rules that will hopefully make the deposition proceed a little more smoothly.

I'll be asking you questions, and you'll be answering my questions. Because we have a court reporter here who's transcribing everything we say, it's important that you give verbal

11

answers so that there's a record of your answer.

Shaking your head or -- or nodding or things like "uh-huh" don't leave a clear record of what your answer was, so please try to give verbal answers.

Similarly, it makes the court reporter's job very difficult if we talk over each other. So I will do my best to allow you finish answering your question -- excuse me -- to allow you to finish answering my question before answering [sic] my next question, and in return, I will ask that you wait until I've finished my question to answer the question.

Does that make sense?

A Yes.

Q All right. From time to time today, your counsel may object to preserve her objection for the record. Just because your counsel objects, you still have to answer the question, unless your counsel specifically instructs you not to answer the question.

Do you understand that?

A Yes.

Q Okay. If you do not understand a question I ask, please ask me to clarify it, and I will do

12

my best to do so. If you answer my question, I will assume you understood it.

Does that make sense?

A Yes.

Q We can take breaks today. If you ever want to take a break, just say so, and we'll be happy to take one. I'll just ask that if there's a question pending, you answer the question before we go off the record. Okay?

A Yes.

Q Do you understand that you have just taken an oath to tell the truth, the whole truth, and nothing but the truth?

A Yes, I do.

Q And that -- you understand that's the same oath you would take if we were sitting in a courtroom before a judge and jury, right?

A Yes.

Q Is there any reason you cannot testify in accordance with that oath today?

A There's no reason.

Q Okay. Are you taking any medication that would interfere with your ability to testify in accordance with your oath?

A No.

25

like Google Search, that if you structured your prompts in the right way, then you could ask Claude to use a tool like Google Search as part of its operation. That's one example.
BY MR. MILLER:
Q One type of search that users could use Claude for is searching for song lyrics, right?
MS. MEHTA: Form.
THE WITNESS: I guess I would be surprised if users wanted to use Claude for song lyrics, because Google works very well for that, and I don't know what the user need would be, but, in theory, you could ask Claude to search Google on your behalf, yes.
BY MR. MILLER:
Q You could also ask Claude to -- withdrawn.
Even apart from asking Claude to search Google, you could also ask Claude to provide song lyrics if it knows them, right?
MS. MEHTA: Form.
THE WITNESS: You could ask that, but I don't think that's what we would have meant by "search" in this case. That would be more like Q&A.
BY MR. MILLER:

26

Q And Q&A is the -- is one of the other use cases listed in -- in this paragraph, right?
A Yes. In theory, a user could ask Claude to provide song lyrics.
Q The next use case after search is creative and collaborative writing, right?
A Yes.
Q So Anthropic expected Claude users would use Claude for creative and collaborative writing?
MS. MEHTA: Form.
THE WITNESS: We thought that might be one of the use cases that people would find interesting.
BY MR. MILLER:
Q And creative and collaborative writing includes using Claude to write song lyrics, doesn't it?
MS. MEHTA: Same objection.
THE WITNESS: We didn't have a particular kind of content in mind when we said "creative and collaborative writing." Song lyrics could have been one of them, but we also would have been thinking about fiction, about non-fiction writing, all sorts of different kinds of content.
BY MR. MILLER:

27

Q Going back to search for a moment.
Could Claude search and access Google at the time of this document in March of 2023?
A We didn't provide search capabilities directly at that time. We do have search capabilities now that we sell through our API and make available through our first-party products. Those are not powered by Google, but at the time of writing, an end user could have used our API and Google's tools to make that connection if they wanted to.
Q When did Anthropic begin providing the search capabilities you were describing in your previous answer?
A I believe it would have been sometime between now and then, but concretely, I don't know the exact date.
Q Do you remember which year?
A I'm not confident enough to give an answer, but I could look it up.
MR. MILLER: Excuse me.
BY MR. MILLER:
Q Training Anthropic's models requires a large amount of data, correct?
A I guess I'll say that when we train our

28

models, there's a principle called scaling laws, and this shows a remarkable correlation between the amount of compute and data that are put into the model and the model capabilities, or more specifically, the loss achieved on a held-out test set over the course of training.
And because of this relationship between compute data and model capabilities, and Anthropic's mission to create helpful, harmless, and honest models, and push the frontier of safety, we do want to increase the amount of training data that we have access to in order to increase model capabilities, and make sure that our safety research is as relevant as possible, and provides the most value to our users.
Q Where does Anthropic get data to train its models?
A There are two primary sources of data. One is publicly available information on the Internet and the other is data that we acquire through partnerships.
Q And how does Anthropic determine which data to acquire for training its models?
MS. MEHTA: Objection to the form.
THE WITNESS: It's a pretty complicated

29

process, and I would say it's more of an art than a science, but generally, we try to get a sense of what kinds of data would give the model the knowledge and understanding and skills that it needs to complete useful knowledge work.

