Exhibit C



HIGHLY CONFIDENTIAL--ATTORNEYS' EYES ONLY

Transcript of Andrew Bray

Date: February 24, 2026
Case: Concord Music Group -v- Anthropic

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

1

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

CONCORD MUSIC GROUP, INC., ET AL., Plaintiffs, v. ANTHROPIC PBC, Defendant. Case Number 5:24-cv-03811-EKL SVK

HIGHLY CONFIDENTIAL--ATTORNEYS' EYES ONLY
VIDEOCONFERENCED and VIDEO-RECORDED
DEPOSITION OF
ANDREW P. BRAY, Ph.D.
San Francisco, California 94111
Tuesday, February 24, 2026

Reported Stenographically by:
MARY J. GOFF
CSR No. 13427
WA CSR No. 21030779
Job No. 620584
PAGES 1-225

2

HIGHLY CONFIDENTIAL--ATTORNEYS' EYES ONLY
VIDEOCONFERENCED and VIDEO-RECORDED DEPOSITION of ANDREW P. BRAY, Ph.D., Volume I, taken on behalf of Plaintiff Publishers, Oppenheim + Zebrak, LLP, at WilmerHale, 50 California Street, Suite 3600, San Francisco, California 94111, beginning at 9:03 a.m. and ending at 4:10 p.m., on Tuesday, February 24, 2026, before MARY J. GOFF, California Certified Shorthand Reporter No. 13427 and WA CSR No. 21030779.

3

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

For Defendant Anthropic PBC
WilmerHale
BY: DISHA PATEL
LOUIS W. TOMPROS
Attorney at Law
60 State Street
Boston, Massachusetts 02109
disha.patel@wilmerhale.com
617-526-6795

ALSO PRESENT: Natalie Naugle, in-house counsel Anthropic
Videographer: Justin Ramirez-Hale

4

INDEX

WITNESS		EXAMINATION
ANDREW P. BRAY, Ph.D.		
Volume I		
	ATTORNEY HOWELL	6
	ATTORNEY PATEL	217

NUMBER	DESCRIPTION	PAGE
Exhibit 1	Anthropic's Supplemental Responses to the Publisher Plaintiffs' First Set of Interrogatories, No. 4	36
Exhibit 2	Lyrics ANTH_SAMPLE000048259	139
Exhibit 3	Lyrics Viva La Vida ANTH_SAMPLE000217578	175
Exhibit 4	Exhibit B to Publishers' Amended Complaint	179

Pages 9-16

9

subsequent draft, the revisions.
And I don't recall for any particular section what that back-and-forth looked like, but there -- the contributions of WilmerHale in terms of drafting the reports were at my direction.
Q Do you have an understanding of the legal issues involved in this case?
A I have some understanding of the legal issues involved. And that understanding was built from conversations with counsel and the review of the documents that they provided to me regarding the kind of parameters of the case.
Q And what is that understanding that you have?
A I understand that the Plaintiffs are -- I believe that there's been copyright infringement on lyrics at a few different phases of Anthropic's product Claude and that Anthropic disputes that characterization.
Q And do you have an understanding of the legal standard for copyright infringement?
A I don't. I don't have a particular -- any particular expert in -- expertise in law. However, I understand that the opinions that I have rendered in these reports do have some bearing on that

10

question.
Q And do you know if Anthropic is asserting any defenses in this case?
A Could you describe to me what you mean by "asserting defenses"?
Q Are you aware of, for instance, Anthropic asserting a fair use defense?
A I -- I have heard the term "fair use" being -- it appears in -- it's appeared in conversations with counsel. It's appeared in Dennis's report. But I don't have a particular precise understanding what that means in terms of legal understand, but I do understand that the opinions that are rendered may have some bearing on that standard.
Q And have you been involved in any copyright cases before?
A I -- I have not. The -- my experience as an expert witness was confined to a case having to do with ballistics.
Q So you have been retained as an expert in one prior case? Is that what you're saying?
A That's correct.
Q Are you currently retained in any other cases?

11

A I am not.
Q And you have never been involved in a case involving song lyrics before?
A I have not.
Q Okay. And what experience -- or what expertise are you applying in your reports here?
A I would say that my center of expertise is around the field of statistics and data analysis, also computing with data.
And this is a fairly broad field. It involves predictive modeling. It involves inferential statistics. But ultimately, it centers on how to draw conclusions from data.
Q And anything beyond the fields of statistics and data analysis?
A I would say that the -- could you perhaps give me an example of what you mean by "an extension"?
Q Any other areas that you have expertise in that you have applied here in this case?
A I would say, yeah, the principal area of expertise that I'm working in is statistics and data analysis kind of generally conceived.
Q And nothing else? No other areas come to mind?

12

A No other areas beyond the interpretation analysis of data.
Q You're not an expert in musicology, for example?
A I'm not an expert in musicology, although I understand that the opinions that I have rendered have some interplay with those that do have expertise in musicology.
Q And you don't have any particular expertise in song lyrics?
A I have no particular expertise in song lyrics, although I -- the opinions that I rendered do some have some bearing on -- on that as a -- kind of a context from which data that I analyzed came from.
Q Are you -- do you consider yourself an expert in linguistics?
A I do not consider myself an expert in linguists, although the opinions that I have rendered in the reports have some interplay with linguists, because that's the context from which the data arose.
Q And you're currently a university professor, correct?
A That's correct.

13

Q Have you ever worked for a generative AI company before?
A I have not worked for a generative AI company.
Q And would that include both direct employment or -- and any consulting work?
A Certainly not direct employment. There's a company who I have done some consulting work for who are, since I have done consulting work for them, starting to get into areas related to generative AI.
Q But the areas you worked on with that company didn't address generative AI; is that fair?
A Not directly. This is a -- it's a software company dealing with data and models, but not LLMs in particular.
Q Given it was only one case, I think the answer probably "No," but have you ever been retained by counsel for Anthropic in any prior cases?
A I have not.
Q Okay. Have you ever published any research on generative AI?
A So the CV, as it was provided, in that CV there's no publications about generative AI. There on other areas of statistics and -- and

14

classification models, models for data. However, there's a paper that is -- a research study that I'm currently engaged in that does bear on that, but the manuscripts and draft phase has not yet been published.
Q Okay. So that manuscript would not be publicly available; is that fair?
A That's fair.
Q Okay.
A We have -- I should say we have submitted -- there are possibly public grant proposals that describe the -- the research.
Q Okay. If I wanted to look for one of those grant proposals, what search terms would I use?
A "Grading accuracy" would be the main one. So the research, it's looking at -- one of applications that's come up for large language models is in grading student work. This is something that's generally done by humans where you have an expert, a teacher devising a rubric for a particular question, and then you have a human -- at Berkeley we use students to apply rubric to student work.
And this has been kind of the only way

15

that we can get grading done at scale. And it has its benefits. It has its downsides. We have perhaps an alternative through LLMs that perhaps LLMs can be useful in this task of grading student work. And so we're simply looking at the question of accuracy of -- the accuracy of student graders versus the accuracy of LLMs as -- as graders. How do they compare, and in what cases do LLMs outperform and in which cases do humans outperform.
Q And do you have any published research where you use regular expression?
A I would have to -- can I take moment --
Q Sure.
A -- to look through? So I think my CV was the first. I can't state for certain, but it's possible that the 2019 paper, "Nature Climate Change," involves regular expressions.
Q Which page of your CV are you looking at?
A This is the first page.
Q This is under "Refereed Articles"?
A Correct.
Q And I'm sorry. What was the --
A The title is "Thermal Extremes Drive Fish Die-offs in Temperate Lakes."
I'm uncertain. The flag for a regular

