Exhibit 4



HIGHLY CONFIDENTIAL
ATTORNEYS' EYES ONLY

Transcript of Benjamin Mann, Corporate Designee

Date: November 20, 2025
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
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

5

Exhibit 192A	Claude Support: Reporting, Blocking, and Removing Content from Claude	166
Exhibit 192B	Claude Support: I think a user is infringing my copyright or other intellectual property. How do I report it?	166
Exhibit 193	Axios, Column/Behind the Curtain: A white-collar bloodbath, May 28, 2025	170
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	187
Exhibit 195	November 3, 2023, Slack message, Bates stamped Anthropic_0000240671 through Anthropic_0000240672	194
Exhibit 196	Excel spreadsheet, hyperlinked in Exhibit 195, Bates stamped Anthropic_0000240674	197
Exhibit 197	Anthropic/hh-rlhf - Datasets at Hugging Face	199
Exhibit 198	Excel spreadsheet, 2025.11.20 Huggingface_Anthropic_Train , 160800 rows listed	203
Exhibit 199	Genius Song Lyrics & Knowledge, Genius TERMS OF SERVICE, Last Updated: June 24, 2020	223

* * *

6

A P P E A R A N C E S

ON BEHALF OF THE PLAINTIFFS, THE MUSIC PUBLISHERS:

Corey Miller, Esq.
Nicholas C. Hailey, Esq.
Michelle Gomez-Reichman, Esq.
OPPENHEIM ZEBRAK LLP
4530 Wisconsin Avenue NW
Fifth Floor
Washington, D.C. 20016
202-596-5043
corey@oandzlaw.com
mgomez-reichman@oandzlaw.com
nick@oandzlaw.com

ON BEHALF OF THE DEFENDANT, ANTHROPIC, PBC:

Sonal N. Mehta, Esq.
WILMER CUTLER PICKERING HALE & DORR, LLP
2600 El Camino Real
Suite 400
Palo Alto, California 94306
650-858-6000
sonal.mehta@wilmerhale.com

- AND -

7

A P P E A R A N C E S C O N T I N U E D

ON BEHALF OF ANTHROPIC, PBC (CONTINUED):

Joseph Taylor Gooch, Esq.
WILMER CUTLER PICKERING HALE & DORR, LLP
50 California Street
Suite 3600
San Francisco, California 94111
628-235-1002
taylor.gooch@wilmerhale.com

- AND -

Jared V. Grubow, Esq. (Via Zoom)
WILMER CUTLER PICKERING HALEK& DORR, LLP
7 World Trade Center
250 Greenwich Street
New York, New York 10007
212-295-6474
jared.grubow@wilmerhale.com

- AND -

8

A P P E A R A N C E S C O N T I N U E D

ON BEHALF OF ANTHROPIC, PBC (CONTINUED):

Sara E. Sampoli, Esq.
LATHAM & WATKINS LLP
555 Eleventh Street NW
Suite 1000
Washington, D.C. 20004
202-637-2200
saa.sampoli@lw.com

ALSO PRESENT:

Devon Hanley Cook, Esq.
Anthropic PBC

Philip Astor, Videographer
Planet Depos

* * *

Pages 21-32

21

A Okay. Thank you.
Q You have this document open in front of you?
A Yes. I'll just quickly skim it to make sure I know what it is.
Okay.
Q And does Exhibit 187 appear to be a copy of the Anthropic web page titled Introducing Claude?
A Yes.
Q And that web page is dated March 14, 2023, correct?
A The post has that date, but the capture has a much more recent date.
Q The publication of this document coincided with Anthropic publicly releasing a version of Claude to the market, right?
A Yes.
Q Could you turn to the third paragraph, please?
A Okay.
Q Do you see the paragraph begins, Claude can help with use cases including summarization, search, creative and collaborative writing, Q&A, coding, and more?

22

A Yes, I see that.
Q So these are the use cases that Anthropic chose to highlight in rolling out Claude, correct?
MS. MEHTA: Objection to the form.
THE WITNESS: These are some of the use cases that we highlighted, yes.
BY MR. MILLER:
Q And the first one -- the first use case that Anthropic highlighted is summarization, right?
A Yes, that is what the document says.
Q And Anthropic intended for Claude users to use Claude for summarization?
MS. MEHTA: Objection to the form.
And, Mr. Miller, can we have an agreement that when I object to form, that includes scope, so I don't have to separately use all those words?
MR. MILLER: Sure.
MS. MEHTA: Thank you.
Objection to form.
Go ahead.
THE WITNESS: I would say these are example use cases. I don't know about ascribing the writer's intention here. I think we were proposing these, but we wanted the market to be

23

able to sort of tell us what the utility of what this model would be. These were some of our hypotheses.
BY MR. MILLER:
Q Well, you weren't -- strike that.
Anthropic wasn't listing use cases that it did not intend users to use Claude for, right?
A I would say that we -- that these would have been acceptable and reasonable ways to use the model.
I guess, can you tell me more about what you mean by "intended"?
Q Anthropic, it expected that users would use Claude for summarization, and designed Claude to have that ability, right?
MS. MEHTA: Objection to the form.
THE WITNESS: I don't know that we explicitly designed Claude to have summarization skills, expressly, but, yes, we did expect users to use the models for summarization purposes.
BY MR. MILLER:
Q And summarization would include summarizing song lyrics, right?
MS. MEHTA: Form.
THE WITNESS: It's possible that some

24

users might choose to use Claude to summarize song lyrics, but that wasn't something that we explicit -- expressly designed into the model.
BY MR. MILLER:
Q The second use case listed here is search, right?
A Yes.
Q What does "search" mean here?
MS. MEHTA: Form.
THE WITNESS: In this context, "search" could mean a couple of different things. One is using search tools, although I don't believe we had function calling at this time. And I can explain what "function calling" means.
It could also be providing long documents and asking Claude to retrieve content within those long documents using its intelligence, but not using tools, and possibly other means.
Q So the first meaning you identified was search tools.
Does that mean responding with facts or information in response to a query?
MS. MEHTA: Form.
THE WITNESS: I guess, more specifically, what I was referring to is, if you have a service

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.

[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.

[redacted]

32

[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.

Pages 37-60

37

BY MR. MILLER:
Q Could you turn to page 13 of this document, please?
A I don't see the page number. Can you tell me the Bates number?
Q Sure. It's Bates number ending in 368.
A Okay.
Q And does this page have a header at the top that says, Anthropic's Approach?
A Yes.
Q And then there's a list of items below that, correct?
A Yes.
[redacted]
A Yes.
[redacted]
MS. MEHTA: Objection. That clearly calls for attorney-client communications and legal advice, and on that basis, I'm going to instruct the witness not to answer.
MR. MILLER: I'm not asking for the substance of any particular legal analysis of a

38

[redacted]
MS. MEHTA: That's still calling for privileged work done by the lawyers of the company, so I don't know what distinction you're trying to draw there. It's also outside the scope.
BY MR. MILLER:
Q Generally, what legal diligence on use of datasets does Anthropic conduct?
MS. MEHTA: Same objection. Same instruction not to answer. If you're asking for what diligence the lawyers of the company are doing, you're asking for privileged information.
BY MR. MILLER:
Q Does Anthropic conduct any non-legal diligence on use of datasets?
A Yes.
Q What non-legal diligence on use of datasets does Anthropic conduct?
A So I can give a couple of examples, but like I said, every dataset is sort of a special
[redacted]

39

[redacted]

40

Those are some of the examples.
MS. MEHTA: Just one question. You said -- you referenced content farms, and then you said with data that is -- did you say useful or not useful?
THE WITNESS: I said not useful.
MS. MEHTA: Thank you.
BY MR. MILLER:
Q In conducting diligence for use of datasets, does Anthropic consider the copyright status of any of the data included in the dataset?
MS. MEHTA: I'm going to again caution the witness not to reveal the contents of any communications that you had with counsel. To the extent you can answer that yes or no without revealing the discussions you had with counsel, you can answer that.
[redacted]

41

MS. MEHTA: Again, you can answer that yes or no, but do not get into any discussions that you had with counsel.

I also object to the form and scope.

THE WITNESS: Well, when we founded the company, we didn't -- we weren't sure that we ever wanted to sell a product, and so -- and we also were a tiny start-up with eight people. So we didn't have formal policies for a lot of things. I believe our first formal policy on this was in our employee handbook.

[redacted]

BY MR. MILLER:

Q When the company was formed, you said there were eight employees?

A Correct.

Q Which one of those was counsel?

MS. MEHTA: You can -- you can answer that question. Again, don't get into any communications with counsel.

THE WITNESS: Although none of those employees were counsel, we did retain outside counsel for many different kinds of matters at the

42

company at the time.

BY MR. MILLER:

Q Who was the outside counsel that Anthropic retained at that time?

MS. MEHTA: Again, you can share the name if you know, but you can't get into the substance of any communications.

Also, form and scope.

THE WITNESS: I know of at least Matan Shacham. I think I'm pronouncing his name correctly.

We'll give you the spelling afterwards.

But there may have been others that I'm not aware of.

BY MR. MILLER:

Q And at the beginning, at the time of the

[redacted]

MS. MEHTA: Okay. Again, you can answer that yes or no, but you cannot get into any communications with counsel.

THE WITNESS: As I mentioned, we didn't have formal policies at the time because we were

[redacted]

43

[redacted]

BY MR. MILLER:

Q And you said that Anthropic didn't have formal policies at that time.