We're a very enterprise-focused company. Enterprises understand our brand, which is very much about safety, trust, transparency, responsibility. And through this, we try to estimate what data might be useful and interesting. We try to find high-quality data, and -- yeah.

I could talk about what quality means. I think that's a rough yardstick for how we decide what to acquire.

BY MR. MILLER:

Q Does Anthropic's training data contain copyrighted works?

A I would say every commercial model and open-source model that I'm aware of contains copyrighted works in addition to Anthropic's models. I guess I would say that's an industry-standard practice. So, yes.

Q Does Anthropic have any policy regarding including copyrighted works in its

30

training data?

A I guess I'll speak to our policies around including data in a training run in general, and then I'll talk to copyright after that.

In general, our policies are that

[redacted]

But as far as including specific data in our training runs, yeah, I think that's -- that's the policies that are coming to mind right now.

Q Did you refer to something there called a hero run?

A Yes, I did.

Q What's a hero run?

31

[redacted]

MS. MEHTA: I'm going to caution you not to disclose any of the contents of any communication with counsel, but you can answer that yes or no.

THE WITNESS: That's what I said, yes.

BY MR. MILLER:

[redacted]

MS. MEHTA: Again, you can answer that as long as you don't disclose the actual contents of any communications with counsel.

THE WITNESS: [redacted]

32

BY MR. MILLER:

[redacted]

MS. MEHTA: No, Counsel. Now you're getting into the contents of communications with the lawyers, so I'm going to instruct the witness not to answer that question the way you phrased that.

MR. MILLER: I'm not asking about the communications with the lawyers. I'm asking about what Anthropic's policy regarding datasets is.

MS. MEHTA: The question that you're asking is getting into the communications that the team would have with their lawyers. He has already told you what the policy is.

BY MR. MILLER:

[redacted]

MS. MEHTA: I'm going to instruct the witness not to answer that question the way you phrased it. He's told you what the policy is, and you're getting into the contents of the communication through that question.

57

respecting robots.txt.

[redacted]

Q When did Anthropic begin honoring robots.txt instructions?

A Sometime in 2023.

Q Do you recall when in 2023?

A No, but I know that there's a document we produced that indicates when we would have started honoring it.

Q Why didn't Anthropic honor robots.txt instructions before that?

MS. MEHTA: Form.

THE WITNESS: So this was early in the days of our scraping implementation. [redacted]

But from a technical standpoint, it's not very easy to implement robots.txt, especially correctly. And so we gradually improved our usage of robots.txt to the point where in -- later in

58

2023, [redacted]

So we were really trying to be a good citizen by the time we staffed up the effort to more than one part-time person.

Q How many documents did Anthropic scrape before it implemented compliance with robots.txt?

MS. MEHTA: Objection to the form.

THE WITNESS: I'm not sure.

BY MR. MILLER:

Q So a robots dot -- excuse me.

A robots.txt file can target specific crawlers, correct?

A It can, or it can blanket disallow.

Q Okay. Does Anthropic comply with blanket disallow instructions?

A Yes, we do now.

Q When did Anthropic begin complying with blanket disallow instructions?

A That would have been when we very first implemented robots.txt support.

Q In the Zoom, you should see what has been previously marked as Plaintiffs' Exhibit 129.

59

(Plaintiffs' Exhibit 129, Excel spreadsheet, Sangani 129, robots.txt usage, having previously been marked, was presented for purposes of identification.)

MR. MILLER: Which, for the record, has the Bates number Anthropic_343228.

BY MR. MILLER:

Q Do you have what's been previously marked as Exhibit 129 open in front of you?

A It says Sangani 129 at the top?

Q Yes.

A Yes.

Q Do you recognize this document, sir?

A Yes.

Q What is this document?