16

expression would be kind of the -- if -- it was in an asserted process text data. And the data that we were looking at there was coming out of fishing game out of the Upper Midwest, and it may have been necessary to apply some text processing in order to analyze that data.
Actually, there is -- there's a piece of software here that I believe uses regular expressions. And that is under -- on the next page, page 2, "R packages, Jayaprakash, Voutov, and Bray, 2024, nemogb."
Q Okay.
A So again, that's a -- this is piece of open-source software that processes text data, and so regular expressions are -- were used to do that in a structured manner.
Q Okay. And for the article, I think you said that article could have used regular expressions, but you weren't sure; is that --
A I wasn't sure, yeah. And it would be in the data processing step, so it almost certainly wouldn't appear in the text of the article itself. It would have been a component that enabled the analysis.
Q All right. Any other articles involving

Pages 25-28

25

A Was -- could you remind me? Is Dr. --
Q He's an expert for -- retained by Anthropic.
A I may have, but I don't recall having done so now. And on a cursory review of my reports, I can't find any reference to that.
Q And no opinions regarding Dr. Altman's analysis come to mind?
A None come to mind.
Q And Dr. Lustig performed an analysis of a sample of the records that Dr. -- sorry -- that Mr. Dennis reviewed.
Did you assess the sufficiency of that sampling method performed by Dr. Lustig?
A This -- it would be -- it would be helpful if I had Dr. Lustig's report so I could review the description of the sample. So I guess without that description of the sample, I -- I'm unsure -- or -- or yeah, I wouldn't want to voice a strong opinion.
Q Do you recall reviewing that section of --
A I do --
Q -- Dr. Lustig's report?
A -- yeah.
Q And you didn't include any opinions regarding that analysis in any of your reports,

26

correct?
A I believe that's the case.
Q And prior to this case in your professional career, have you ever been asked to assess whether a document contains a request for lyrics?
A Not to my knowledge, no.
Q And prior to this case in your professional career, have you ever been asked to assess whether something contains regurgitation of lyrics?
A Not to my knowledge.
Q And your affirmative analysis -- so I'm referring to your opening report and your reply -- you analyze prompts and outputs that came from the Claude sample; is that right?
ATTORNEY PATEL: Objection, vague.
A Can you please restate the question?
Q (BY ATTORNEY HOWELL) You understand that Anthropic produced a sample of five million records containing Claude prompts and outputs in this case, correct?
A My understanding -- yes, my understanding is that Anthropic produced, pursuant to a court order, a set of five million records, and those

27

files were the subject of my analyses.
Q And I think you called those five million records "the sample" in your report; is that fair?
A I believe that's true.
Q If I refer today to "the sample," will you understand that I'm referring to those five million records?
A Yes.
Q All right. And in your affirmative analysis, you rely on prompts and outputs that come exclusively from the sample, correct?
A I believe so, yes.
Q And we already mentioned there's five million records in the sample, correct?
A That's correct.
Q How many of those records did you personally read?
A Here I would have to estimate and ask: Do you mean read in full? Read the part?
Q Let's start with "read in part."
A I would put it on the order of magnitude of hundreds, although that's an estimation.
Q Okay. And "read in full," how many would you estimate in read in full?
A I would say less than a half of that, but

28

probably still hundreds.
Q And when you say "hundreds" for "read in part," what -- is that something under a thousand; is that fair?
A That would be my guess --
Q All right.
A -- under a thousand.
Q And let's see. And you understand that sample came from a wider universe of records?
A That's my understanding.
Q And that covers -- and that universe covers the same time period?
ATTORNEY PATEL: Objection, vague.
Q (BY ATTORNEY HOWELL) It is vague. I'll set it up.
The sample records all come from September 22, 2023 to March 22, 2024, correct?
A Correct.
Q Okay. And those sample records come from a wider universe of records covering that same time period, correct?
A Just one moment, please.
Q Paragraph 16 might be helpful. That refers to the metadata in the universe.
A Great. So this states the -- paragraph 16

Pages 45-68

45

formatting in spaces, line breaks, white spaces. It can, of course, search for characters or character strings verbatim.

It can allow for certain logical conditions. So: I'm looking for this, as long as it's not followed by that thing; I'm looking for this thing or that thing or this thing; I'm looking for this thing at the beginning of a sentence; I'm looking for this thing partway through a sentence; I'm looking for this thing at the end of a sentence.

So it's a -- it's a tool that's used to very flexibly search through text data.

Q And would it be fair to describe regular expressions as a way to search for a particular set of patterns of words, spaces, line breaks?

A That's a fair characterization. Often the -- the formulation of the regular expression itself is called the pattern that you seek to match to.

And that pattern can contain verbatim words or it can contain kind of variables, things -- you can build up, for example, the notion of an email address because you know its structure. You know that an email address has some number of acceptable characters pre -- that precede an @ sign and then a series of one or more of a set of other

46

characters and then a "dot "and then three letters that are pulled from a particular set.

Q Do regular expressions look for a particular pattern in text?

A They do.

Q Okay. And you used regular expressions here in a deterministic way?

A That's correct. Yeah, regular expression if you have this, you know, one set of documents. And you run regular expressions over them and then come back the next day and run the same set of regular expressions, you'll get the same response.

Q Okay. And if a document matched the pattern you were looking for for a particular regular expression, you included that document in your results, correct?

A Yes. To be more precise, the -- there are two steps to the analysis. Step 1 involved the determination of which files were valid files that illustrate user behavior. That determined which documents I would be searching over. Regular expressions were used in that process along with a few other techniques.

The secondary stage was of those documents that constitute valid records, I think what I called

47

the adjusted sample, each of those documents was subject to a regular expression. And all the documents that matched that regular expression that were responsive to it are in the numerator of each of the proportions that are provided in the answers to the research questions in the summary.

Q You didn't perform a manual review of documents that matched the pattern to see if they, for example, actually contained a request for lyrics, correct?

ATTORNEY PATEL: Objection, vague.

A Could you describe -- could you restate it AND describe what you mean by "manually review"?

Q (BY ATTORNEY HOWELL) You didn't -- you didn't read through the documents that matched the pattern for, for example, your regular expressions for Research Question 1 to assess whether, in fact, they should be included?

A I did. I read through a reduced spot-check. So after it would pass, you know, every -- you can -- I think about a regular expression kind of as a sieve and throwing all the valid documents through the sieve -- one particular sieve. And then what comes out the other side is a collection of documents. And I would, you know,

48

after every run through a sieve, take a look at the documents and determine, I guess, at -- one level is the -- is the structure of the regular expression behaving as I intended and then is it capturing the -- kind of semantically is it capturing what we wanted this regular expression to capture.

Q Once you finalized your regular expression for Research Question 1, you ran that regular expression and looked for documents that matched that pattern, correct?

A That's correct.

Q And if a document that you assessed matched that pattern, you included it in the list of files that were responsive to Research Question 1, correct?

A That's correct.

Q You didn't take that list of files and do a manual review to see if they actually contained a request for lyrics, correct?

A Similar to the answer to the first version of this question, I did spot-checks after every pass of the sieve, and so, for example, I was aware that there remain false positives in the collection of documents that were responsive to Question 1. And here defining a false positive as a document that

49

the regular expression matched to, but that clearly do not -- were not in this case seeking any amount or form of lyric regurgitation.

Q And so you -- you recognize that there are false positives in the list of files that match your regular expression?

A That's correct.

Q You didn't remove those files from the list of documents that you found to be responsive to each research question correct?

A That's correct. They remain in the -- in the collection of documents that were responsive to that regular expression.

Q Why wouldn't you remove false positives?

A The approach that was used here is a standard approach in building a classifier model.

Whenever you're building a classification model -- a two-class classification model, there's the possibility that that classifier will generate false positives. There's the possibility it'll generate false negatives.

It's, in many cases and certainly in this case, impossible to write a classifier -- realistically write a classifier that could do perfect sorting that wouldn't have either false

50

positives or wouldn't have false negatives, so the question then becomes: How flexible should my model be -- work -- can I build a model? And here by "model," I mean the regular expression that will capture the -- the sort -- sort of text documents that I'm seeking without incurring too many false positives.

And so it's trying to balance the -- as proportions, these are the recall and the precision of the classifier, in this case thinking about just having a balance of false positives and false negatives.