Subsequent to the formation of Anthropic,

[redacted]

MS. MEHTA: Again, you can answer that yes or no. You cannot get into any communications with counsel.

THE WITNESS: I believe we do have a formal policy, but I don't know the exact phrasing of what it is.

BY MR. MILLER:

Q When was that policy adopted?

A I don't have the exact date.

Q What year was that policy adopted?

MS. MEHTA: Form. Scope.

THE WITNESS: I can't remember the exact year, but I am pretty sure it's in the employee handbook, and I believe that was furnished for

44

this case. So if you would like to furnish the document, we can check.

BY MR. MILLER:

Q Is that policy referring to -- reflected anywhere else besides the employee handbook?

A It's possible. I guess I'll note that every employee's required to read and -- and sign something that they -- that they have read and understand the employee handbook.

Q Have you read the employee handbook?

A Yes, I have.

Q Have you signed a document stating that you read and understand the employee handbook?

A I did an e-sign then, yes.

Q When did you read the employee handbook?

A I've read it multiple times, but the first time was when we first created it. In fact, I may have written -- read an early draft before it was actually, like, formally turned into company policy, but I don't remember the exact date.

I'll also note that there have been multiple revisions of the employee handbook, including a recent one.

[redacted]

45

[redacted]

MS. MEHTA: Scope. Form.

THE WITNESS: I don't think it's phrased exactly that way. I don't remember the exact phrasing. I haven't memorized the employee handbook. But something to that flavor is certainly in the handbook.

BY MR. MILLER:

Q And has the employee handbook's statement of that policy been revised or amended in any of the subsequent versions?

MS. MEHTA: Objection to form.

THE WITNESS: Can you be more specific about what -- I guess, if you want me to comment on a specific thing, it would be helpful for me to see the thing that you'd like me to comment on.

BY MR. MILLER:

Q Mr. Mann, are you familiar with the term "de-risking a dataset"?

A Yes, I am.

Q What does it mean to de-risk a dataset?

A So, earlier in my testimony, I described how Anthropic is a very empirically driven company, and when we use new techniques or

46

interventions to train models, we attempt, to the best of our ability, to validate or verify that that new technique will have the desired effect, and "de-risking" is simply another way of saying this validation process; I guess the risk being that if we were not to sufficiently validate the dataset or intervention, then we might waste a lot of resources, both human, compute, time, et cetera, on something that doesn't work.

Q Mr. Mann, are you familiar with the term "crawling"?

A Yes, I am.

Q What is crawling?

A Well, I have two small children, so there's certainly many meanings of the word "crawl." But in -- is there a specific context that you're referring to?

Q In the context of the Internet or collecting data from the Internet.

A Yes. So a web crawl, specifically, is the idea that web pages contain hyperlinks to other web pages. And so if you are trying to, for example, build an index, like a search engine might of the Internet, then you can metaphorically crawl these links, crawling from one page to

47

another. That, eventually, gives you access to large parts of the Internet.

Q Does -- are you familiar with the term "scraping" a web page or "scraping" on the Internet?

A Yes, I am.

Q What is scraping?

A Scraping, as I would colloquially describe it, would involve making a copy of the contents that are downloaded when visiting a web page.

Q Is there a difference between crawling and scraping?

A Yes. So I would say crawling is just this act of going from page to page, and it doesn't necessarily involve scraping.

So, for example, before doing a scrape, our crawlers check robots.txt to make sure that the page -- that the webmaster is allowing the scraping of those pages.

Whereas, for crawling, typically, you're -- you're simply following the links. And it might be for creating of some, like, ranking signals or connectivity graph of pages.

For some use cases, this might be interesting and important; certain academics look

48

at web graph connectivity independent of scraping.

So, yes, there's a pretty significant difference there.

Q So is it fair to say, then, that crawling is building an index or a map of the links and relationships between websites, while scraping is collecting the actual data present on the website?

MS. MEHTA: Objection to the form.

THE WITNESS: I mean, these are terms of industry, and so it could be that when some people use the word "crawl," they implicitly also imply scraping. But at a high level, I think that what you said is sufficient.

BY MR. MILLER:

Q Has Anthropic used bots to crawl the Internet?

A We explicitly say on our help website docs.anthropic.com --

THE COURT REPORTER: I'm sorry. I missed that.

THE WITNESS: Sorry. support.anthropic.com. That we do use -- we have a bot that scrapes the Internet with instructions on how to interact with it.

BY MR. MILLER:

49

Q When did Anthropic begin using bots to crawl the Internet?
A To the best of my knowledge, it was sometime around 2023.
Q Do you recall when in 2023?
A We've done some investigations on exactly when, and it was hard to say. Some of the people who were doing the work have since left the company, so we were actually unable to find some of that information. But we're fairly confident that it was sometime in 2023.
Q And has Anthropic used bots to scrape content from the Internet?
A Yes.
Q And when did Anthropic begin using bots to scrape content from the Internet?
A Again, sometime in 2023.
MR. MILLER: Now might be a time for a break?
MS. MEHTA: We have a long day ahead of us, so I would say, unless you need a break, maybe we can go another half an hour and then we can take a break then. Otherwise, we'll be here well into the night.
THE WITNESS: I'm good to keep going.

50

MR. MILLER: Okay.
BY MR. MILLER:
Q What's the name of the bot that Anthropic uses to crawl the Internet?
A I believe it's ClaudeBot, but I think we actually now have -- we have multiple user agents, but I think we only use one of them now for scraping.
Q What is a user agent?
A A "user agent" is an industry standard technical term which refers to a -- when you make an HTTP request to a website, which is how would you visit a website typically -- that's what your browser does on a computer, but you can also do it without a browser -- there is a part of that request that includes what's called a user agent, and that user agent contains metadata, such as the version of the software that you're using to make the request, and, optionally, an identifier, and we've decided to include this identifier to give webmasters the ability to identify us and contact us if there's any problem with what we're doing.
Q Has Anthropic used any other user agents besides ClaudeBot to crawl or scrape the Internet?
A I believe we have, yes.

51

Q What are the other agents besides ClaudeBot that Anthropic has used to crawl or scrape the internet?
A I don't recall the specific user agent, but I believe it has "Anthropic" in the name somewhere.
Q How many other user agents besides ClaudeBot has Anthropic used to crawl or scrape the Internet?
A It would be a small handful if -- if it's not just the two that I mentioned so far.
Q Besides the two that you mentioned, are you aware of any others?
A No.
Q How does Anthropic choose which -- which websites to crawl or scrape?
A This is a fairly complicated and evolving process. So I guess I can talk about what we do today. Is that what you're asking for?
Q When Anthropic began its crawling activities in -- sometime in 2023, how did Anthropic choose which websites to crawl or scrape?
MS. MEHTA: Objection to the form.
THE WITNESS: So, to the best of our

52

[redacted]
So Common Crawl has very large list of URLs associated with it, and we specifically wanted these different types of content.
Just for the purposes of this case, I will say that these kinds of documents typically do not contain song lyrics. It just so happened that that's kind of an artifact of what we were looking for, which is things to help with knowledge work. It wasn't something we expressly designed. We -- yeah, that's how we started scraping.
I don't know if there were other particular domain signals, but if I looked at the code, then I could be more specific.
BY MR. MILLER:
Q How do you know that the types of documents you're describing don't typically contain song lyrics?
A From my personal experience, I can't remember a time that I saw song lyrics in a PDF, other than in the document produced for this case.
Q But you didn't conduct any investigation

53

or review of the documents that were being collected to determine whether or not they contained song lyrics, did you?

A Actually, I have looked at quite a lot of training data over the course of my career, and while I wasn't specifically looking for song lyrics, that's something that I probably would have noticed over that time.

Q Did you specifically review or investigate the data collected in the scrape in 2023 you were just describing?

A I have looked at a large number of PDFs that came from the Common Crawl project over the course of my career, both at Anthropic and elsewhere, and so I have a pretty good understanding of the distribution of data that's represented here.

Q What is the Common Crawl project you referred to?

A Common Crawl is an extremely large-scale, open-access scrape of the web that is used across industry and academia for a wide variety of purposes. I believe it's a nonprofit foundation that does this and makes this data available to anybody who wants it.

54

Q Is Common Crawl project a third-party organization?

A Are you asking me if they're affiliated with Anthropic?

Q I'm just asking what it is.

A Does "third party" refer to Anthropic as first party?

Q Yes.

A So you are asking if they're affiliated with Anthropic?

Q Yes.

A No, they're not affiliated with Anthropic.

Q You said that Common Crawl is a nonprofit foundation.

Has Anthropic donated money to Common Crawl?

MS. MEHTA: Form.

THE WITNESS: I don't know.

BY MR. MILLER:

Q When Anthropic scrapes data from the Internet, what format is the data in?

A So it strongly depends on the data source. So, typically, when we scrape, we get whatever form the data is in from the hosting provider. So we make an HTTP request and then something comes back, and that's the thing we get when we scrape.

So that might be PDFs, HTMLs, CSS, JavaScript, MOV, made-up, pretend file extensions that don't actually exist, binary data, things without file extensions that we actually don't even know what they are. It could really be anything.

55

Q And when Anthropic scrapes data, does it copy it to some location on Anthropic servers?

A Typically, we store the data from our crawls in a durable cloud blob storage.

THE COURT REPORTER: Cloud blob?

THE WITNESS: Cloud blob storage.

BY MR. MILLER:

Q You referred a couple times already to robots.txt. Is robots.txt the file name for used for implementing the Robots Exclusion Protocol?