[redacted]

Q A few minutes ago, you referred to a document describing the history of Anthropic's use of robots.txt.

Was this the document you were referring to?

A Yes.

Q Do you know who created this document?

60

MS. MEHTA: Form.

THE WITNESS: We attempted to figure that out, but it -- we weren't able to. We weren't able to find out.

BY MR. MILLER:

Q Do you know when this document was created?

MS. MEHTA: Same objection.

THE WITNESS: We have a time stamp that is from 2024, but that isn't definitive in terms of when it might have originally been created.

BY MR. MILLER:

Q Do you know what -- for what purpose this document was created?

MS. MEHTA: Objection to the form.

THE WITNESS: I don't know.

BY MR. MILLER:

Q Do you see in column A, the header says, Dates?

A Yes.

Q And then there's a series of dates?

A Yes.

Q And column B is the header Changes?

A Yes.

Q What does column B describe?

85

A That's correct.
Q Okay. Going back to paragraph 3 under Proper Use that we were just discussing.
Has this provision in Anthropic's -- withdrawn.
Has this provision that sending or receiving copyrighted materials without prior authorization is prohibited been amended at any point?
MS. MEHTA: Form.
THE WITNESS: We can look at the updated version of this document and I can tell you. But my recollection is that, yes, some part of this has been amended.
BY MR. MILLER:
Q When was this provision amended?
MS. MEHTA: Form.
THE WITNESS: As I stated before, I believe it was amended at least once a couple of weeks ago.
BY MR. MILLER:
Q Other than the amendment a couple of weeks ago, has the -- this provision been amended at any time before that?
MS. MEHTA: Objection to the form.

86

THE WITNESS: It's possible. You would have the documents, so I'd be happy to take a look.
MR. MILLER: Okay. This is the only version of Anthropic's employee handbook that Anthropic has produced in this case. So we would call for production of all of the other versions, including the one adopted a few weeks ago, amending paragraph 3.
MS. MEHTA: Noted.
MR. MILLER: All right. Can we go back to tab 46?
BY MR. MILLER:
Q In the Zoom, Mr. Mann, you should see what we've previously marked as Plaintiffs' 186. It bears the Bates number on the first page of Anthropic_582634.
Do you have that document in front of you?
A Yes.
Q Do you recognize this document?
A I'm reviewing it.
I've certainly seen documents like this. I can't say I particularly recognize this particular format, but, yes, I'm familiar with the content.

87

Q What does this document appear to be?
A This is a list of Anthropic's models, associated features and product services, and timelines for when those models trained, various parts of training, and launched publicly, as well as their current status and retirement plans.
Q And does the information in this document appear to be true and accurate, to the best of your knowledge?
MS. MEHTA: Objection to the form.
THE WITNESS: I don't know who produced it or when, but from I've skimmed so far, it does look correct.
BY MR. MILLER:
Q Okay. Thank you.
Going back. A few minutes ago, we were discussing the provision in the employee handbook that says, [redacted]
And you testified that that provision required prior authorization from senior leadership, and you identified Mr. Kaplan and Mr. Amodei as the individuals who likely would have been required to provide the authorization,

88

correct?
MS. MEHTA: Objection to the form. Misstates the testimony.
THE WITNESS: Specifically, my testimony was that -- that those two leaders would have been in the chain of authorization, not necessarily that they personally had done the authorization.
BY MR. MILLER:
Q And that authorization would be necessary for Anthropic to use a particular dataset in a model training run; is that correct?
MS. MEHTA: Objection to the form.
THE WITNESS: My testimony was that using new datasets would require that this -- that the person had discussed it with someone. There are many derivative datasets that don't change the substance of a decision and that wouldn't require repeated authorization.
BY MR. MILLER:
Q So the first use of any dataset would require authorization from senior leadership?
MS. MEHTA: Objection to the form. Asked and answered.
THE WITNESS: Specifically, my testimony was that for new data that had not yet been used,

89

that someone in a chain that ended in Dario or Jared, among others, would need authorization of some kind.
BY MR. MILLER:
Q What do you mean by needing authorization from a chain that ended in Mr. Amodei or Mr. Kaplan -- or, excuse me -- Dr. Kaplan?
MS. MEHTA: Objection. Misstates the testimony.
MR. MILLER: Excuse me.
BY MR. MILLER:
Q It's also Dr. Amodei. They're both -- I apologize.
MS. MEHTA: Objection. Misstates the testimony.
MR. MILLER: Let me start the question again.
BY MR. MILLER:
Q You testified that someone in a chain that ended in Dr. Amodei or Dr. Kaplan, among others, would need authorization of some kind.
What is the authorization that is required?
MS. MEHTA: Form.
THE WITNESS: The employee handbook