Q Do you have a sense for how many false positives there are in your set of documents for Research Question 1?

A I haven't reviewed enough of them to have kind of a systematic estimate. I would -- I think it's fair to say just many of the documents are false positives.

Q Can you provide any more precision regarding what percentage of the records you found to be responsive to Research Question 1 you considered to be false positives?

A I wouldn't provide any more precision. It was based on a spot-check and the determination that

51

indeed there were still false positives in the -- you know, upon seeing false positives, there's the option of making the regular expression less flexible.

So the original regular expression was not very flexible. It had only matched to a few very explicit types of files. And the goal here was to have a classifier that had high precision, which is to say of those that we were classifying as positive, we wanted to be sure that a high proportion of them were positive.

The second stage -- and this is what describes the refinement of the regular expressions -- this was to address the recall side of the equation. So initially the model was built to be kind of deliberately strict, knowing that it would be excluding certain ways that people would, you know, write down a -- when they're seeking any amount or form of lyric regurgitation, knowing that we are -- you know, there are more subtle forms of asking for -- more indirect forms of asking for it.

And knowing that we would miss those in the original regular expression, I had Claude -- we can get into this more explicitly, if you would like -- surface examples of these false negatives

52

that allowed me to find ways that I could broaden that regular expression.

So it's a -- you have these competing forces. You have -- if you have a lot of false positives, there's a -- kind of a need to constrain your regular expression, make it a little bit more strict. If you have it too strict, you are going to be having more false negatives, and so there was a back-and-forth trying to determine the -- the -- kind of the Goldilocks zone here of regular expression that was just the right level of flexibility that would capture most of the text documents that were seeking any amount or form of lyric regurgitation without inducing a flood of false positives.

Q Do you have a sense for the false positives rate for Research Question Number 2?

A So Research Question 2 is looking at those that were responsive to Research Question 1 and then further refining kind of what it was looking for. It's looking for the, I believe it's, song title or artist.

And if -- if it found those within particular parameters of the regular expression, it was responsive. In terms of the -- the -- similarly

53

to Question 1, I spot-checked the results and can't recall -- I don't recall whether there were false positives in -- in this.

Q So sitting here, you're not aware of what the false positive rate was for Research Question 2?

A That's correct.

Q Research Question 3 "assessed the percentage of outputs that provided any amount of -- any amount or form of lyric regurgitation."

Do you recall that?

A Sorry. Could you restate the question?

Q Sure.

Research -- your Research Question 3 "assessed what percentage of outputs provided any amount or form of lyric regurgitation"?

A That's correct.

Q And you found around 300 files that matched that -- the pattern you were looking for for that research question?

A That's correct.

Q Do you have a sense for the false positive rate for those relating to those files?

A I -- so for those 325 files, I also spot-checked it, and I do recall finding false positives. I don't recall what the -- I wouldn't

54

try to be precise about, you know, precisely what that false positive rate was.

Q Okay. Research Question 4 "assessed what percentage of outputs provide any amount or form of lyrics to the works-in-suit"?

A Um-hum.

Q And you found, between your opening report and your reply report, around a hundred files?

A Ah. Between -- oh, right. So I should clarify that I've been looking at the -- my original summary opinions that were -- are you referring to these numbers or are you referring to those from the reply report, just so we're consistent?

Q Sure. So I guess let me understand first.

You -- as part of your reply, you reran your analysis to address a non-English language exclusion --

A Right.

Q -- issue?

A That's correct.

Q Are you intending to offer at trial, the opinions that you disclosed in your opening report regarding your research questions or does the reply report over -- override those?

A The precise numbers in the reply report

55

would be the ones that I would offer up at trial, although substantively they're -- they're largely the same, at least the proportions are, as those that are found in the opening report.

Q So for Question 4, in your reply it should be at paragraph 63, if you want to look at it there. You ultimately found, for Question 4, 109 responsive records?

A Question 4? Correct. 109 of the 313 outputs were providing any amount or form of lyric regurgitation, as kind of articulated in the regular expression that was used there.

Q And do you know what the false positive rate is for those 109 files found responsive to Question Number 4?

A Similar to the previous ones, I didn't manually review all of them. I spot-checked them. I don't recall -- I don't recall -- I don't recall whether -- whether or not there were false positives there. It's maybe worth noting that the regular expression was unchanged, you know, from the opening report to the reply report in the case of Questions 1 and 2 and 3, but the regular expression in Question 4 was -- there was a typo in the regular expression that led to some of the classification

56

errors that Mr. Dennis flagged, and I was able to -- it was one character that -- a typo in the regular expression for Question 4 that, once corrected, led to this number.

Q And do you disclose the corrected regular expression for Question 4 in your reply report?

A One moment. Actually, it's probably -- I do. It's on page 9 of -- oh, gosh, Table A. It's the first page. It's labeled "Pattern Corrected."

And it can be compared to the analogous code that was provided in the opening report where that variable was called just "Pattern."

And the distinction is the -- at the end of the regular expression.

Q And sorry. This is Exhibit A to your --

A Table A. Sorry.

Q Table A.

A One of the last pages. Page 9 of Table A. Were you able to find it?

Q I'm not. This is an exhibit to your reply or is it --

A This is a table. Table A in the reply. Table A, page 9. And it may have shown up -- may have shown up in the body. Yeah, 22. Right back -- yeah. So it's on page 9 of Table A.

57

Q Okay. And you mentioned spot-checking for false positives for your Research Questions 1 through 4?
A That is correct.
Q Did you do anything beyond spot-checking to assess the false positive rate related to your Research Questions 1 through 4?
A Sorry. Could you restate that?
Q Yeah.
Did you do anything other than the spot-check you described to assess what percentage of records you deemed responsive to Research Questions 1 through 4 were false positives?
A The -- the main way that false positives were -- utilized a spot-check of false positives was during the refinement process where I -- actually, let me show you.
In the opening report, Table -- or there it is. Table A, page 3 -- it's labeled "Q1 v2" -- it was clear from the Claude refinement process there were false negatives that were caused by the insertion of, you know, just additional words that are occurring between what I'm characterizing as kind of this verb request part and the object, the thing that they're requesting. People are using

58

different -- additional words in there. So wanting to provide for that flexibility, that's what that option two-word buffer was.
Originally that was set -- I don't recall if it was at a five-word or a four-word. I think I found an example of a false negative flagged by Claude and adapted the buffer to exactly match that and then ran that filter, ran the regular expression, and then spot-checked the results.
And you can tell from that basically that each modification, are you getting a flood of false positives? Are you kind of getting more chaff for wheat?
And so that's -- the false positives helped me determine when the regular expression was becoming too flexible.
Q Okay. And so there you're describing the spot-checking you were doing as you were iterating your regular expression?
A That's correct.
Q And for the final regular expressions, you, again, did a spot-check for false positives?
A It was, yeah, a spot-check of the records. And to -- you know, it was after every kind of -- "try a slightly different calibration, a regular

59

expression, either more flexibility or less flexibility after each one," wanting to determine: Okay, I have made it more flexible. Does it -- has this induced a flood of false positives?
And if it does -- if it did, then that amount of flexibility would lead to a poor classifier with poor, you know, error characteristics, so it would keep it more constrained, less flexible.
Q Okay. And for your final set of regular expressions, did you do anything to systematically assess -- systematically assess the false positive rate associated with those regular expressions?
A I did a spot-check of the files that were responsive to the final one. And -- and in this here, it would depend on the precise question. I recall in certain questions, viewing false positives and making the determination that I couldn't constrain the regular expression to exclude them. And if I were to constrain it farther, it would create more false negatives than the false positives that were being induced, so in that case it would stay the same.
But your question about the -- you know, offering a precise error rate, I didn't do an

60

exhaustive manual review of all of the ones that were responsive after every run of the regular expression.
Q And so you don't provide a precise error rate associated with false positives related to any of your final regular expressions, correct?
A That's correct.
Q Okay. And what did you do to assess false negatives beyond the Claude refinement?
A So I guess -- you know, the notion of a false negative, this is when, you know, there's an expression that would correspond to the sort of behavior that I'm seeking to capture. Someone seeking any amount or form of lyric regurgitation in the case of Question 1.
THE COURT STENOGRAPHER: Wait.
THE DEPONENT: Sorry.
THE COURT STENOGRAPHER: Slow down.
Amount of lyric regurgitation?
A Any amount or form of lyric regurgitation for Question 1. Sorry. I lost my train of thought. Could you ask the question again?
Q (BY ATTORNEY HOWELL) Sure. What did you do to assess false negatives beyond the Claude refinement?