A I don't know if there's an official thing called Robots Exclusion Protocol, but that sounds like the thing that we do, yes.

Q And Robots Exclusion Protocol, is it fair to say, is a standard used by websites to indicate to visiting web crawlers or scrapers which portions of the website the crawlers or scrapers are allowed to visit and use?

56

MS. MEHTA: Form.

THE WITNESS: Are you reading from a specific source? Are -- is that a quote?

BY MR. MILLER:

Q I'm asking if you agree with that description.

A It sounds roughly accurate, but I guess knowing the context of the quote would be helpful to know if that's, like, an official W3C thing.

I guess I'll say I think robots.txt are included in WC3 standards. And so there is an official definition, but I don't know if the one you gave is the official one or not.

Q When Anthropic crawls and scrapes the Internet, does it currently honor robots.txt instructions?

A We attempt to honor robots.txt, yes.

Q Has Anthropic always honored robots.txt instructions?

A There have been periods of time when we didn't due to bugs, and there have also been periods of time in the very early days when we were only crawling the specific document types that I described, where, to the best of our knowledge, it seems possible that we were not

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

[redacted]

58

[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?

Pages 69-80

69

A That's right.
Q And for both of those changes made on
[redacted]
MS. MEHTA: Form.
THE WITNESS: It does say that robots.txt support is basic, yes.
BY MR. MILLER:
Q Do you know what the "basic" means in this context?
[redacted]
Q And as of May 11, 2023, was Anthropic honoring general robots.txt disallow instructions?
A Yes. If we said "basic" here, that would have included not just our user agent, but also general robots directives, such as disallow star.
[redacted]
A Yes.

70

Q And it indicates the date for that change was August 8, 2023?
A Yes.
Q And in cell B-9, the document indicates,
[redacted]
A Yes.
Q And the document indicates the change for that date is August 11, 2023?
A Yes.
Q What is crawl-delay support?
A My understanding is that in the robots.txt protocol, there is a crawl-delay directive, which indicates how frequently a crawler can scrape a given page, and this is to prevent abuse scenarios, like hitting the same web page hundreds or thousands of times in a single day.
My guess is that that just wasn't relevant for us because we were scraping at such a small scale. However, as we scaled up, it became more important for us to respect this because there would be a chance that we would go over the maximum set by the webmaster.
Q So prior to August 11, 2023, Anthropic was not honoring crawl-delay instructions?

71

MS. MEHTA: Objection to the form. Misstates the testimony.
THE WITNESS: It's possible that we did honor those directives, but not from a technical enforcement standpoint.
So I guess, to directly answer your question, it's unclear -- like, it would be very hard for me to know whether or not we were following those call directives implicitly or not.
BY MR. MILLER:
Q Okay. So the last entry in this document is dated August 11, 2023, correct?
A That's right.
Q Have there been any changes to Anthropic's practices with respect to robots.txt since August 11, 2023?
MS. MEHTA: Objection to the form.
THE WITNESS: So we don't use the repository that is specified in this document
[redacted]

72

[redacted]
could look at the code, too, to verify.
MR. MILLER: Is this a good point for a break?
MS. MEHTA: Sure.
THE VIDEOGRAPHER: We are going off the record at 10:38.
(A brief recess was taken.)
THE VIDEOGRAPHER: We are back on the record at 10:54.
THE WITNESS: Before we get started, can I make a clarification for the record?
MR. MILLER: Go ahead.
THE WITNESS: So, before, in the earlier period, I mentioned that I couldn't quite remember the date when we started the Crabby project,

73

crawling the Internet. After thinking about it [redacted]

So that's a more exact time period than I was able to recall earlier.

BY MR. MILLER:

Q Mr. Mann, in the Zoom, you should see Plaintiffs' Exhibit 129, which we were looking at before the break, and we discussed the last dated entry in here is for August 11, 2023.

Are you aware of any scraping activities that -- or, crawling activities that Anthropic engaged in prior to August 11, 2023, that are not reflected in this document?

MS. MEHTA: Objection to the form and asked and answered.

THE WITNESS: Your question was prior to August 11?

BY MR. MILLER:

Q Yes.

MS. MEHTA: Same objections.

THE WITNESS: There may have been other

74

[redacted]

BY MR. MILLER:

Q And when you say "a limited scale," what do you consider a limited scale?

MS. MEHTA: Objection to the form.

THE WITNESS: A limited scale would be, for example, a single website. Yeah.

BY MR. MILLER:

Q Okay.

MR. MILLER: Is this 189?

MS. GOMEZ-REICHMAN: Yes.

(Plaintiffs' Exhibit 189, Anthropic, PBC, employee handbook, effective April 15th, 2021, Anthropic_0000429433 through Anthropic_0000429458, was marked and presented for purposes of identification.)

BY MR. MILLER:

Q In the Zoom, you should see what's being marked as Plaintiffs' Exhibit 189, which bears the Bates number on the first page of Anthropic_429433.

75

Do you have Plaintiffs' Exhibit 189 open in front of you, sir?

A The scrolling isn't working. I can only scroll, like, huge amounts.

THE WITNESS: Any tips?

THE VIDEOGRAPHER: Are you able to hit the scroll on the far-right side? Nothing?

Does the down arrow do one page at a time?

THE WITNESS: Oh, now the arrow's working.

THE VIDEOGRAPHER: Okay.

THE WITNESS: Okay. Okay. It's working now. It stopped working.

Okay. Bates number 429433?

BY MR. MILLER:

Q Yes.

A Okay.

Q And what's been marked as Plaintiffs' Exhibit 189 is a copy of Anthropic's employee handbook, dated April 15, 2021, correct?

A Sorry. The scroll keeps breaking.

Yes, it is dated as effective April 15, 2021.

Q Is this the first version of Anthropic's employee handbook or was there a version that predated this one?

76

MS. MEHTA: Objection to the form.

THE WITNESS: My recollection is that this is the first official version of the Anthropic employee handbook, although I am certain that there are many drafts that were created before it was made official.

BY MR. MILLER:

Q And when you were testifying earlier about Anthropic's employee handbook, this is the document you were referring to, correct?

A This is one of the couple of editions that I was referring to, yes.

Q Can you please identify where in this document it instructs Anthropic employees to consult with counsel before using a new dataset for a training run?

MS. MEHTA: Form.

THE WITNESS: I'll look through.

[redacted]

77

[redacted]
BY MR. MILLER:
Q And to be clear, which page of the agreement are you looking at -- or, excuse me -- the handbook are you looking at?
A This would be page 18, which has Bates number ending in 9451.
Q And there's -- on page 18, there's a heading that says, Proper Use, correct?
A Yes.
Q And then under that -- under the heading [redacted]
A Yes.
[redacted]
A Correct.
Q Yes. And then there's a list of five items.
A Yes.
Q And you were pointing specifically at item number 3 on that list, correct?
A That's right.
Q And item number 3 states, as you were

78

[redacted]
A Yes.
Q And this is the provision you're pointing to as instructing employees to consult with counsel before using a new dataset in a model training run?
MS. MEHTA: Objection to the form.
THE WITNESS: This is likely one of the things I was remembering.
BY MR. MILLER:
Q This item number 3 doesn't say anything about datasets or training runs, does it?
MS. MEHTA: Objection to the form.
THE WITNESS: I guess copyrighted data would certainly, for anybody working on model training, immediately bring to mind datasets.
BY MR. MILLER:
Q Does using a dataset for a model training run involve sending or receiving copyrighted materials?
MS. MEHTA: Objection to the form.

79

[redacted]
lawyer, so how exactly the legal definition of "copyrighted materials" is defined would be outside my purview.
BY MR. MILLER:
Q How are copyrighted materials sent or received in connection with using a dataset for a model training run?
MS. MEHTA: Objection to the form.
THE WITNESS: Sorry. I lost the -- let me get back there.
Can you ask your question again?
BY MR. MILLER:
Q Yes.
How are copyrighted materials sent or received in connection with using a dataset in a model training run?
MS. MEHTA: Objection to the form.
THE WITNESS: So in the process of scraping the web, for example, or using materials that are contained in some -- in a dataset, like one that Common Crawl makes available, it is possible that there is copyrighted material within

80

those datasets. And in order to use a dataset like Common Crawl, an employee would have to obtain that data from Common Crawl by making a request to Common Crawl's hosting servers, which, in their case, is on Amazon S3, which is, again, a blob storage provider, and that, I guess, I would consider receiving.
And then in terms of sending, you could consider taking that downloaded information and putting it on Anthropic servers to be a form of sending that data as well.
But, again, what data inside those datasets might be copyrighted or not is a -- is often undefined, and so any employee working on this kind of thing would understand that even if it wasn't sort of guaranteed that any copyrighted materials would be there, that they should get -- that they should get authorization.
BY MR. MILLER:
Q And this provision in the employee handbook requires employees not just to consult with counsel, but to get prior authorization before engaging in sending or receiving copyrighted materials, correct?
MS. MEHTA: Objection to the form. Vague.

Pages 85-108

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 [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.

93

THE WITNESS: Could have included a lot of people. Some people I would have expected to give these kinds of approvals would have included Nick Joseph, Pranay Sangani -- Pranay Sangani, more recently Sambavi, who leads the tokens team; I can't remember her last name. I have given approval at least once here and there as a cofounder and as a leader of the team that produce data.

Those are some off the top of my head.

BY MR. MILLER:

Q Anyone else?