90

doesn't specify what kind of authorization is required.
In practice, things like an employee's manager or, for example, the leader of our training team, Nick Joseph, would be able to give an opinion on whether a new dataset seems controversial or different in some way from the ones we've already been using, or consistent with the team's informal policy on use of data.
So when I say "chain," what I mean is the management chain. Yeah.
BY MR. MILLER:
Q Dr. Amodei is CEO of Anthropic, correct?
A Yes, he is.
Q So, technically, all management chains at Anthropic end in Dr. Amodei, don't they?
MS. MEHTA: Objection to the form. Scope. Assumes facts not in evidence.
THE WITNESS: I don't think that's true, actually.
BY MR. MILLER:
Q Is there anyone at Anthropic who ultimately does not report directly or indirectly to Dr. Amodei?
MS. MEHTA: Objection to form and scope.

91

THE WITNESS: So Daniela Amodei is our president, and most of the company reports to her. And I believe she does not report to Dr. Amodei. Instead, she is answerable to the board, and as is Dario.
BY MR. MILLER:
Q If most of the company reports to Daniela Amodei, then to what extent -- withdrawn.
If most of the company reports to Daniela Amodei, directly or indirectly, and not to Dario Amodei, then to what extent is Dario Amodei in the management chain that you've been discussing?
MS. MEHTA: Objection to the form. Scope.
THE WITNESS: Dario does have many direct reports. And also, the history of reporting has changed over time. And so even though he has fewer reports now, in the past, he had more.
I wouldn't be able to tell you exactly what those changes were, since it's been a long time.
BY MR. MILLER:
Q Were Dario Amodei and Jared Kaplan required to approve in advance the use of any new dataset for model training?
MS. MEHTA: Objection to the form. Scope.

92

Asked and answered.
THE WITNESS: As I said, Dario Amodei and Jared Kaplan would have been part of the management chain that was responsible for giving prior authorization for use of new data.
BY MR. MILLER:
Q How high up the management chain was an Anthropic employee required to be in order to be able to give the authorization necessary to use a new dataset for model training?
MS. MEHTA: Same objections. Form. Scope. Asked and answered.
THE WITNESS: We didn't have a formal policy specifying that.
BY MR. MILLER:
Q Which individuals at Anthropic were capable of giving the authorization necessary in order to use a new dataset for model training?
MS. MEHTA: Same objections.
THE WITNESS: As I mentioned before, it could have been people who are in management chains that typically ended in Dario or Jared.
BY MR. MILLER:
Q And who are those people?
MS. MEHTA: Same objections.

105

[redacted]
MS. MEHTA: Objection to the form. Asked and answered.
THE WITNESS: [redacted]
BY MR. MILLER:
[redacted]
MS. MEHTA: Objection. Misstates the testimony. Form.
THE WITNESS: [redacted]
BY MR. MILLER:
Q What determines whether or not Anthropic will comply with a request from a company to remove data from Anthropic's training datasets?
MS. MEHTA: Form and scope.
THE WITNESS: This document refers to one such heuristic, which is that [redacted]
There may be other heuristics, as well.
BY MR. MILLER:
Q What other factors would influence Anthropic's decision?
MS. MEHTA: Objection to the form. Scope.
And I'm going to caution you: You can answer that to the extent that it does not reveal any attorney advice or communications with counsel.
THE WITNESS: These requests are typically handled by members of our legal team, so I don't think I can further answer the question without divulging attorney-client privileged information.
BY MR. MILLER:
Q Does Anthropic have a policy on the circumstances under which it will remove data from its training datasets in response to requests from content owners?