61

A Right. Thank you.
So yeah, first being clear what a false negative was, there's an expression that is not currently being captured by the regular expression. That notion informed the original formulation of the regular expression. So that -- that first form that shows up in -- it's not labeled View 1, but under the Q1 section, that more paired back regular expression, you know, even that is keeping in -- it has some amount of flexibility built in to catch as many genuine expressions of this behavior that we're seeking to capture.
But beyond that, the main tool that was used to address the possibility for false negatives and wanting to just the regular expression accordingly was the -- this process of Claude refinement.
Q Sorry. You said the Claude refinement was the main tool that you used to assess false negatives. Was -- anything else in particular you used to assess false negatives associated with your regular expression?
A I think, you know, limiting the false -- addressing false negatives, wanting to capture -- you know, multiple ways to express this particular

62

behavior, you know, that -- and that -- that was the goal of part of the regular expression. The other part being restricting the false positives.
And the way that that concern was materialized was in, you know, my own formulation of that first version of the regular expression and the subsequent refinement using Claude.
Q Okay. For Research Question 1, are you able to quantify how many false negatives there are in the sample that you didn't find?
A No. So the extent to which, you know, false negatives -- the way that I dealt with false negatives was through kind of holistic, not data informed -- or not directly data informed, not algorithmically informed, articulation of the regular expressions in the process of Claude refinement.
But the -- you know, an assessment of the false negative rate would require -- in this case because the regular expression were run over, I guess, roughly 2.5 million, the adjusted sample, it would require an exhaustive -- like, actually knowing that rate would review an exhaustive review of those records. And doing, you have to consider -- one -- I guess how many -- which records

63

do you look at? Do you look at all of them? Do you look at a subset of them in trying to assess that error rate?
And then secondly you look: How do I determine what the ground truth is?
So you know, the notion of a false negative is premised on the notion that, like, I can -- I can show you a file and I can tell you beyond a shadow of a doubt that that expresses -- that should have been responsive, because that does capture this thing that I'm trying to capture.
And you know, methods that you could try to get at that ground truth are, I would say, manual review. That's one method.
Another method is kind of an analog using large language models to do that review. My experience working with Claude in that process is that, you know, it -- it was not providing me the ground truth. It was providing suggestions.
I disagreed with its suggestions sometimes. I -- it flagged certain -- certain files that were not kind of concordant with what I was seeking in those research questions.
Similarly, the manual method that was used by Mr. Dennis, that, too, is shown to be

64

error-prone. And you know, he has -- he does his own error analysis. And then Mr. Lustig does an error analysis on -- on that error analysis and found that even that manual process had its own errors.
So the difficulty of establishing what a ground truth would be, particularly at scale in order to estimate that false negative rate, would give kind of a misleading sense of -- of certainty because it's -- it's -- as the -- this case has shown and Mr. Dennis's categorization followed by Dr. Lustig's, it's -- it's very difficult to determine for any particular file whether or not it should be responsive to a particular research question.
Q So it's practically impossible to determine the false negative rate for at least Research Question 1?
ATTORNEY PATEL: Objection, mischaracterizes the witness's testimony.
A I would say the -- so quantifying the false negative rate would require a sense of ground truth. And practically getting a ground truth -- so using the method that Mr. Dennis used, which is to say manual review followed by -- well, I will move

65

to say manual review, that process would be challenging.

But the even -- even after the -- like, the practical aspects of it were dealt with, which is to say just the volume of the data, you run into the problem of: We don't have a reliable ground truth model here. Right. We have -- we have, you know, multiple opinions by multiple humans, Mr. Dennis, Mr. Lustig, discussing the questionable state of just two documents.

And in this case the -- to get that false negative rate, you would have to -- you should have that level of scrutiny over 2.5 million documents. And even then there's still some debate over: For every given file, is that actually -- what is it capturing?

So the advantage of the regular expression is it's known that it has false positives, false negatives because it -- you have to make some sort of determination on the flexibility of the regular expression.

The determination I made was based on balancing the false positives and the false negatives, kind of a goldilocks-style model. But this has the advantage of being very transparent in

66

what it's capturing. So you know, the regular expression, it exists there in a few lines of code. And you can -- you know, in this case it's computer code. But I could translate this exactly into a paragraph that articulates exactly the sorts of files that it would return.

So we know precisely what the regular expression is giving. It's very transparent. It's also rep -- replicable. So in this case, regular expression -- Mr. Dennis had access to all the code. He could run the regular expression. He would get the exact same answer.

I suspect that if I were to have research associates review and taxonomize the files I see in his opening report, they would come to different conclusions. I suspect if Mr. Dennis sent those same files -- the 6,847 in the case of Exhibit 5 -- back to his own associates to categorize, that they would come up with different categorizations.

So in this case I favored a method that was transparent and replicable.

Q Okay. I just want to break that down. I appreciate the response. I want to make sure I have what you were saying.

You didn't consider the Claude refinement

67

to be providing a source of ground truth?

A That's correct. I used Claude to propose files that might be characterize -- might be examples of false negatives. And kind of as suggestions for me, as there -- there might be another common way that this sort of intention of seeking any amount or form of lyric regurgitation, there might be a common way that that is phrased that wasn't captured by my original regular expression, and so those suggestions could be used -- were -- were used to adapt the regular expression to capture that sort of expression, but in a way that didn't induce too many false positives. But it didn't serve as the ground truth.

Q And I think you said "to determine the false negative rate, you would need to apply that level of scrutiny," and I think you were discussing it in the context of Mr. Dennis's analysis and Mr. Lustig -- or sorry -- Dr. Lustig.

So do I have that right, that to determine the false -- the actual false negative rate, you would need to manually review all records in the -- in the valid, adjusted sample?

A Are you asking if I wanted to get the false negative rate for, like, this last query to

68

Question 1, Version 2, yes, it -- like, getting the precise false negative rate, you would need two things. You would need to run it over all of the files, and then you would need to have a source of the ground truth. And to the extent that you don't have either of those, it's an estimated error rate.

And you know, probably the strongest form of -- if you want to find a classifier that has the lowest error rate, I would say it's an expert in music lyrics. And -- and any -- so you're -- you know, you need that.

And for the first part, you know, running it over -- if you have to estimate this on -- if you don't have Mr. Lustig or another expert look over the 2.5 million files, you need to subsample those 2.5 million files. And the fact of the matter is, this is a very rare type of phenomenon, and that has its own challenges with estimations. So there are actually just -- just in trying to -- so let me back up.

The precise provision of the false negative rate here is extremely challenging because you need to run it on all 2.5 million, roughly, in the adjusted sample, and you need really the ground truth. And it -- but in reality, you have estimates

Pages 73-80

73

Q And do you believe for Research Question 1, that the number of false positives matches the number of false negatives?

A That was my -- I would say the likelihood that it exactly matches, I think that's unlikely. I think that they're roughly of the same magnitude. I do believe that.

Q Okay. And what did you do to assess whether the false negatives and false positives were roughly equivalent?

A So there were two sides of it. I would see -- assess false positives by -- after running a given regular expression, looking at the results, and determining: Was this flooded by false positives?

Let me give you an example of this. This is on page 3 of the opening -- sorry. This is on page 3 of Table A in my opening report where the regular expression -- the final regular expression here is -- is written out.

And I want to draw your attention to in particular -- actually, the -- close to the very end of it, kind of this middle bit, we have the creation of something called "object." And then you have a regular expression this is the -- what -- what they

74

are seeking.