A There are likely others.

MS. MEHTA: Objection to the form. Scope.

THE COURT REPORTER: I missed your answer.

THE WITNESS: There are likely others.

BY MR. MILLER:

Q Who else?

MS. MEHTA: Same objections.

THE WITNESS: Those are the people who I'd be confident of. Off the top of my head, I don't specifically remember others.

BY MR. MILLER:

Q You said that you have given approval at least once here and there for use of new dataset in model training, correct?

MS. MEHTA: Form.

THE WITNESS: Yes, that's what I said.

BY MR. MILLER:

Q What dataset did you approve for use in model training?

MS. MEHTA: Form.

THE WITNESS: I guess, to be clear, the period of time in which I was directly involved in training data was something like the first two years of the company, and then a roughly one-month period in -- in 2024, and since then, I haven't been tightly involved in those -- those things, although I did serve as a mentor for the leader of the tokens team for some time.

So between now and then, it -- I wouldn't remember. I mean, I don't remember, unfortunately, yeah.

Q Do you recall whether you approved the use of more than one dataset or just one?

MS. MEHTA: Form.

THE WITNESS: I don't recall.

BY MR. MILLER:

Q Are authorizations to use a new dataset in model training recorded or documented in any way?

94

95

MS. MEHTA: Objection to the form. Scope.

THE WITNESS: We don't have a formal documentation process that I'm aware of. We do have a code review process. And as I mentioned, [redacted]

So we do have a form of approval, but not exactly the kind you're describing.

BY MR. MILLER:

Q Did the data Anthropic used to train Claude include the plaintiffs' song lyrics?

MS. MEHTA: Objection to the form.

THE WITNESS: I don't have a way of knowing that.

BY MR. MILLER:

Q Did the data Anthropic used to train Claude include the works-in-suit in this case?

MS. MEHTA: Objection to the form.

THE WITNESS: Sorry. Can you clarify the difference between this question and the last one? I just want to make sure I'm answering correctly.

BY MR. MILLER:

Q Yeah.

96

So the first question, I asked about publisher song lyrics as a general matter. The second question, I asked specifically about the lyrics to songs that are at issue in this case.

Does that make sense?

A Yes.

Q Okay. So why don't I ask the two questions again, and then you can answer them, and we'll have a clear record?

Did the data Anthropic use to train Claude include publisher song lyrics, generally?

MS. MEHTA: Objection to the form.

THE WITNESS: It's likely that some lyrics from some publishers' songs were included in Anthropic's model training.

BY MR. MILLER:

Q Did the data Anthropic used to train Claude include the works-in-suit that are at issue in this litigation?

MS. MEHTA: Objection to the form. Vague.

THE WITNESS: So I believe that the works that are cited in the suit, we got a PDF with the specific lyrics that are -- that are being claimed.

I don't know if those specific lyrics in

97

their specific presentation as given to us in the PDF matched exactly what was in our training data. And from a technical standpoint, it would be quite difficult to verify that.
BY MR. MILLER:
Q Does Anthropic have the technical capability to exclude publisher song lyrics, generally speaking, from its training data?
MS. MEHTA: Objection to the form. Scope. Vague.
THE WITNESS: I would say, in general, from a technical standpoint, it is very difficult to exclude any specific training data from our -- from our corpus, not just for us, but for anybody doing large-scale training, because of the sheer scale of data involved, the amount of processing that we go through, and the complexity of what the definition of a given piece of content is.
And so I would say from a -- from a current, state-of-the-art standpoint, it is not technically feasible to reliably remove specific data.
BY MR. MILLER:
Q If a content owner requests it, will Anthropic remove that content owner's data from

98

datasets Anthropic uses to train models?
MS. MEHTA: Objection to the form.
THE WITNESS: I guess it would depend on the details of what that content owner is requesting.
(Plaintiffs' Exhibit 190, Defendant Anthropic PBC's Amended and Supplemental Responses and Objections to Plaintiffs' First Set of Interrogatories (Nos. 9, 11, 12, was marked and presented for purposes of identification.)
BY MR. MILLER:
Q You should see in the Zoom what has been marked as Plaintiffs' Exhibit 190.
MS. MEHTA: Does he have control of the document?
THE WITNESS: Oh, this one doesn't have a Bates number.
BY MR. MILLER:
Q No.
Does Exhibit 190 appear to be a copy of Defendant Anthropic PBC's Amended and Supplemental Responses and Objections to Plaintiffs' First Set of Interrogatories, numbers 9, 11, and 12?
A Yes.
Q Could you go to the last page, please?

99

A Certificate of service?
Q Yes.
A Yes.
Q And do you see, on this page, it states that the above document was served on November 12, 2025?
A I see that.
Q Could you go to page 8, please?
A Okay.
Q Scroll down to the middle of the page.
A Okay.
Q Oh, sorry. I guess it would be page -- the page with the number 8 on it, which would be 10, I think, in the PDF.
A Okay.
Q I'd like to direct your attention to the second paragraph on this page.
A Okay.
Q The one -- do you see there that this document states, Additionally, if a company reaches out to remove certain discrete corpora from Anthropic's training datasets, Anthropic will consider that request and may work with the company to exclude that data from new training runs that have not commenced.

100

Do you see that?
A Yes.
Q Do you agree with that statement?
MS. MEHTA: Form.
THE WITNESS: Yes. We will certainly consider the request and may work with the company.
BY MR. MILLER:
Q Have any companies or individuals reached out to Anthropic, requesting to remove certain discrete corpora from Anthropic's training datasets?
MS. MEHTA: Objection to the form.
THE WITNESS: Yes.
BY MR. MILLER:
Q How many have done so?
[redacted]
Q Okay. And you're relying on a -- for that answer on a document that you brought with you to the deposition?
A That my counsel has brought for me, yes.
Q We'll mark that document as Exhibit 190 -- 191, please.
MR. MILLER: Counsel, do you have that

101

document?
MS. MEHTA: Yes. Let me try to find it.
(Plaintiff's Exhibit 191, Ben Mann 30(b)(6) Topics 47 and 50 - Communications with Rights Holders, was marked and presented for purposes of identification.)
BY MR. MILLER:
Q And what has been marked as Plaintiffs' Exhibit 191 is a two-page document titled Ben Mann 30(b)(6), topics 46 and 50, Communications with Rights Holders; is that correct?
A Yes.
Q And this -- and Exhibit 191 identifies 10 organizations, correct?
A That's my count, yes.
Q Has Anthropic removed -- withdrawn.
[redacted]
MS. MEHTA: Objection to the form. Scope. Calls for speculation.
THE WITNESS: I haven't read all of these, but I believe they've all been produced, so if you want to look at any in particular, or all of them,

102

[redacted]
MS. MEHTA: Objection to the form. Outside the scope.
[redacted]
MS. MEHTA: Objection to the form. Scope.
[redacted]
BY MR. MILLER:
Q So I'll ask the question again.
[redacted]

103

MS. MEHTA: Objection to the form. Outside the scope. And vague.
[redacted]
BY MR. MILLER:
Q Do you see, towards the bottom of page 1,
[redacted]
A Yes.
[redacted]
MS. MEHTA: Objection. Outside the scope. Form.
THE WITNESS: Well, I -- actually, I want
[redacted]
So just to make sure I'm being accurate on the details of the question, can you ask your

104

question again?
BY MR. MILLER:
[redacted]
MS. MEHTA: Objection. Outside the scope. Form.
THE WITNESS: I would better be able to answer that question if we reviewed the Anthropic
[redacted]
BY MR. MILLER:
[redacted]
MS. MEHTA: Objection to the form. Outside the scope.
THE WITNESS: I did review some of these letters. I don't recall reviewing that one in particular.
BY MR. MILLER:
[redacted]

105

[redacted]

MS. MEHTA: Objection to the form. Asked and answered.

[redacted]

MS. MEHTA: Objection. Misstates the testimony. Form.

[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

[redacted]

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

Pages 113-120

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.
[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

117

not behind a paywall, all of which we respect, then those things would incidentally get picked up in our crawl and end up in the training data.
And so, similarly, we may also have the styles of Nike shoes that exist. I also don't particularly care whether those pieces of data are in our training datasets, but it would be quite difficult for us to go filter those out, and it doesn't make a huge negative impact on the models' capabilities for the things we care the most about. And so we haven't done the effort to filter them out.
BY MR. MILLER:
Q A minute ago, you said that there's enough content in the different kinds of data about songs that Claude could have a strong understanding of what song lyrics are and be able to write new ones, even if it didn't have -- even if Claude didn't have song lyrics in its training data.
That was your testimony?
A Yes, that's what I said.
Q Has Anthropic tested that hypothesis?
MS. MEHTA: Objection to the form.
You can answer that to the extent that you know the answer outside of any attorney

118

communications or work product.
THE WITNESS: So, in general, we have done work studying generalization of models from training data in a couple of different ways. And, also, researchers outside of Anthropic have studied large language model in generalization, and we find that with the right training setup, generalization is not only possible, but it occurs frequently.
Internally to Anthropic, we've seen evidence of cross-language knowledge transfer. So for example, if a fact is only expressed in one language, and now you ask a question in a new language, Claude is able to express that new fact without ever having seen it in that other language.
Also, with skills like between math and computer science, that -- there is a convex relationship between those skills, even though they don't explicitly overlap from a contents standpoint. So we do have evidence that generalization in large language models is possible.
BY MR. MILLER:
Q Anthropic hasn't specifically tested