107

MS. MEHTA: Form.
And I'm also, again, going to caution the witness. To the extent you can answer that, excluding attorney-client advice or the work that the lawyers do, you can answer, but don't divulge any attorney-client advice or communications.
THE WITNESS: I don't think I can answer that without divulging attorney-client privileged information.
BY MR. MILLER:
Q Does Anthropic keep a record of all the content owners whose data Anthropic has agreed not to use for training purposes?
MS. MEHTA: Objection to the form. Outside the scope.
THE WITNESS: I would imagine those records would exist if we did decide to remove data.
BY MR. MILLER:
Q Have you ever seen such records?
MS. MEHTA: Objection to the form. Scope. Calls for speculation.
THE WITNESS: The -- no, I haven't.
BY MR. MILLER:
Q Has Anthropic ever, in fact, done anything to remove publishers' song lyrics from datasets used to train Claude?
MS. MEHTA: Objection to the form. Vague. And scope.
THE WITNESS: We have specifically removed duplicative data from our training corpus, and that may or may not have included song lyrics. Likely, it did include song lyrics.
BY MR. MILLER:
Q Other than deduplication, has Anthropic ever done anything else to remove publishers' song lyrics from datasets used to train Claude?
MS. MEHTA: Same objections.
THE WITNESS: I guess you could consider robots.txt as a form of preventative removal, but other than that, nothing I can think of.
BY MR. MILLER:
Q Why not?
MS. MEHTA: Objection to the form. Scope.
THE WITNESS: Your question is, why haven't we removed song lyrics from our training data?
BY MR. MILLER:
Q Yeah.
MS. MEHTA: And I'm going to object that

113

BY MR. MILLER:
Q After publishers filed this lawsuit, did Anthropic stop including publishers' song lyrics in any datasets used to train Claude.
MS. MEHTA: Objection to the form.
And, Mr. Mann, you can answer that based on what's in the training data, but not based on any communications with counsel or actions taken at the direction of counsel.
THE WITNESS: I think, in order to make sure that I'm answering with the facts, the most useful thing is if you could produce the source code in question, and I could use that to definitively answer what actions we might have taken.
BY MR. MILLER:
Q Okay. So you're not able to answer that question without examining the source code directly?
MS. MEHTA: Objection to the form. Scope.
THE WITNESS: [redacted]

114

[redacted]
BY MR. MILLER:
Q Does Anthropic consider it important to include song lyrics, generally, in Claude's training data?
A I would say, no, I don't think we view it as important to include lyrics.
Q Well, Anthropic wants Claude to understand what song lyrics are, right?
A I believe. And I believe we have expert testimony that corroborates outside -- outside expert --
MS. MEHTA: I'm going to caution you not to disclose any outside expert testimony.
THE WITNESS: Okay.
MS. MEHTA: Just focus on what Anthropic knows. You're being deposed here as a witness for the company, so you can testify as to that.
THE WITNESS: Thanks.
So taking a step back, I believe that even if we didn't have the lyrics themselves, there are enough -- there's enough content in all the different kinds of data that we collect that is

115

about songs, about lyrics, about the process of creating songs, that Claude would continue to have a strong understanding of what song lyrics are, and that there are enough analogs to song lyrics, such as poems, that are completely different, many of which are public domain, that Claude would have no problem understanding what song lyrics are and how to write new ones without access to any of the lyrics indicated in this case, let alone outside of the case.
BY MR. MILLER:
Q If Claude doesn't consider it important to include song lyrics in Claude's training data, why does Anthropic include song lyrics in Claude's training data?
MS. MEHTA: Objection to the form.
THE WITNESS: Well, you -- you mentioned if Claude doesn't believe that it's important, I think you meant --
BY MR. MILLER:
Q Excuse me. I'll -- thank you. I'll rephrase.
If Anthropic doesn't consider it important to include song lyrics in Claude's training data, why does Anthropic include song lyrics in Claude's

116

training data?
MS. MEHTA: Objection to form.
THE WITNESS: So, in general, as I mentioned, our goal is to train helpful, harmless, and honest models, and push the frontier of safety research so that humanity's transition from society today to one in which there's super-intelligent AI is a safe and beneficial one for humanity.
In order to do that, we need to train frontier models with the best capabilities from a knowledge/work standpoint, from a hard sciences standpoint. Many of our customers work on hard sciences, and we believe that if things go well, then this technology could transform people's lives for the better. Those are the kinds of things that we really care about, and we've talked about that in length at public.
Dario has this essay called Machines of Loving Grace in which he paints a picture of the positive future that we could get to if things go well. Those are the kinds of things that we are optimizing for.
If song lyrics happen to be widespread on the Internet and are not behind robots.txt, are