And at the end you have this syntax that -- towards the end you see words, but that is preceded by some syntax that includes "slash slash D."

And what this is articulating -- this -- this is actually a really good demonstration of considering false negatives and false positives. So Claude flagged for me -- you know, Claude is doing a much more expansive search than I could do manually, 46,000 records that it combed through and served some up to me as examples records that might be, you know, following the directions of the particular prompt that I gave it.

And a very common thing in that was just the inclusion of asking for words. Because if you look at my original regular expression, we had words, but the words was tethered directly to the request. So it was, like, "Give me the words."

But after adding this optional two-word buffer, if you run the regular expression of that and you have words in there, what you find is a really common use case that I ran across. I'm not going to recall the exact parameters of it, but something to the tune of "write me a" -- "write me

75

300 words that."

They almost look like a student with an essay prompt or something like that. Some of them also look, like, programmatic queries because they were very -- there were many of them following the exact same syntax that may have been an example of an API query that's, like, programmatically trying to get consolidated -- you know, reviews of a particular larger document.

So that was an example of being -- trying to make it more flexible by adding the optional two-word buffer to capture the false negatives that Claude had surfaced after looking at 46,000 --

THE COURT STENOGRAPHER: Wait.

THE DEPONENT: Sorry.

A -- and finding that there was a form of people -- the way that people are using Claude that was responsive to that -- the original regular expression, but that was all false -- false positives in this case. Like, it is easy to say these were not people that were seeking any amount or form of lyric regurgitation.

So that was adjusted to -- that thing that precedes it essentially just says: You can use words -- you can seek words, as long as it's not

76

immediately preceding a quantifier, because that's a -- that was a common false positive rate.

There was analogous one with -- with Bible versus where after putting the optional two-word buffer in there, we're getting lots of requests for Bible verses. You know: What verse talks about charity? What verse talks about -- so that regular expression permits still the inclusion of verses, as long as it's not preceded by "Bible."

Q (BY ATTORNEY HOWELL) Okay. And I believe you testified earlier you have used regular expressions prior to this case?

A That's correct.

Q Have you ever used, prior to this case, regular expressions to determine if a document contained a request for lyrics?

A Prior to this case, no.

Q And prior to this case, have you ever used regular expressions to determine if a document contained regurgitation of lyrics?

A Prior to this case, I did not.

Q Are you aware of any peer-reviewed publications that discuss using regular expressions specifically to identify whether something is a request for lyrics?

77

A I wasn't able to identify examples of publications where the regular expressions were used for lyric regurgitation.

I found examples of this notion of using, because it's fairly standard, regular expressions to comb through a corpus of documents and identify a subset of those documents that exhibited a particular characteristic and then thinking about ways to do that in a manner that would balance the precision of the recall to kind of have a classification model that's robust.

Q Are you aware of anyone else who has used regular expressions to identify whether something is a request for lyrics?

A I -- I'm not aware -- I didn't have any examples in a peer-reviewed literature of -- of such thing, and I'm unsure if Mr. Dennis ran my code. If -- if he did, then he would have been using it in a similar capacity.

Q Okay. But outside of peer-reviewed literature and -- or actually, sorry.

Outside of potentially Mr. Dennis, you're not aware of anyone else who has used regular expressions to identify whether something is a request for lyrics?

78

A That's correct. This particular, you know, application, "prompts that seek any amount or form of lyric regurgitation" is quite specific, and so I'm not -- I'm not aware of that exact use case being used before, but I'm certainly familiar with cases of -- using regular expressions to comb through a body of documents and find a subset that exemplifies some sort of behavior.

Q And we talked earlier about the fact that you excluded certain records in the sample from your analysis, correct?

A That's correct. There was a subset of the five million samples that had certain characteristics that made sense to exclude from the prevalence statistics that were reported in the research questions.

Q And by the time we get to -- I'll direct you to your reply, paragraph 62. And you say:

"After making these adjustments, the new set of valid records were the corrected adjusted sample and included 2,205,304 files."

Do you see that?

A Yes.

79

Q So that's the population of records that you ran your regular expressions on?

A That's the set that I used to run the regular expressions for Question 1 and Question 3.

Q Okay. Did you use a different set for Questions 2 and 4?

A 2 and 3 used a subset of that set.

Q Okay. Yeah. And so you excluded just under 2.8 million records from the sample from your analysis?

A That's correct. The 2.8 million were found to have a few characteristics that, if used, would not allow me to estimate these prevalence statistics. That was the whole goal of the analysis.

Q But you excluded a majority of the sample from your analysis; is that fair?

A I excluded the majority of the analysis -- well, I excluded 2.8 million files while calculating the prevalence statistics. And that's because -- for -- well, for many reasons.

But those 2.8 million files -- for example a huge chunk of them were title records, which don't -- are not the original record of a user's prompt. It's kind of an echo record that's created

80

on the back end by Claude that will have the same text, but is not a unique illustration of that, in this case, kind of lyric regurgitation-seeking behavior.

So if they hadn't been excluded, you would -- this would be overcounting. And there are -- there are other characteristics of files that similarly, had they been included, would make these -- would make it not possible to actually estimate the prevalence of prompt -- of user behavior in these prompts and outputs.

Q Okay. But ultimately you would agree you excluded the majority of the records in the sample from your regular expression analysis?

A I excluded a majority of the records from the regular expression because, had I included them, it would have led to very misleading results.

Q And you excluded non-English files, for example?

A That's correct.

Q You describe in your reply, a vendor issue regarding metadata that was provided to you?

A That's correct.

Q I just want to understand. You had meant to run a language detection software or have --

Pages 97-100

97

Q And for Question 1, did you distinguish between partial and full requests for lyrics?
A One moment. Could you -- I'll find it. Let me just get to my opening report.
ATTORNEY PATEL: For that pending question, objection, vague.
A Here we go. Okay. To the part of my report that references this, could you ask your question again?
Q (BY ATTORNEY HOWELL) Sure.
Did you Research Question 1 distinguish between partial and full requests for lyrics?
ATTORNEY PATEL: Objection, vague.
A So the text of the research question is "what percent of prompts seek any amount or form of lyric regurgitation," and the regular expressions reflect that.
Q (BY ATTORNEY HOWELL) So you were seeking to capture both partial and full requests for lyrics?
ATTORNEY PATEL: Objection, vague.
A I was seeking to capture any amount or form of lyric regurgitation.
Q (BY ATTORNEY HOWELL) Is that -- sorry.
Is that different from what I said?

98

A Yes. So I believe you just said "lyrics" and not "lyric regurgitation."
Q Okay. So thank you.
So were you seeking to capture both full and partial request -- requests for lyrics regurgitation?
A Yes, any amount or form of lyric regurgitation.
Q Okay. And in your analysis, did you identify any true positives that -- relating to Question 1 that Mr. Dennis didn't find?
ATTORNEY PATEL: Objection, vague.
A Yeah, could you state your question again? So I'll let you...
Q (BY ATTORNEY HOWELL) Sure. As part of your -- or strike that.
You're aware that Mr. Dennis also assessed whether records contained requests for lyrics?
A I believe he had a different definition than the definition than I operated under.
Q But you understand he was also looking at requests for lyrics -- lyrics to songs?
A I interpret Mr. Dennis's approach as him being provided with a set of documents that contained references to lyrics and songs. And his

99

task was to taxonomize those and categorize them into different ways that users in this particular subsample of documents that he has provided, how he would categorize their behavior.
Q And did you identify any true positives responsive to your Question 1 that Mr. Dennis didn't consider?
ATTORNEY PATEL: Objection, vague.
A So here when you say "Mr. Dennis did not consider," my -- I'm not sure what -- what do you mean by "his consideration"?
So his -- it -- because he looked at several datasets. There was some overlap between them. So I guess it would be helpful if --
Q (BY ATTORNEY HOWELL) Sure.
A -- you could tell me precisely which analysis of his you're referring to.
Q Sure. Did you identify any true positives responsive to your Question 1 that were not in the group of records from the sample that Mr. Dennis analyzed?
ATTORNEY PATEL: Objection, vague.
A So here are you referring to -- he has an Exhibit 5 that's a set of 6,849 that he purports were drawn from the five million sample -- selected

100

for him from the five million sample.
And is your question whether one of my true positives -- so a document that was in the numerator of the proportions that are provided in Section 1 -- whether or not one of those true positives was a part of his 6,849, his Exhibit 5?
Q (BY ATTORNEY HOWELL) Are you aware of any true positive that was not included in that set of --
A Ah --
Q -- 6,000?
A -- I -- I didn't -- didn't perform that analysis.
Q But sitting here today, you're not aware of any particular document that meets that criteria?
A I -- I didn't run that analysis.
Q How did you design your initial regular expression for Question 1?
A Let me go to the original.
So the original regular expression, it was informed by the research question, which is to say the prevalence of prompt seeking any amount or form of lyric regurgitation, just kind of a sense of what that would look like.
And then also the first phase of my data

Pages 125-128

125

and the output side, but at least one of them, I think it was on the prompt side at least was requests for translation of lyrics.