119

whether removing song lyrics from the model's training data affects the model's ability to understand what song lyrics are or write new lyrics, has it?
MS. MEHTA: Objection to the form.
And I'll give you the same caution.
THE WITNESS: I don't think I can answer that without divulging attorney-client privilege.
BY MR. MILLER:
Q Apart from any work conducted in connection with this litigation, which I don't want you to answer about, apart separately from that, Anthropic has not done -- Anthropic hasn't specifically tested whether removing song lyrics from the model's training data affects the model's ability to understand what song lyrics are or write new song lyrics; is that correct?
MS. MEHTA: Form.
THE WITNESS: I don't believe we've done such tests outside of anything that might have done it -- that might divulge attorney-client privilege. And I guess, furthermore, I'm not sure why we would do that.
BY MR. MILLER:
Q How many song lyrics has Anthropic used to

120

train Claude?
MS. MEHTA: Objection to the form. Scope.
THE WITNESS: I don't think anyone has a reliable method of knowing that, even from a technical feasibility standpoint.
Sorry.
MS. MEHTA: We've been going -- I'm just checking the time. We've been going for about an hour and 10 -- 12 minutes. Do you want to -- do you want to finish this document and take a break, or do you want to take a break now?
MR. MILLER: I think we should keep going.
MS. MEHTA: Okay. Let's take a break in about five minutes, if we could.
MR. MILLER: Well, I'd like to go a little bit longer. I think we -- when does -- do you know when the lunch will be here?
MS. MEHTA: I don't, but I'd like to go to the bathroom. So if you want to go another five minutes or so, we can, and then we'll -- we'll need to take a bio break.
MR. MILLER: Okay. You can just use the bathroom. We can take a break now.
MS. MEHTA: Okay. Should we take 10?
MR. MILLER: Sure.

Pages 125-128

125

datasets if a company submits a takedown pursuant -- a takedown notice pursuant to the DMCA policy?
MS. MEHTA: Objection to the form. Misstates the testimony.
THE WITNESS: In the interrogatory, we state that we will consider removing such data if it is technically feasible and we will work with the companies who request it to try our best to comply with those requests. There may be other avenues than DMCA. But, yes, we -- we try to be a good public steward in general.
(Plaintiffs' Exhibit 179, Excel spreadsheet, Struggle sheet + Orowa edit, Bates stamped Anthropic_0000458491, was marked and presented for purposes of identification.)
BY MR. MILLER:
Q If you turn to the Zoom, you should see what has been previously marked as Plaintiffs' Exhibit 179, which bears the Bates number Anthropic_458491.
Do you have that document open in front of you?
A Yes.
Q Do you recognize this document?

126

A I'll take a minute to skim it.
Okay. I've skimmed the document.
Q Do you recognize this document?
A I don't recall seeing it.
Q Does this document appear to be measuring Claude's capabilities in alignment on a number of different criteria against two other AI models?
MS. MEHTA: Objection to the form. Scope.
THE WITNESS: That could be a fair characterization.
BY MR. MILLER:
Q And the two other models included in the analysis here appear to be InstructGPT and ChatGPT; is that correct?
MS. MEHTA: Form and scope.
THE WITNESS: Yes.
BY MR. MILLER:
Q And both of those -- withdrawn.
Both InstructGPT and ChatGPT are developed by OpenAI, correct?
MS. MEHTA: Same objections.
THE WITNESS: Yes. OpenAI previously developed models with those names.
BY MR. MILLER:
Q Why was Anthropic comparing Claude's

127

performance to the performance of OpenAI's models?
MS. MEHTA: Objection to the form. Outside the scope.
THE WITNESS: This probably would have been a competitive analysis to try to understand where Claude was strong and where it was relatively weak compared to our competitors.
BY MR. MILLER:
Q Can you please go to the Harmlessness tab?
A Yep.
Q And this tab reads, The harmlessness of responses by the three models being evaluated to a range of different prompts, correct?
MS. MEHTA: Objection to the form. Foundation. Speculation. And scope.
THE WITNESS: I'm taking a look.
I don't know specifically how these were evaluated, but they do appear to be in response to a prompt: What are the three models' responses?
BY MR. MILLER:
Q And in doing this kind of evaluation, Anthropic would only include prompts that it thought were important to understanding the relevant models' capabilities, right?
MS. MEHTA: Objection to the form. Calls

128

for speculation. And outside the scope.
THE WITNESS: I don't know how these prompts were determined.
BY MR. MILLER:
Q Do you know who created this document?
MS. MEHTA: Objection to the form. Outside the scope.
THE WITNESS: I'm sure Orowa Sikder was involved, given his -- mention of the name on the other page, but beyond that, I don't know who was involved.
BY MR. MILLER:
Q Do you see, on this tab, there's a color coding?
A Yes.
Q Do you know what the various colors represent?
MS. MEHTA: Objection to the form. Outside the scope. Calls for speculation.
THE WITNESS: Let me take a look here.
It appears that the colors are the document authors or some other entity's judgment of the level of harmfulness of the response, where green indicates the response is fine, yellow indicates some problematic answer, and red is very

Pages 133-152

133

evaluation is for that definition of "harmfulness."
Whether we have some opinion on copyright would be independent of this evaluation.
BY MR. MILLER:
Q Anthropic's definition of "harmfulness," for the purpose of evaluating its models' performance and training its models, does not include compliance with copyright law?
MS. MEHTA: Objection. Misstates the testimony. Asked and answered. Outside the scope.
THE WITNESS: In this particular spreadsheet by an unknown author, I would say this was not considering copyright to be inside the scope of harmlessness.
But certainly, we have taken great lengths to not output copyrighted content in a very elaborate, multilayered system that stretches from the first pre-training in our models all the way to our extensive guardrail systems, including our CCS system, our Prompt Shield system, and -- and others.
BY MR. MILLER:
Q What lengths does Anthropic go to in the

134

pre-training of its models to prevent the models from outputting copyrighted content?
MS. MEHTA: Form.
THE WITNESS: This is part of our respecting of robots.txt. So if the content creator doesn't want their content to be output by the model, then it's up to them to follow the instructions on our website, or industry best practices, to prevent the scraping of that content by any scraping bot.
BY MR. MILLER:
Q Does Anthropic exclude -- withdrawn.
So we talked a little bit earlier about Anthropic's practices with regards to robots.txt compliance.
Does Anthropic exclude from its training data any data that has -- in its possession that has been collected without compliance with robots.txt protocols?
MS. MEHTA: Objection to the form.
THE WITNESS: Can you say more by what you mean "without compliance"?
BY MR. MILLER:
Q Sure.
If Anthropic learns that some data that

135

it's -- in its possession was collected even though there was a robots.txt instruction telling -- instructing for that data not to be collected, will Anthropic then remove that data from its training datasets?
MS. MEHTA: Objection to the form. Assumes facts not in evidence.
THE WITNESS: I think, if we had evidence that we mistakenly scraped some data that was contemporaneously disallowed by a robots.txt, then we would work to make sure that that didn't happen again, and see if it was feasible to remove that data from future training. But specifically, I'm not aware of cases of that happening.
BY MR. MILLER:
Q Well, Anthropic collected a significant amount of data via crawling and scraping before it implemented robots dot -- robots.txt compliance, correct?
MS. MEHTA: Objection to the form. Assumes facts not in evidence. Misstates the prior testimony.
THE WITNESS: To the best of our knowledge, we believe we did some crawling without robots.txt implemented, but we don't know

136

specifically what data was involved there, despite our efforts to figure that out. And we do know that when we started crawling in earnest with a significantly re sourced team, that we had robots.txt protections from day one.
BY MR. MILLER:
Q Has Anthropic ever tried to remove from its training data the data it collected before it was complying with robots.txt protocols?
MS. MEHTA: Objection to the form. Scope. Misstates the prior testimony. And assumes facts not in evidence.
THE WITNESS: Well, given that we don't even have a good handle on what data we did collect at that time, and what happened to it, I think it would be very hard to take efforts to remove it. So I would say that's -- that wouldn't have been technically feasible.
BY MR. MILLER:
Q What is Claude's constitution?
A Claude's constitution is part of a technique that Anthropic invented called Constitutional AI. And the constitution itself is a set of natural language principles that we have published on our website that describe what rules,

137

ethical principles, et cetera, Claude is expected to abide by.

And we have a training process that we use to bake those principles into the model so that at runtime, it tries its best to obey these principles. And we've subsequently done a lot of improvements on these techniques.

It's sort of an evolving society-wide conversation that we hoped to kick off around what the ethical principles of large language models should be.

Q And who wrote Claude's constitution?

A A couple of different researchers collaborated on it. I don't know exactly who wrote which parts, but the people involved would have been listed on the paper that we wrote describing the techniques.

And many of the parts of the constitution were borrowed from publicly available sources, like the U.N. Declaration of Human Rights, Apple's terms of service, I think, or some Apple principles document, and a couple of other places, as well as some principles that we generated ourselves, because we believe they are important.

Q Is it important that Claude's constitution

138

be public?

MS. MEHTA: Objection to the form.

THE WITNESS: Part of our goal for the Constitutional AI project was to create more transparency in what values and principles our models would uphold to create better trust for people who are using those models, and to try to inspire other people in the industry to similarly make these principles transparent, because it -- far be it for a small group of people in San Francisco to hold the keys to the values and ethics that what, in the future, we believe, will be some of the most powerful pieces of technology that have ever been created.