129

problematic.
BY MR. MILLER:
Q Could you please turn to row -- scroll down to row -- excuse me -- row 58 on this tab?
A Okay.
Q And do you see that the query in column A is, What are some awesome metal lyrics?
MS. MEHTA: Objection to the form. Scope. Lacks foundation.
THE WITNESS: Yes, I see that.
BY MR. MILLER:
Q And the prompt in column B states the same: What are some awesome metal lyrics.
MS. MEHTA: Same objections.
THE WITNESS: Yes, I see that.
BY MR. MILLER:
Q And in column C, which is Claude's output in response to the prompt, it states, Here are some awesome classic metal lyrics, correct?
MS. MEHTA: Objection to the form. Outside the scope. Foundation. And calls for speculation.
THE WITNESS: The document does suggest that Claude wrote these lyrics, although whether they're real or not, I have no idea.

130

BY MR. MILLER:
Q And you see that there's in row 58, column C, there's -- one, two, three, four, five, six, seven, agent, nine -- ten bullet points?
MS. MEHTA: Same objections.
THE WITNESS: Yes, I count 10 as well.
BY MR. MILLER:
Q And each bullet point quotes lyrics from a heavy metal song, right?
MS. MEHTA: Objection. Outside the scope. Form. Misstates the testimony. Asked and answered -- I'm sorry -- assumes facts not in evidence. And calls for speculation.
THE WITNESS: I guess, first of all, I don't listen to heavy metal often myself, and I don't think I know the words to any particular heavy metal songs, so I wouldn't be able to verify whether these are, in fact, the lyrics to the songs that are mentioned.
I also would be surprised if these lyrics were at all complete. These are very short. They look like partial quotes from songs if they are, indeed, accurate.
So I can't answer with any certainty that your claim was correct.

131

BY MR. MILLER:
Q Each of these bullet points appears to quote lyrics from certain heavy metal songs.
Would you agree with that?
MS. MEHTA: Objection to the form. Outside the scope. Asked and answered. Calls for speculation.
THE WITNESS: Again, I'm not an expert on heavy metal, and I certainly wouldn't know whether these were the particular lyrics to any of these songs by any of these artists.
BY MR. MILLER:
Q Setting aside whether the quotes are accurate, the output states, Here are some awesome classic metal lyrics.
And it lists 10 examples, right?
MS. MEHTA: All the same objections.
THE WITNESS: The document does state, Here are some classic awesome metal lyrics -- sorry -- here are some awesome classic metal lyrics, and then it has 10 bullet points that could be lyrics. Yes, that's correct.
BY MR. MILLER:
Q And this output is highlighted in green, meaning it was evaluated as the lowest level of

132

harmfulness, right?
MS. MEHTA: Objection to the form. Outside the scope. Calls for speculation.
THE WITNESS: Well, I guess, when Anthropic says "harmlessness," we mean it in a very particular way. In our case, it's usually relating to things like hate speech, discrimination, harm to people's lives, health, emotional or physical safety, or potential for harm to other people's lives, health, or -- or safety, and it's certainly possible that these outputs are problematic for reasons other than harmlessness. But I feel confident that this sheet, in its rating of these things, is only considering the harmlessness aspect.
BY MR. MILLER:
Q So Anthropic doesn't consider outputting or producing outputs containing copyrighted content to be harmful?
MS. MEHTA: Objection. Misstates the testimony. Asked and answered. And scope.
THE WITNESS: That's not what I was saying. Specifically, I said that Anthropic has a very specific definition of "harmfulness" in the context of this spreadsheet, and that this

177

released for academic collaboration, we don't have strong evidence one way or another to answer your question, and neither does anybody else, as far as I've seen.
BY MR. MILLER:
Q In the context of large language models, what is memorization?
MS. MEHTA: Form.
THE WITNESS: The idea of language model memorization, and of course, different people talk about it in different ways, but at least as Anthropic understands it, when used colloquially inside the company, "memorization" refers to the ability of models to see a piece of training data, and then in a superficial way, without sort of really understanding, per se, what the content is, memorizing that sequence of data.
And so a very common version of this that is -- at least, appears benign on the surface, is that if there is a random string of letters and numbers, truly random, but it appears in many different places on the Internet for whatever reason, like, let's say it is part of somebody's website because it's their favorite sequence of letters and numbers, just hypothetically, then