Q And Mr. Dennis also assessed requests for generation of new -- new songs?

A Yeah, I don't have his exact criteria in front of me, but I do recall something about, yeah, new lyrics versus existing lyrics in here. I guess that would be new lyrics.

Q And you didn't assess the number of records in a sample that contained requests for generation of new lyrics, correct?

ATTORNEY PATEL: Objection, vague.

A Yeah, my research questions were confined to this notion of lyric regurgitation. And you know, here lyric regurg -- if, say, Claude is performing lyric regurgitation, I'm thinking about this as, you know, he's -- he's -- the algorithm is providing text that contained lyrics.

And these lyrics come from the -- kind of the training process and the -- the understanding of the weights of Claude itself, the weights of the Claude model. And that the result of those lyrics is kind of the -- the content of the lyrics themselves as opposed to an incidental inclusion of

126

lyrics pursuant to, say, Claude trying to describe the meaning of a song or a history of a song or something like that.

Q (BY ATTORNEY HOWELL) Okay. And you didn't seek to assess the number of records in the sample that used song lyrics in those broader ways that were assessed by Mr. Dennis?

ATTORNEY PATEL: Objection, vague.

A Yeah, so if you could point to what you mean, like, in his report -- I may have it in my rebuttal -- what he referred to here as "broader ways," because the -- the -- what I was focusing on was any amount or form of lyric regurgitation, as I have described to you my kind of definition of what that is.

Q (BY ATTORNEY HOWELL) Take a look at your rebuttal report, paragraph 81.

A Oh, great. Thank you. It has the --

Q And do you see that Mr. Dennis assessed, for example, requests for -- Number 8, "Request for new song lyrics based upon provided artist(s), song title(s), or lyric(s)"?

Do you see that?

A Right. Yes, I can see that.

Q You didn't assess the number of requests

127

for new song lyrics based on provided artist(s), song title(s), or lyric(s) in the -- in the sample, correct?

ATTORNEY PATEL: Objection, vague.

A Yeah, give me a moment to read through what he may have been trying to capture here.

Yeah, so it's difficult -- you know, his categories don't perfectly overlap with my categories, so it's kind of difficult for me to ascertain the extent to which this corresponds to the particular question that I was looking at because I'm -- this is a challenge about -- his method is -- we have these, you know, 12 words that describe what's going on here, but we don't have the -- kind of the rules by which his reviewers were actually deciding whether or not a document fell into this category or not, so I guess I can't be clear whether or not my regular expressions would capture, or not, the documents that he was seeking to capture here.

Q But you weren't seeking to assess the number of requests for new song lyrics based upon provided artist(s), song title(s), or lyric(s) as part of your analysis?

A The focus of the analysis was on lyric

128

regurgitation. And you know, "lyric regurgitation" defined as the provision of lyrics whole hog and that was the -- kind of the objective of Claude's output -- that's the core of the output that was provided. If there were lyrics that were provided that were incidental to some other goal, then that would not have -- I wouldn't consider that lyric regurgitation.

So the extent to which your understanding of "new song lyrics" falls into the "not lyric regurgitation" camp, then that was not what I was seeking, seeking lyric regurgitation.

Q You considered creating original songs to not be lyric regurgitation, correct?

A Correct. I guess -- so like, the creation on the prompt side -- so somebody who is seeking any amount or form of lyric regurgitation, they're -- these are looking for kind of regurgitation of -- of lyrics from the model architecture into -- into the output.

Q Okay. And by "lyric regurgitation," you're talking about existing -- lyrics to an existing song?

A Here I'm just returning to the precise wording of the questions.

Pages 153-168

153

order to be responsive to Question 2?
A That's incorrect.
Q It's incorrect? It needed only one of those fields?
A Yeah, so could I direct you to -- this is Table 1 of Table B.
Q Sure.
A So you see this code here under the "Q2" block, if you look in particular at this "combine artist and song into an alternating regular expression."
That paste, what it's doing there it's taking a long list of all the artists' names and a long list of all the song titles, and it's concatenating them into a single long list that are each separated by that vertical bar that you see.
And in a regular expression that vertical bar is "or." So in order to pass through this phase of the regular expression, you would have to match this version of the first artist's name or this other version of the artist's name or this other version of that same artist's name. Or the second artist, this version of their name. Or another version of that second artist's name, exhausting many different permutations of artists' names. Or

154

you can match the first song title of a work-in-suit or this variant -- variation on the song title or -- and so on and so forth for all the works-in-suit -- song titles of the works-in-suit.
Q So if someone said "Give me the lyrics to "American Pie" without specifying it was by Don McLean, your understanding is that would be covered by this regular expression?
A I would have to, like, check, you know, the exact syntax of what you just told me to be sure that it matches the characters that would appear in text, since this is, you know, based on that, but I guess in general just providing the title would be sufficient to be responsive to Q2.
Q Okay.
A And in this way, I guess -- you know, this was a trade-off. Similar to the others of thinking about false positives and false negatives and regular expressions, some of the song titles in particular in -- that are part of the works-in-suit are quite generic song titles.
And under this regular expression, you know, it's being responsive if -- you know, if it says "just that song title," even if it says, you know, in a subsequent thing "by," this is actually a

155

different version of -- this is a different song, has the same title, but by a different artist; it's not a work-in-suit, it would still be responsive.
So in that way even that component, by allowing for only one of the two, would have some amount of false positives, but there was the worry that, you know, most people will not provide both, so requiring both would incur more false negatives than the false positives that we would incur by just requiring one of the two.
Q And you ran your regular expression for Question 2 on the records that were responsive to Question 1?
A That's correct.
Q And would you agree that any errors in Question 1 would affect your analysis in Question 2?
A So I would say that the -- I mean, the formulation of the regular expression is -- you know, it is largely recycling the same regular expression in the first, so that same -- the same files that were missed, say, as false negatives are included as false positives would -- you know, that -- that same issue would perpetuate through to Question 2.
Q And if there's a false negative for

156

Question 1, you wouldn't search that file as part of Question 2, correct?
A That's correct.
Q And if there's a false positive for Question 1, that means that record does not actually contain a request for lyrics -- lyric regurgitation?
A That's right.
Q And so that -- by definition, that false positive can't also include a request for lyrics to a work-in-suit?
A That's --
ATTORNEY PATEL: Objection, vague.
A -- so the -- if a -- the records that have made it into Question 2, the roughly 130 -- 1,034, are all responsive to Question 1, and some of those are false positives, the -- it's still possible that it would be responsive to Question 2 because we know at least that the first component of the regular expression would match.
Because the regular expression used in Question 2 is largely just a concatenation of that using Question 1 -- Question 2. We already the first component has passed, so in that way, they would be guaranteed to match that component.
But if -- you know, most false positives

157

that aren't seeking any amount or form of lyric regurgitation, I think it's unlikely that it would match the second component.