BY MR. MILLER:

Q In the Zoom, you should see what has previously been marked as Plaintiffs' Exhibit 84.

A Mm-hmm.

(Plaintiffs' Exhibit 84, ANTHROP\C, Announcements: Claude's Constitution, May 9, 2023, Bates stamped Anthropic_000000485 through Anthropic_000000498, having previously been marked, was presented for purposes of identification.)

BY MR. MILLER:

139

Q Do you recognize the Plaintiffs' Exhibit 84?

A I'll just take a minute to skim this document.

Okay. I've skimmed the document.

Q Do you recognize Plaintiffs' 84?

A Yes, I do.

Q This is a -- Plaintiffs' Exhibit 84 is a copy of the website post publishing Claude's constitution on Anthropic's website, correct?

A Yes, this appears to be a copy of our blog post announcing our constitution.

Q And it's dated May 9, 2023, correct?

A Yes, that's right.

Q And the constitution itself begins on page 6, right?

A Checking.

Yes.

Q Does Claude's constitution include a principle regarding copyright infringement?

MS. MEHTA: Objection to the form. And vague.

THE WITNESS: Yes. I -- I would say that the 11-17 on page -- PDF page 7, which states, Please choose the response that is most respectful

140

of everyone's privacy, independence, reputation, family, property rights, and rights of association -- I think that could be construed to refer also to copyright.

BY MR. MILLER:

Q Would the model being trained on that principle understand that principle to mean that the model should not output copyrighted content?

MS. MEHTA: Form.

THE WITNESS: I would be speculating in answering that question. However, I can tell you that it -- around this time, we did have a specific constitutional principle around copyright that is not included in this list.

As I mentioned, our constitution was evolving, and although this is the early version of it, there's another version that does explicitly talk about copyright.

BY MR. MILLER:

Q Has Anthropic ever published a revised version of the constitution, of Claude's constitution?

MS. MEHTA: Form.

THE WITNESS: I believe there are more constitutional principles in the paper, but I

141

don't know whether those are different from these.
And we also did another paper called Collective Constitutional AI in which we crowd-sourced constitutional principles from people all over America, and I believe other countries, and then tried to figure out what was different between the principles that those people came up with and the ones that we had suggested.
And so it's possible that there's some more information in that paper, as well. But beyond that, I'm not aware of any specifics.
BY MR. MILLER:
Q So the -- the constitution for Claude, as reflected in people's -- withdrawn.
The -- Claude's constitution, as reflected in Plaintiffs' Exhibit 84, does not reflect the current version of Claude's constitution that Anthropic uses to train Claude?
MS. MEHTA: Objection to the form. Misstates the testimony. Asked and answered.
THE WITNESS: This blog post talks about our first constitution, and I have every reason to believe that it was accurate at the time of publishing. But as I mentioned, our constitution, and also our more general techniques of creating

142

Claude's character and inculcating the values that we believe it should have, have evolved over time, and Constitutional AI is no longer the primary way that we do this work.
BY MR. MILLER:
Q When you say "do this work," what do you mean?
MS. MEHTA: Form.
THE WITNESS: I guess, early in my last statement, I mentioned things like Claude's character and the moral and ethical principles that it would hold -- uphold. That's the -- baking those into the model is the work that I'm referring to.
BY MR. MILLER:
Q So Anthropic's techniques of creating Claude's character and inculcating the values that Anthropic believes it should have no longer primarily rely on Constitutional AI?
MS. MEHTA: Objection to the form. Misstates the testimony.
THE WITNESS: Constitutional AI would certainly be a component, but, yes, we now have many other techniques that we use to make sure that Claude is acting in the way that we hope.

143

BY MR. MILLER:
Q Do those other techniques rely on a fundamental set of principles, like those reflected in Exhibit 84?
MS. MEHTA: Form.
THE WITNESS: We do have a set of principles, but they differ from the constitutional principles in some ways. But I would need to look at our code to give you the details.
BY MR. MILLER:
Q Has Anthropic ever published that set of principles?
MS. MEHTA: Objection to the form.
THE WITNESS: We've published many sets of principles, but I don't know if the specific one currently being used in production is -- has been published.
BY MR. MILLER:
Q Has Anthropic ever published anything stating that it no longer uses the Constitutional AI principles reflected in Plaintiffs' 84?
MS. MEHTA: Objection to the form. Misstates the prior testimony.
THE WITNESS: We do -- do still use these

144

principles.
BY MR. MILLER:
Q But not the Constitutional AI -- withdrawn.
But not, primarily, the Constitutional AI technique?
MS. MEHTA: Objection to the form. Misstates the prior testimony.
THE WITNESS: We do still use these [redacted]
BY MR. MILLER:
Q Does Anthropic disclose what those are?
MS. MEHTA: Objection to the form.
THE WITNESS: We have done things like podcasts with our researchers, Amanda Askell and Stuart -- I can't remember his last name -- that talk about how we do our character training, and I'm sure they give examples of what they mean when they say "character."
So I can't say for sure if it's in writing anywhere, but certainly, we have talked publicly

145

about our new techniques.
BY MR. MILLER:
Q Is there somewhere on Claude's website or other where -- or anywhere else where -- excuse me. Withdrawn.
Is there somewhere on Anthropic's website where Anthropic has published a list of the principles it currently uses for training Claude's character and values?
MS. MEHTA: Objection to the form.
THE WITNESS: I am not aware of such a list, although it might exist.
BY MR. MILLER:
Q Why isn't that list published?
MS. MEHTA: Objection to the form. Assumes facts not in evidence.
THE WITNESS: Our research is constantly evolving, and we want to make sure that when we announce different research and -- and major updates, that it is relevant to the people who will hear and understand the things that we're saying. I think, if we were to do some publication every time anything changed, it would be very noisy. That's one of the reasons that we have not published.

146

There may be other reasons, but I would have to speak with the team doing that work to understand what their research roadmap is, what major breakthroughs they expect, what their plans are for future publications, to fully answer your question.
BY MR. MILLER:
Q Well, a minute ago, you talked about the importance of creating more transparency in the values and principles your models uphold. How is -- and in broadening the conversation beyond just a small group of people in San Francisco.
How are other people outside of Anthropic supposed to participate in that process if Anthropic doesn't publish the principles it's using to train its models' values?
MS. MEHTA: Objection. Misstates the testimony. Form.
THE WITNESS: The thing I would say here is, these principles in this document are still active and we still use them. They are core to Claude's personality, I am sure.
To give a very trivial example of something else that I'm pretty sure is not in [redacted]

147

[redacted]
That's not interesting. I'll just come right out and say that. I don't think that needs to be a society-wide conversation about what ethics and moral principles an AI agent should have. Many of the character amendments that we've made over the years have that granularity of difference.
And so we do strongly believe in transparency and a society-wide conversation. We're continuing to conduct research on how to make these things good for people. We recently published a paper on detecting bias in language models, which is something that I would certainly expect to be an important thing in aligning language models with human values and ethics. So I think there is strong evidence that we continue to uphold transparency and a society-wide conversation.
But I also trust our researchers and our communications team to choose which of our research seems interesting and valuable for the public to know about, rather than publishing, exhaustively, every change that we might make to

148

Claude's character.
BY MR. MILLER:
Q You said that Anthropic has developed other techniques besides Constitutional AI training to train Claude -- or, inculcate Claude's character and values.
Do those other techniques have a name?
MS. MEHTA: Objection to the form.
THE WITNESS: There are two that I'm aware of, sitting here today. One is called character training.
BY MR. MILLER:
Q Mm-hmm.
[redacted]
Q Does character training for Claude include any training with regard to copyright infringement?
MS. MEHTA: Objection to the form. Outside the scope.
[redacted]
BY MR. MILLER:

149

[redacted]

MS. MEHTA: Same objections.
THE WITNESS: It -- I think, to answer that question effectively, I'd have to describe how the character training works. Should I? It's kind of a complicated and involved process. Should I explain that?
BY MR. MILLER:
Q Sure.
MS. MEHTA: Same objections.
Go ahead.
[redacted]

150

[redacted]

BY MR. MILLER:
Q In the reinforcement learning environment, what counts as a reward? Like, if I'm training my dog and I want to reward her for good behavior, I give her a treat because she likes treats. But models are not -- don't necessarily want treats.
So how do you -- how do you reward a model for choosing a preferred outcome?
A Yeah. So the -- the reward is determined by the preference model that I described previously.
Q But what counts as a reward to a model?
MS. MEHTA: Objection to the form.
[redacted]

151

[redacted]

BY MR. MILLER:
Q So you described how character training works.
[redacted]
MS. MEHTA: Objection to the form.
[redacted]
BY MR. MILLER:
Q Yes.
[redacted]

152

[redacted]

Q Is there a specific trait regarding reproduction of copyrighted information on which the model is trained?
MS. MEHTA: Objection. Asked and answered.
THE WITNESS: I believe there are such traits, but I wouldn't be able to tell you specifically what they are without looking at the source code.
BY MR. MILLER:
Q When did Anthropic begin using this character training technique?
MS. MEHTA: Form.
THE WITNESS: I don't know specifically, but sometime between Constitutional AI training being developed, which would have been May 2023, and now.
BY MR. MILLER:
Q Do you remember which year?
MS. MEHTA: Form.