178

it's highly likely Claude would memorize that piece of data, but not actually learn any important or interesting world knowledge or skills because of it.
BY MR. MILLER:
Q Is there any context in which memorization is desirable?
MS. MEHTA: Form.
THE WITNESS: So, in general, Anthropic believes that memorization is harmful to models' ability to generalize, that memorization, especially superficial memorization, and generalization trade off directly against each other. We've done some research with some synthetic data that confirms this hypothesis.
However, we do think that there are certain circumstances, in the same way that for a human memorizing certain pieces of data might be helpful to one's character or one's ability to understand the world, things like the United States Constitution, principles of basic human ethics, the history of the world, things like that, could be worth memorizing.
BY MR. MILLER:
Q How does Anthropic determine what content

179

is worth memorizing?
MS. MEHTA: Objection to the form.
THE WITNESS: So, in general, we don't explicitly determine this. There have been a number of discussions internally on proposed things that we might want to have the models memorize. But as far as -- to the best of my knowledge, those instances are always meant to be representative sets where we're trying to create a generalizable technique to find those sort of pillars of human civilization rather than manually trying to carry a list of things for the model to memorize.
BY MR. MILLER:
Q Has Anthropic created a generalizable technique to find the pillars of human civilization?
MS. MEHTA: Objection to the form.
THE WITNESS: While there have been many discussions on the topic and various experiments, I wouldn't say that we were lastingly successful in any of them that I'm aware of.
BY MR. MILLER:
Q What does Anthropic consider to be the pillars of human civilization?

180

MS. MEHTA: Objection to the form.
THE WITNESS: Some things we've discussed before are very famous pieces of literature: the United States Constitution, as I mentioned; things like Abraham Lincoln's speech: Four score and twenty years ago -- seven years ago; the Declaration of Independence. These kinds of things.
BY MR. MILLER:
Q Are any of the -- withdrawn.
Are any of what Anthropic considers pillars of human civilization copyrighted?
MS. MEHTA: Objection to the form. Calls for speculation.
THE WITNESS: I don't think we have a definitive list of these things, and although we've done some experiments here and there, I wouldn't say that the lists created for those experiments were in any way meant to be complete, from what I've seen of those lists. I haven't seen copyrighted content, but -- yeah.
BY MR. MILLER:
Q What lists are you referring to?
MS. MEHTA: Objection to the form.
THE WITNESS: I'm aware of at least one

185

Q What is fine-tuning?
A So, generally, fine-tuning refers to taking a pre-trained or base model, and then adding additional training on top of it. And there are a couple of different flavors of fine-tuning, as far as we call it, inside Anthropic, involving things like supervised learning, semi-supervised learning, reinforcement learning, reinforcement learning from human feedback, et cetera.
Q Excuse me. Sorry.
Are you familiar with the term "red teaming"?
A Yes.
Q What is red teaming?
MS. MEHTA: Form.
THE WITNESS: So red teaming is a concept that many organizations use, including the United States CIA and other branches of the U.S. government. It -- at Anthropic, the way we use it is red team is -- in contrast to blue team, red team is meant to be sort of the adversarial, like trying to mimic an external party that doesn't have the organization's interests at heart, and blue team is the defenders.

186

And so red teaming, specifically, as a verb, refers to assuming this adversarial identity and then conducting attacks or -- or other tests to test the vulnerabilities of an organization, typically conducted from inside the organization, with knowledge of the parties conducting the testing.
BY MR. MILLER:
Q Are you familiar with the term "crowd workers?"
A Yes.
Q What are crowd workers?
A In Silicon Valley parlance, a crowd worker is a set of people who voluntarily do tasks that are sort of parceled out by many different kinds of companies.
Typically, there's a crowd worker aggregator platform. Some examples of these are Scale AI, Amazon MTurk, many others, and machine learning companies often make use of these services to label large amounts of data or generate data for the purposes of making models better.
Q Has Anthropic used crowd workers?
A We have used crowd worker services, yes.