Q (BY ATTORNEY HOWELL) Okay. And your Question 3 analysis did not rely in any way on your Question 1 and 2 analysis, correct?

A The -- I mean, informed -- they're very similar analyses. The distinctions were -- Question 3 is focused on the outputs, on prompts, so there it's looking at kind of a different component of the data.

But the -- you know, the -- the method of using regular expressions, of using Claude to refine and the -- those 46,000 records that were used by Claude, those were shared across all of the research questions.

Q Okay. And if you turn to paragraph 122 of your opening report, the last sentence says:

> Based on my review, when Claude provided any amount or form of lyrics in an output, the output would start with the phrase "Here are the lyrics to."

Do you see that?

A Yes, I see that.

158

Q And do you stand by that statement?

A No. I think this is a misstatement.

Q How so?

A I guess there are few things here. I would say this is the most common of the form that I saw, but it wasn't the only.

And I have seen examples, well, that don't use this language and that the regular expression reflects a much broader sense than -- than is reflected here.

But the other distinction is that the output would start with this phrase, so that's -- that doesn't reflect my current understanding. I have seen examples that, you know, are responsive to the regular expression for Question 3, but where that component of the output is not actually at the start of the output, but also in the output.

And the regular expressions actually throughout the four research questions, none rely upon any sort of positional -- absolute positioning of where in the full text document the regular expression -- the matching component of the regular expression would live.

So here the lyrics would qualify for Question 1. It's not the only thing that would qualify for it. And I have seen examples that use different phrasing, and they confined it anywhere in the text string. It doesn't have to be at the beginning.

159

Q Okay. So this statement does not include an assumption that you rely on as part of your analysis?

ATTORNEY PATEL: Objection, mischaracterizes the witness's testimony.

A Yeah, I guess I would say that this just is a -- based on my current understanding is an inaccurate statement.

Q (BY ATTORNEY HOWELL) Okay. And your understanding is your regular expression for Question 3 broader than this one phrase?

A It certainly is, yeah. I don't recall -- at some point I think I may state the number of different kind of static phrases that could be matched. But it's far more than one. It's many different phrases that could be matched, and they could match anywhere in the output -- the first paragraph, the second paragraph, in the middle of the paragraph, at the very end.

Q Did this understanding inform how you designed your regular expression for Question Number 3?

A So I guess I would say that, you know, the -- designing the regular expression for Question 3 took a similar form as Question 1, which is to say I'm looking at the content of the research question, thinking about how Claude might do that. I'm also looking through the database at examples of records where Claude is providing lyrics and noticing patterns.

And those most common patterns are what are expressed in the first version of the regular expressions for Question 3.

Q And you performed a -- you used Claude as part of the -- refining your Question 3, regular expression, correct?

A That's correct.

Q And the prompt you used with Claude is -- as part of this Question 3 refinement is on page 36 of your opening report.

Do you see that?

A Yes, this was the system prompt that was sent to Claude.

Q Okay. And who drafted this system prompt?

A I drafted this system prompt after I went through a few versions of it. This was primarily

160

161

written based on kind of my understanding of how LLMs work in particular using an LLM essentially for data analysis, asking for a categorization, a classification in reply.

And then while looking through the documentation, Anthropic has documentation with suggestions for how to write effective prompt -- system prompts specifically. And if there are classification tasks, to give it actual examples, positives and negatives, to kind of ground some of its reasoning. That -- that can improve the overall effectiveness of the classification.

Q Did anyone assist you in drafting this prompt?

A I likely discussed it with counsel. I can't remember at what stage the -- I drafted the initial version of it. I revised it after looking through the documentation to include the examples for and against.

And at some point I shared it with counsel. I don't recall at what point. And I don't recall if their feedback led to any changes.

Q Do you recall if any language in this prompt originated from counsel?

A I would say -- I mean, this notion of,

162

like, "use regurgitation" is a term that we have used throughout this analysis. And I don't recall if -- if I was the first one to use that word or if they were the first ones to use that word, but that is certainly a component that is -- I guess would stand out as something that we have talked quite a bit about.

The rest of it, I don't suspect any -- any or much of it would have originated with counsel.

Q Okay. And as part of your prompt to Claude, you wrote:

"A message is considered to be regurgitating lyrics if it consists primarily or entirely of consecutive lines from a song, whether verbatim or with minor variations."

Do you see that?

A Um-hum.

Q Okay. Would you agree that that is instructing Claude that lyric regurgitation requires something more than just any amount of lyrics from an existing song?

A I think --

ATTORNEY PATEL: Objection, vague.

163

A -- yeah, there's -- I think of it as two separable things. There's the notion of, you know, what is -- how am I thinking about regurgitation and how am I thinking about any amount or form.

And regurgitation -- the instructions that I'm trying to give to Claude is so that it can capture, by definition of regurgitation here, which is not, you know, the wholesale provision of lyrics from the architecture of Claude for the sole purpose of providing those lyrics.

And kind of one of the reasons for that is, you know, there are other examples of Claude output that include lyrics, but where it's not lyric regurgitation. And you know, I think this happens in a few cases.

One way that it can happen is, you know, it's -- it's -- a user requests, you know, "What is the name of the song -- "what is the meaning behind this particular song?"

And in responding to it, it might, you know, use the song; it might use a snippet of lyrics from the song to -- in furtherance of that objective of providing a sense of a meaning. I don't want it -- I don't want Claude to be giving me those. Those wouldn't be considered regurgitating lyrics.

164

The other thing that I wanted to guard against with Claude in terms of false positives was that -- you know, the lyrics of songs tend -- tend to have quite a bit of overlap. So in all of human language, you know, there's a particular distribution of words and phrases that are used.

Within songs and song lyrics, there are certain words, certain lyrics that come up far more often like "love" and "life." And there could be sets of those words that seem kind of song-like because they -- they share these -- you know, those words that are commonly found in song -- song lyrics, but aren't actually song lyrics.

So those provisions that you see here describing the affirmative case and the kind of negative case all are trying to get at this notion of: Regurgitation is providing lyrics for the sake of the lyrics and not kind of incidental -- the provision of lyrics incidental to some -- to some other use case.

Q (BY ATTORNEY HOWELL) So for you, lyric regurgitation requires Claude providing lyrics for the sake of the lyrics?

A I would define a -- so I would define "lyric regurgitation" as it is providing sets of

165

lyrics where its goal is to provide just those lyrics. And what I would not consider lyric regurgitation is if it is -- say a user has asked for a song meaning, song history, and it provides that, and in doing so, excerpts a bit of the lyrics, that would not be, to me, regurgitating lyrics.

Q And so this description of "consisting primarily or entirely of consecutive lines from a song," that's consistent with your understanding of what it means to be regurgitating lyrics?

A I think this captures kind of the cleanest description that I could give Claude that would capture this notion of what lyric regurgitation means without inducing, you know, too many false positives.

Q Okay. All right. Turning to Question 4, you assessed "what percentage of outputs provide any amount or form of lyrics to the works-in-suit," correct?

A Right.

ATTORNEY PATEL: Objection --

A The percentage of --

ATTORNEY PATEL: -- just hold on one second. Objection, vague.

You can proceed.

166

A -- so the goal of Question 4 was to:

"Assess the prevalence of outputs that were providing any amount or form of lyric regurgitation to the works-in-suit."

Q (BY ATTORNEY HOWELL) Okay. And as part of this, you performed an N-gram analysis?

A That's correct.

Q And that analysis required -- for a record to be responsive, required there to be a match of at least 5 percent of the N-grams from a given work-in-suit; is that fair?

A Right, so let me describe exactly what this "5 percent" means. If the -- each one of the works-in-suit was divided up into N-grams -- four grams in particular -- kind of short snippets of the song that were overlapping and for every -- any given output that was responsive to Question 3, in the N-gram phase of this analysis, I checked to see -- for each one of those songs, I have the set of N-grams, and what proportion of those N-grams can be found in the output. And if that proportion exceeded 5 percent, then it was classified as responsive to Question 4 by the N-gram component of

167

this methodology.