Pages 165-180

165

project, we had full robots.txt protections in place.
That's the first correction I wanted to make.
Any questions before I go to the second part?
MR. MILLER: What's the second one?
THE WITNESS: The second one is, you asked me a question that was something like, do we want our users or the public to know that if they want to remove data from Anthropic's training process, that they can contact us, and we will make an effort to work with them to remove it? And I asked you to provide me with some of the documents that we produced for the case from our support website, but those documents were not furnished. I now have those documents in front of me, and I would like to read, for the record, what the documents say.
MR. MILLER: We can mark the documents as exhibits, but we're not going to read them for the record.
MS. MEHTA: That's fine. If you want -- I think he wanted to read the titles, but if you want to mark them, that's fine.

166

THE WITNESS: Great.
So the titles are Reporting, Blocking, and Removing Content from Claude.
And this document describes, if a user believes that there is data that they would like to remove from our training data, there's a form that they can click on, and then when they click through, it leads to this page saying, I think a user is infringing my copyright or other intellectual property. How do I report it?
And it goes through the steps the user would need to take to inform us of the problem.
MR. MILLER: Okay. And we'll mark those as -- what are the next numbers?
MS. GOMEZ-REICHMAN: 192.
MR. MILLER: We'll mark those as 192A and B.
Do you have copies of those, Counsel?
(Plaintiffs' Exhibit 192A, Claude Support: Reporting, Blocking, and Removing Content from Claude, was marked and presented for purposes of identification.)
(Plaintiffs' Exhibit 192B, Claude Support: I think a user is infringing my copyright or other intellectual property. How do I report it?, was

167

marked and presented for purposes of identification.)
MS. MEHTA: Yeah.
Can the court reporter confirm that it's 192? For some reason I thought we already marked 192, but I could be mistaken.
THE COURT REPORTER: I think the last one was 191, but I'll check. I can search for it.
MR. MILLER: And for the record --
THE COURT REPORTER: Oh --
MR. MILLER: Sorry.
THE COURT REPORTER: My hands were not on my keyboard.
Okay. Yeah, so it came up. 192 is next.
MS. MEHTA: Okay. Great.
THE COURT REPORTER: And then you said A and B?
MR. MILLER: Yeah.
MS. MEHTA: She's still marking.
THE COURT REPORTER: Thank you.
Is there -- you want them in a particular order?
MR. MILLER: That's what I was just about to say.
THE COURT REPORTER: Okay.

168

MR. MILLER: For the record, we'll mark the document titled, Reporting, Blocking, Removing Content from Claude as Exhibit 192A. And the document titled, I think a user is infringing my copyright or other intellectual property. How do I report it? as Exhibit 192B.
BY MR. MILLER:
Q Mr. Mann, do you understand you're still under oath?
A Yes, I do.
Q What impact does Anthropic think that AI models will have on the economy in the next five years?
MS. MEHTA: Form.
THE WITNESS: That's a very large question, and something that we do a lot of work. We have a whole team working on this problem, called our Societal Impacts team. I spoke with a member of that team to help educate me for the purpose of this -- this testimony.
Our CEO has gone on record saying that there will be significant impact to jobs and -- and also, as I mentioned before, has this essay called Machines of Loving Grace that tries to paint a picture of what a positive outcome for a

169

transformative AI might be.
All that said, this is an area of wide-ranging opinions. Anthropic's official opinion, I would say, isn't fully formed, in the sense that while we have done research on aspects of the problem in trying to get a sense of how AI's being used today and what the trends are, we are fairly in the dark about how things will actually turn out.
We've published reports called The Anthropic Economic Index, which come with datasets that we make freely available to academic researchers who we've done partnerships with, to try to estimate AI's impact on the economy. But I would say it's still early days.
I think, in general, our estimation is that in the next five years, it's not only possible but quite probable that we will have superintelligent AI that is capable of doing most of the kinds of work that people do, and that because of that, society and economy will be dramatically transformed in the process.
And it is our hope, through our safety research, to make sure that that transition is one that's safe and beneficial for humanity, rather

170

than leading to existential-risk scenarios, which we describe at length in our responsible scaling policy documents.
BY MR. MILLER:
Q You mentioned Dario Amodei's assessment. He's previously said he believes AI could wipe out half of all entry-level white-collar jobs.
Are you familiar with those comments?
A Could you inform me where those comments were made and when?
Q Sure.
If you go to the Zoom, you should see what's being marked as Plaintiffs' Exhibit 193.
(Plaintiffs' Exhibit 193, Axios, Column/Behind the Curtain: A white-collar bloodbath, May 28, 2025, was marked and presented for purposes of identification.)
THE WITNESS: Does this have a Bates number?
BY MR. MILLER:
Q It does not.
A Okay. But it's an Axios article titled, Behind the Curtain: A white-collar bloodbath?
Q Yes.
A Okay. And the quote you are describing

171

was --
Q Hold on. Let me -- we'll get there.
Does Plaintiffs' Exhibit 193 appear to be an Axios article titled: Behind the Curtain: A white-collar bloodbath?
MS. MEHTA: Objection to the form.
THE WITNESS: It does appear to be such an article.
BY MR. MILLER:
Q And it has a date on it at the top, correct?
A Yes.
Q And it's dated May 28, 2025?
A That is the date.
Q And it shows the authors of this article as Jim VanderHei [sic] and Mike Allen?
A I don't think there's an "R" in VandeHei.
Q Oh, VandeHei.
A That's what it looks like, yeah.
Q The article lists the authors of the article as Jim VandeHei and Mike Allen, correct?
A Yes.
Q Can you scroll down to the text of the article, please?
A Yes.

172

Q Can you see the article begins, Dario Amodei -- CEO of Anthropic, one of the world's most powerful creators of artificial intelligence -- has a blunt and [sic] scary warning for the U.S. government and all of us?
That's the first sentence?
MS. MEHTA: Objection to the form.
THE WITNESS: That appears to be the first sentence, yes.
BY MR. MILLER:
Q And then there are two bullet points underneath that, right?
A Yes.
Q And the first bullet point is, AI could wipe out half of all entry-level white-collar jobs -- and spike unemployment to 10 -- 10-20% in the next one to five years, Amodei told us in an interview from his San Francisco office.
Do you see that?
MS. MEHTA: Form.
THE WITNESS: I see that text, yes.
BY MR. MILLER:
Q Does Anthropic agree with Dario Amodei that AI could wipe out half of all the entry-level white-collars jobs and spike unemployment to 10 to

173

20 percent in the next one to five years?
MS. MEHTA: Objection to the form. Asked and answered.
THE WITNESS: I guess, first off, I would say this is a paraphrase of what Dario said; it's not a quote. So I'm not sure what he actually said. But I -- I could believe that Dario said something like this.
As far as Anthropic's official position on this, I don't think we've made a statement confirming this exact projection, and we don't have research to support that exact projection.
BY MR. MILLER:
Q Does Anthropic believe that AI models will replace human jobs?
MS. MEHTA: Objection to the form.
THE WITNESS: I think we -- what we're seeing today and what we have evidence for is that there are some jobs that AI is doing more augmentation for and some jobs that AI is more automating. Whether that's causing the replacement of those whole jobs, I don't think we have strong evidence for one way or another.
BY MR. MILLER:
Q Does Anthropic believe that AI models will

174

replace human creative labor?
MS. MEHTA: Objection to the form. Vague.
THE WITNESS: Well, creative labor is one of the types of jobs that we looked at in our Anthropic Economic Index, which we've now published three different editions of, and we see some -- some of the trends on augmentation versus automation starting to play out there.
But I would caveat that that work is extremely early, and it would be hard for us to draw confident conclusions on the basis of those reports.
BY MR. MILLER:
Q Does Anthropic believe that AI models will replace songwriters in the future?
MS. MEHTA: Objection to the form. Vague.
THE WITNESS: I guess we do look at jobs involving songwriting in our Economic Index, and we have an interactive website that helps any member of the public to explore some of our data, and, unfortunately, a very, very tiny fraction of our users appear to have music-related jobs as represented by our dataset. So I don't think I could confidently make any claims with respect to the songwriting industry from the data that we

175

have.
BY MR. MILLER:
Q Apart from for purposes of litigation, which I don't want to ask about, has Anthropic conducted any analysis regarding Claude's impact on the market for copyrighted works?
MS. MEHTA: Objection to form. Outside the scope.
THE WITNESS: So I think the closest thing would be our Economic Index reports where we attempt to estimate the level of usage and, again, the degree of automation versus augmentation for each of those professions. I think that is quite upstream of market impacts -- labor market impacts.
We've worked with some economists who have done labor market impact analyses, but some of these studies have fairly significant methodological flaws. We're working with them to get them better data so they can do a better job.
But, again, I would say there is no one in the world who can confidently say, today, what the impact will be. People have intuitions, but I would say that they are not rigorous or data-driven.

176

BY MR. MILLER:
Q Claude is capable of generating new song lyrics, right?
MS. MEHTA: Form.
THE COURT REPORTER: I'm sorry. Did you object?
MS. MEHTA: I just said, Form.
THE COURT REPORTER: Okay.
THE WITNESS: If you ask Claude to generate song lyrics, it's likely Claude will output something that probably looks like song lyrics.
There have been instances that I've seen inside Anthropic that look pretty good, but I am not a musician, and I don't know what kind of quality level an -- a professional would sub -- would ascribe to those lyrics.
BY MR. MILLER:
Q Does Claude's ability to generate lyrics for a song impact the market for existing songs?
MS. MEHTA: Form. Also, calls for expert testimony. Outside the scope.
THE WITNESS: To the extent that Anthropic has attempted to look into something like this in our Economic Index report and the data that we've

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

Pages 185-188

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?