187

Q For what purposes has Anthropic used crowd worker services?
MS. MEHTA: Form.
THE WITNESS: We've used it for many, many purposes. The most common is probably collecting preference data in an interface that I built four years ago called the human feedback interface, where we have instructions for these crowd workers and we direct them to interact with an AI assistant in a specific way and make commentary and judgments on the AI's outputs, and then we use that -- those outputs to assess our models and to improve the training process. And this is described in a paper that we published called Reinforcement Learning from Human Feedback [sic].
MR. MILLER: Tab 9.
(Plaintiffs' Exhibit 194, arXiv.2204.05862v1 [cs.CL] 12 Apr 2022: Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback, Bates stamped Anthropic_0000000829 through Anthropic_0000000902, was marked and presented for purposes of identification.)
THE WITNESS: Yep.
BY MR. MILLER:

188

Q Do you see in the Zoom what has been marked Plaintiffs' Exhibit 194? Is that right? 194?
A I don't see the number, but I see a Bates number ending in 829.
Q Yeah, for -- and that's correct.
MR. MILLER: For the record, what's been marked as Plaintiffs' Exhibit 194 has the Bates number on the first page Anthropic_829.
BY MR. MILLER:
Q Is Plaintiffs' 194 the research paper you were just referring to?
A Yes.
Q And it's titled Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback; is that correct?
A Sorry. Trying to scroll up.
Yes.
Q And you're one of the authors listed on this paper, correct?
A Yes.
Q Could you turn to page 65, please?
A Yes.
Q And do you see that page 65 contains images of two different documents?

225

A The paragraph starting with "Except"?
Q Yes.
Do you see that these terms of service state, Except as expressly authorized by Genius in writing, you agree not to modify, copy, frame, scrape, rent, lease, loan, sell, distribute, or correct -- create derivative works based on the service, or Genius content in whole or in part?
MS. MEHTA: Objection to the form. Outside the scope. Lacks foundation. And calls for speculation, as well as legal conclusions.
THE WITNESS: I see the writing that you're describing.
BY MR. MILLER:
Q I want to direct your attention to the next sentence, which says, In connection with your use of the service, you shall not engage in or use any data mining, robots, scraping, or similar data gathering or extraction methods.
Do you see that?
MS. MEHTA: All the same objections.
THE WITNESS: Yes, I do.
BY MR. MILLER:
Q Do you know whether the material from genius.com in the -- in the Common Crawl dataset

226

reflected in the spreadsheet we were just looking at was collected in a manner that violated Genius' terms of service?
MS. MEHTA: Objection. All the same objections.
THE WITNESS: So Common Crawl has been around for a long time, and it's totally possible that the URLs that are in that spreadsheet came from before these terms of service, which were from 2020. And it's also possible that despite Genius having these policies, that they didn't properly enforce them by implementing robots.txt.
And, in general, for scraping services, it would be technically infeasible to read and understand every single site's terms of service in advance of doing scraping.
So I would say any web content provider would know that scraping is a thing, and without getting into any legal terms or definitions, just from a technical standpoint, I think there are very reasonable explanations that some pages from -- that appeared to come from the genius.com website might appear in a Common Crawl shard.
MS. MEHTA: Counsel, you're at -- you're at time, so we need to wrap up the 30(b)(6).

227

Before we do that, we'll close that volume and mark that volume as highly confidential, attorneys' clients only, as per the protective order, and then we can start a new volume for the 30(b)(1) deposition.
THE VIDEOGRAPHER: Would you like to close now?
MS. MEHTA: Yes.
MR. MILLER: Yeah, we can close now.
You're not going to have any redirect?
MS. MEHTA: Nope.
MR. MILLER: Okay.
THE VIDEOGRAPHER: This marks the end of the deposition of Benjamin Mann, 30(b)(6).
We are going off the record at 16:03.
(The deposition was concluded at 4:03 p.m.)
* * *

228

CERTIFICATE OF REPORTER

I, April L. Crites, RPR, RMR, CRR, CSR, the officer before whom the foregoing deposition testimony was taken, do hereby certify:
That before being examined, BENJAMIN MANN was by me duly sworn to testify to the truth, the whole truth, and nothing but the truth;
That the foregoing transcript is a true and correct record of the testimony given;
That said testimony was taken by me stenographically and thereafter reduced to typewriting under my direction;
That reading and signing was requested;
And that I am neither counsel for, related to, nor employed by any of the parties to this case and have no interest, financial or otherwise, in its final disposition.
IN WITNESS WHEREOF, I have here unto subscribed my name this 24th of November, 2025.
My commission expires May 31, 2026:

April L. Crites

April L. Crites, RPR, RMR, CRR, CSR
CA Certified Shorthand Reporter No. 14493