And the threshold of 5 percent is used for a similar reason that the -- I don't recall the -- the baseline noise ratio that Dr. Lustig talked about in his report, this notion that there's just this, like, baseline similarity in language. And that if you don't set a threshold, say, of "a single N-gram matches," if that were the criteria, this would be a very faulty technique that would lead to all sorts of false positives.

There's an example of this in my reply report. Here we go. Page 22. Yeah. So you know, one of the works-in-suit, Hurricane, has, as one of its N-grams, "no matter how many."

And so you know, that's one of the N-grams. So if that were the criteria for being responsive to Question 4, then you can imagine most of the use cases that would be responsive have nothing to do with lyrics from Hurricane because this is just too common a phrase in common English. So that's what the threshold was.

The lowest one that I could use, the most permissive one without getting into a zone where I'm just picking up on background noise. Nothing related to the particular content of these -- the

168

lyrics of the works-in-suit.

Q And did you assess different thresholds other than 5 percent?

A I don't recall. I likely did. There's a similarity between setting this threshold and determining how flexible the regular expression should be.

The higher the threshold, the less permissive it is, which is to say I set a very high bar. For those lyrics that I see in the output are actually matching lyrics to the works-in-suit.

The lower that it is, the lower the bar that I have for saying "this is this particular work-in-suit."

And so there likely was some amount of calibration there to determine the right balance there between -- I want to capture examples of the lyrics of the works-in-suit, but not to the extent that I'm catching a bunch of false positives as well.

And this -- the difficulties of this method was one reason for the use of the second method in addition to this one.

Q Do you recall testing any other specific thresholds?

Pages 177-184

177

Q Do you recall Mr. Dennis identifying this document?
A Right. Yes, I do.
Q And do you see the System Prompt Date at the top says "February 19, 2024"?
A February? Yes, that's what I see.
Q And so this would be a record that was created after November 12, 2023, correct?
ATTORNEY PATEL: Objection, vague.
A Yeah, I'll say this is a record -- assuming this is the -- I guess totality of the record -- is this the totality? It was just one prompt and one output?
That -- I think within a few days is -- that date is correct, because of the reasons that we talked about previously about using the system prompt as the authoritative date field. It has some error in it.
Q (BY ATTORNEY HOWELL) Okay. So this record was created sometime in February of 2024?
A Yes.
Q And that would certainly be after November 12, 2023, correct?
A It would.
Q And are you aware that "Viva La Vida" is a

178

work-in-suit?
A I am aware of that. And I'm also aware that, you know, in Question 4, it was specifically looking for "any amount or form of lyric regurgitation to the works-in-suit." So even if this is a work-in-suit, this doesn't appear to meet my definition of "lyric regurgitation."
Q Okay. Do you see the second paragraph under "completion?
A Um-hum.
Q It says:
"The Lyrics tell the story of a character who once ruled powerfully over others."
Do you see that?
A I do.
Q And it says, "I used the rule the world."
A I do.
Q And then if you go to the third -- the third bullet point there, the refrain:
"One minute I held the key. Next, the walls were closed on me."
Do you see that?
A I do.
Q Do you recognize -- or first, are you

179

familiar with the song "Viva La Vida"?
A Vaguely.
Q Okay. Do you recognize the lyrics "one minute I held the key; next, the walls were closed on me" as being part of "Viva La Vida"?
A I couldn't say whether or not they are.
"I used to rule the world" seems somewhat familiar.
(Exhibit 4 was marked for identification and is attached to the transcript.)
THE DEPONENT: Thanks.
A Oh, there it is.
Q (BY ATTORNEY HOWELL) So you have been handed what's been marked as Exhibit 4. This is "Exhibit B to Publishers' Amended Complaint."
I believe in your opening report you said you relied on this document for the lyrics to the works-in-suit; is that correct?
A Yeah, I believe that -- I don't know if it was this document precisely. It was some version of this document that was -- you can see references in my code to having "red" in it when it's doing the N-gram matching -- matching.
Q And if you turn to the third actual page. I should say this is an excerpt of Exhibit B.

180

It's -- the pagination is page 155.
You'll see the start of the lyrics to "Viva La Vida"?
A Right.
Q And do you see the first line under "Genuine Lyrics" reads "I used to rule the world"?
A I do.
Q And if you go to the third paragraph, it reads "one minute I held the key; next, the walls were closed on me."
Do you see that?
A I do.
Q Do you dispute that Exhibit 3 contains lyrics from "Viva La Vida"?
A I -- I see the same sequence of words that show up in this output as appear in this document that you have provided on the "Genuine Lyrics" side. However, the -- the context of these lyrics is really important. So you know, that -- that bullet point that you pushed to -- point to to the -- that uses the phrase "I use the world" -- "I used to rule the world," it says, "The lyrics tell the story of a character who once ruled powerfully over others."
And then it's a parenthetical. So this -- this, to me, is a clear version of -- there are

181

lyrics that appear, but they're in the service of the goal of providing meaning.

And that's, to me, unambiguous from the first sentence. And the second sentence of the output, it says "Viva La Vida by," and then there's some retracted -- "is generally interpreted as being about the rise and fall of power and success. Here is a quick summary of the main themes and meanings."

So not a summary of the lyrics, but it's -- this is all centering around the meaning of lyrics --

Q Okay.

A -- sorry -- meaning of a song.

Q And I just want to break it down.

You agree that Exhibit 3 is a record that was created after November 12, 2023?

A I agree.

Q And you agree "Viva La Vida" is one of the works-in-suit?

A That's my understanding.

Q And you agree that Exhibit 3 contains some of the same lyrics as appear in "Viva La Vida"?

A I agree it contains some of the snippets of the lyrics from the full lyrics to "Viva La Vida."

182

And it also contains the context of providing those lyrics very unambiguously that are provided to provide, you know, kind of context to the intent of the -- the theme of the overall output, which is about themes and meanings, as is shown in the output.

Q Okay. And because Exhibit 3 contains lyrics as context discussing the theme of this particular work, you do not consider this to be regurgitation of lyrics?

A That's correct. The regurgitation of lyrics -- something would not be considered lyric regurgitation if it is incidental to some other aspect of the song or the lyrics. In this case the meaning and the theme is -- in fact, it's a good example of how a document can contain snippets of lyrics and yet not be, at least in my definition, regurgitating lyrics, like trying to whole hog say: Hey, here are all the lyrics.

Q Is your definition of "lyric regurgitation" based on any understanding of the legal standard for copyright infringement?

A I have no particular expertise in law or understanding the shades of distinction of what would and would not constitute copyright infringement. It is just kind of based on this notion of -- general notion of lyric regurgitation.

I know it has some bearing on legal questions of copyright infringement, but I have no expertise on where that line lies.

183

Q Okay. You can put that aside. If you turn to page 50 -- or sorry -- paragraph 52 of your opening report.

A Sorry. 52?

Q Yes.

A Yes.

Q And there's a footnote there, Footnote 22. Do you see that?

A 4 -- 48?

Q It should be paragraph 52. It's on page 14 of your opening report.

A Sorry. Okay. Yeah. If you could give me just a moment to read it. Prompt output pair. All right. Okay.

Q Okay. The first sentence in Footnote 22, you write:

"I understand that during the sample's time period of September 22, 2023 to March 22, 2024, it was possible for a user to

184

put words in Claude's mouth or prefill part of Claude's responses -- responses in order to shape Claude's response."

Do you see that?

A I do.

Q And you cite a web page here.

Are you relying on anything other than the website listed here for this understanding?

A Right, so this documentation describes exactly this feature. I believe this also came up in discussions with the data -- the Anthropic data scientists. This -- you know, was this a -- indeed was this a feature? And do we have metadata that would track whether or not a user was using this feature of the model?

Q Okay. And anything else that you were relying on for this understanding?

A I believe that's the extent of it.

Q Okay. And the website you're -- let me ask first -- you mentioned talking to Anthropic data scientists?

A That's correct.

Q Who did you speak with at Anthropic?

A One of their names was Alfie. I don't