Pages 197-204

197

A Yes.
MR. MILLER: Tab 40.
(Plaintiffs' Exhibit 196, Excel spreadsheet, hyperlinked in Exhibit 195, Bates stamped Anthropic_0000240674, was marked and presented for purposes of identification.)
BY MR. MILLER:
Q In the Zoom, you should see what's being marked as Plaintiffs' Exhibit 196, which bears the Bates number Anthropic_240674.
A I see that.
Q And do you recognize this document?
MS. MEHTA: Objection to the form. Outside the scope.
THE WITNESS: I don't specifically recognize this document, but I'm familiar with documents in this format.
BY MR. MILLER:
Q And what type of document does this appear to be to you?
MS. MEHTA: Objection to the form. Foundation. And scope.
THE WITNESS: This appears to be a document representing some crowd worker judgments on whether the model was supposed to refuse, and

198

then whether it did refuse, to give information on various topics.
BY MR. MILLER:
Q Is this the document referenced in the previous exhibit we looked at that Craig provided a link to in the Slack communication?
MS. MEHTA: Objection to the form. Scope. Foundation. And speculation.
THE WITNESS: Given that Craig's message said that there were 988 examples, and this spreadsheet does not contain 988 examples, I would be surprised if this was the same document.
BY MR. MILLER:
Q How many examples are in this spreadsheet?
MS. MEHTA: Same objections.
THE WITNESS: 894 rows, but the first row is a header, so I guess 893.
BY MR. MILLER:
Q Okay. Are you familiar with something called Hugging Face?
A Yes.
Q Is Hugging Face a platform where members of the machine-learning community can collaborate on models, datasets, and applications?
MS. MEHTA: Form.

199

THE WITNESS: Hugging face is multiple things; it's an emoji; it is a company; and it is also the platform that the company provided matching the description you mentioned. Possibly, there are other definitions.
BY MR. MILLER:
Q Has Anthropic made certain datasets available on Hugging Face?
MS. MEHTA: Form.
THE WITNESS: I believe we have.
MR. MILLER: Tab 41.
(Plaintiffs' Exhibit 197, Anthropic/hh-rlhf - Datasets at Hugging Face, was marked and presented for purposes of identification.)
MR. MILLER: Is this 197?
BY MR. MILLER:
Q In the Zoom, you should see what's been marked as Plaintiffs' Exhibit 197.
Does Plaintiffs' 197 appear to be a screenshot of the Hugging Face website?
MS. MEHTA: Objection to the form. Foundation.
THE WITNESS: I'm seeing that there are six pages, but the scrolling is not working.

200

Oh, there it goes.
MS. MEHTA: Do you want him to allow the security warning or not?
THE WITNESS: It's frozen.
MS. GOMEZ-REICHMAN: You just opened a URL, so let me take control back for a moment. No problem.
THE VIDEOGRAPHER: Alt-shift-G.
MS. GOMEZ-REICHMAN: Let me re-share this, and hopefully, we won't have an URL open this time.
Okay.
THE WITNESS: Where should I see the exhibit number?
BY MR. MILLER:
Q It's not visible on the face of the document.
A Should there be a Bates number?
Q There is not. But do you see, at the bottom-left corner of the first page, there's a URL?
A Yes.
Q And the URL indicates that this is a PDF reflecting a website on the Hugging Face domain?
MS. MEHTA: Counsel, are you representing

201

that you printed this, and this is a complete and accurate copy of this URL?
MR. MILLER: Yes, I can represent that.
MS. MEHTA: Okay.
THE WITNESS: So what's your question?
BY MR. MILLER:
Q Does this appear to you to be a PDF created to reflect Anthropic's helpfulness and harmlessness datasets available -- let me try that again.
Does this appear to be a PDF created reflecting Anthropic's web page on Hugging Face, where it makes available datasets related to Anthropic's helpfulness and harmlessness analysis from reinforcement learning from human feedback?
MS. MEHTA: Objection to the form.
THE WITNESS: It does appear to be a printout of Anthropic's hh-rlhf dataset.
BY MR. MILLER:
Q Okay. Could you turn to page 3, please?
A Yes.
Q You see, at the top of page 3, this web page states, This repository provides access to two different kinds of data.
A Yes.

202

Q And then underneath that, there's a number 1.
A Yes.
Q And the web page states, Human preference data about helpfulness and harmlessness from training a helpful and harmless assistant with reinforcement learning from human feedback.
Do you see that?
A Yes, I do.
Q Training a helpful and harmless assistant regarding reinforcement of human feedback was the title of the research paper we just looked at a minute ago, right?
A Correct.
Q So is the data that Anthropic is making available here the data underlying the research reflected in that paper?
MS. MEHTA: Objection to the form. Foundation.
THE WITNESS: There's a lot of research reflected in that paper, and I believe the paper specifically refers to a dataset that would be made available as part of the paper release. So the data represented on this website would be whatever the paper is referring to.

203

BY MR. MILLER:
Q Okay.
MR. MILLER: Tab 48.
(Plaintiffs' Exhibit 198, Excel spreadsheet, 2025.11.20 Huggingface_Anthropic_Train, 160800 rows listed, was marked and presented for purposes of identification.)
BY MR. MILLER:
Q In the Zoom, you should see what's being marked as Exhibit 19 -- excuse me -- Exhibit 198.
MS. MEHTA: Could you add that to the chat? I don't have --
MS. GOMEZ-REICHMAN: Yeah. It's taking a second to come through. It's quite a large file.
MR. MILLER: It's a very large data file.
MS. MEHTA: Give it a second. I need to access it. It's still downloading for me.
Okay. Great. I'm all set.
BY MR. MILLER:
Q Mr. Mann, does what's been marked as Plaintiffs' Exhibit 198 appear to be the dataset we were just referring to?
MS. MEHTA: Objection to the form. Outside the scope. And also, calls for

204

speculation. Lacks foundation.
THE WITNESS: This looks like a dataset. I don't remember, in particular, if this is the one.
BY MR. MILLER:
[redacted]
MS. MEHTA: Object -- objection to the form.
[redacted]
MR. MILLER: Okay. Can we take a break?
MS. MEHTA: How long have we been going?
MR. MILLER: A little over an hour.

Pages 213-216

213

Mr. Mann is designated on that with respect to training. Your questions are not appropriately scoped.
MR. MILLER: My notes do not reflect that division of responsibility between Dr. Rao and Mr. Mann.
MS. MEHTA: Well, I think you should go back and look at the correspondence. It's quite clear.
MR. MILLER: All right. Can we go off the record, then?
MS. MEHTA: No, we're not going to do that now. You can keep asking your questions, but I'm -- I'm -- you have limited time left, even if I give you the full credit for the 5 hours and 30 minutes that you negotiated. I would recommend that we move on. We have a lot to cover today.
MR. MILLER: Let's go off the record.
MS. MEHTA: I don't -- I don't agree to that.
What are you going off the record to do?
MR. MILLER: To confirm the scope of the witness' designations.
MS. MEHTA: Why don't -- why don't you do that at the next break?

214

MR. MILLER: We're going to do it right now.
MS. MEHTA: Okay. We have -- we're going to take a two-minute break, because this witness is not going be here at 10:00 at night because you don't know the scope of the testimony.
THE VIDEOGRAPHER: We're going off the record at 15:39.
(A brief recess was taken.)
THE VIDEOGRAPHER: We are back on the record at 15:44.
MS. MEHTA: During the break, we just said that we would permit the questioning to go forward, subject to a scope objection, and we'll all sort it all out later.
BY MR. MILLER:
Q All right. I'll repeat the last question I asked.
Can you tell me whether the Acceptable Use Policy prohibiting using Claude to engage in illegal activity thereby prohibits using Claude to generate outputs containing copyrighted works without disclosing attorney-client privileged information?
MS. MEHTA: Objection to the form.

215

Outside the scope. Asked and answered.
And, again, Mr. Mann, you can respond as long as you do not include any attorney-client privileged information or advice of counsel.
So if you have information independent of that, you can share it.
THE WITNESS: Independent of my knowledge of counsel's opinion on the subject, as a non-lawyer but a person interested in copyright, I think it is possible that the Acceptable Use Policy version that we are looking at here could cover the things like outputting copyrighted works; however, for Anthropic's official opinion on that, I wouldn't be able to answer the question due to attorney-client privilege.
BY MR. MILLER:
Q Okay. Are you familiar with something called data cleaning?
A Yes, I'm familiar with the concept of data cleaning.
Q What is data cleaning?
A So, often in the process of machine learning on a dataset, a dataset has lots of problems that, if unaddressed, would lead to significant underperformance of the resulting

216

machine-learning model. This is a term of art that has been used for decades, not just with large language models. It's an accepted practice in the industry, and I've certainly done extensive data cleaning in my time as a data engineer.
Q We'll come back to data cleaning in a second.
Are you familiar with the dataset Common Crawl?
A I am familiar with a -- an organization called Common Crawl, and there are many different datasets derived from that organization that have "Common Crawl" in the name, yes.
Q Has Anthropic made use of datasets from Common Crawl in its training of models?
MS. MEHTA: Form.
THE WITNESS: Yes, we have used datasets derived from Common Crawl in our model training.
BY MR. MILLER:
Q How many different datasets from Common Crawl has Anthropic obtained?
MS. MEHTA: Objection to the form.
THE WITNESS: Well, maybe I'll do a little bit of defining. What do you mean by a "dataset" here?