Exhibit 20



Planet Depos®
We Make It Happen™

HIGHLY CONFIDENTIAL
ATTORNEYS' EYES ONLY

Transcript of Vinay Rao, Ph.D., Designated Representative

Date: November 7, 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
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
Anthropic PBC, by and through its Designated
Representative,
VINAY RAO, Ph.D.
San Francisco, California 94111
Friday, November 7, 2025

Reported Stenographically by:
MARY J. GOFF
CSR No. 13427
WA CSR No. 21030779
Job No. 606055A
PAGES 1-159

2

HIGHLY CONFIDENTIAL--ATTORNEYS' EYES ONLY

VIDEOCONFERENCED and VIDEO-RECORDED DEPOSITION of VINAY RAO, 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 9:18 a.m. and ending at 2:54 p.m., on Friday, November 7, 2025, before MARY J. GOFF, California Certified Shorthand Reporter No. 13427 and WA CSR No. 21030779.

3

APPEARANCES:

For Plaintiff Publishers
Oppenheim + Zebrak, LLP
BY: COREY MILLER
TIMOTHY S. CHUNG (NYC office)
Attorneys at Law
4530 Wisconsin Avenue
5th Floor
Washington, D.C. 20016
202-957-5043
corey@oandzlaw.com
tchung@oandzlaw.com

For Defendant Anthropic PBC
WilmerHale
BY: TAYLOR GOOCH
Attorney at Law
50 California Street
Suite 3600
San Francisco, California 94111
taylor.gooch@wilmerhale.com
628-235-1026

ALSO PRESENT: Ian Chen, In-house counsel Anthropic
Ellen Stephano, Esquire, WilmerHale (remote Boston)
Videographer: Philip Astor

4

INDEX

WITNESS		EXAMINATION
VINAY RAO, Ph.D.		
Volume I		
	ATTORNEY MILLER	7, 161
	ATTORNEY GOOCH	149

NUMBER	DESCRIPTION	PAGE
Exhibit 80	Usage Policy Anthropic_000437165-170 Tab 47	21
Exhibit 81	Anthropic's Consumer Terms of Service, 11-28-2025 Anthropic_0000437176-190 Tab 48	24
Exhibit 82	Acceptable Use Policy Anthropic_0000013675-679 Tab 46	27
Exhibit 83	4-27-23 email, Sanderford Tab 4 Anthropic_00000010907	37
Exhibit 84	CLAUDE's Constitution 5-9-23 Anthropic_000000485-498 Tab 13	54
Exhibit 85	Defendant Anthropic PBC's Supplemental Responses and Objections to Plaintiffs' First Set of Interrogatories Number 12 Tab 78	59

Pages 45-48

45

ATTORNEY GOOCH: The same instruction.
And I will direct the witness not to answer.
(Instruction not to answer.)
Q (BY ATTORNEY MILLER) The question was whether -- was whether he's going to accept your instruction.
You're instructing him not to answer whether he's going to answer?
ATTORNEY GOOCH: Oh, I'm sorry. I misunderstood.
You were just asking if he's going to follow my instruction?
ATTORNEY MILLER: Yeah.
You should answer "Yes" or "No" if you're going to follow my instruction.
A I will follow the instruction.
(Instruction not to answer.)
Q (BY ATTORNEY MILLER) Are you aware that CLAUDE's outputs can sometimes contain copyrighted materials?
ATTORNEY GOOCH: Objection, scope; and misstates facts not in evidence.
A So let me state it the way I understand it, and then you can tell me if it aligns with the

46

question.
So the -- the way CLAUDE is set up is it has all of these detections and mitigations along in -- in -- on both sides, meaning at the prompt side and at the completion side.
So is the question that it is -- contains copyrighted information when CLAUDE generates or outside or like on -- outside of the mitigations?
THE COURT STENOGRAPHER: Did you say perceptions or detections?
A So it's both. It's the detection. It detects. And then a mitigation applies, and so that's how it protects.
Q (BY ATTORNEY MILLER) Okay. I'll clarify and ask a couple of separate questions.
Despite the mitigations that Anthropic uses to prevent CLAUDE from outputting copyrighted materials, CLAUDE is still capable outputting copyrighted materials, correct?
ATTORNEY GOOCH: Objection, vague; and scope.
A So there are -- there are mitigations at the prompt level, which very briefly -- and we can go into the details, if we want -- is a prompt classifier where the classifier flags or detects if

47

there is -- if the classifier was trying to figure out if the completion is going to have copyrighted content.
And then there is a mitigation associated with that. And there is a completion-level detection and mitigation. And it could -- that will also detect if there is copyrighted material there. And if that fires, then it will prevent the copyrighted material from being streamed further.
So for use of CLAUDE, these two are always turned on.
Q Are those -- are the mitigations you described 100 percent effective?
A In a system of this kind, it is hard to make anything 100 percent effective.

48

Q Are you aware whether CLAUDE -- withdrawn.
When you refer to "being streamed," do you mean included in an output that --
A Included in an output, yeah.
Q Okay. Are you aware whether CLAUDE has ever streamed or included in an output, copyrighted materials?
ATTORNEY GOOCH: Objection, scope.
A Yes, I am aware.
Q (BY ATTORNEY MILLER) And so there have been instances where CLAUDE included in an output, copyrighted materials, correct?
ATTORNEY GOOCH: Same objection.
A So is it possible? The answer is "Yes."
Is it easy? The answer is: It has to get past a series of detections and mitigations.
So it has to go through each of these mechanisms, which is kind of what we have thought of

Pages 69-72

69

Q (BY ATTORNEY MILLER) But right sitting here right now, you don't know whether that data exists; is that correct?
ATTORNEY GOOCH: Same objection.
A So again, when we talk about data existing in kind of any system, there are mechanisms that you do. You run a test, and that test has multiple objectives; and it -- you can draw multiple inferences depending on how the test is done. And also, there are a bunch of, you know, conclusions that are developed and written down.
What exactly was done to -- whether it answers your specific question is like two different people. Like, this is a question, and some other data may have been done -- computed and so on. So it is probably best to go back to the primary sources and look at it, if there is data that supports the precise question that you're asking.
Q (BY ATTORNEY MILLER) Do you know whether Anthropic has done any tests or studies to determine how effective constitutional -- excuse me -- to determine how effective fine-tuning on a constitutional copyright principle is at preventing users from writing creative prompts to get around more traditional auxiliary systems?

70

A I'll say the same thing: To get a precise answer to your question, we will have to go back to the -- to whatever evaluations that were done and if they support the statement that you are making.
Q Okay. Do you see at the end of line 20, this response cites to an "artifact"?
A Is this the statement -- the line that says "constitutional.r1 Artifact" --
Q I'll ask you not to read the actual citation to keep it out of the record.
A Oh, I see. Okay. Sorry.
Q But do you see that citation?
A Yes.
Q And is that a citation to a portion of CLAUDE source code?
A Yes, so I spoke with Kevin Kelley about this. So from my understanding of what he said is that this is a subset of the Git repo where this update is recorded.
THE COURT STENOGRAPHER: The Git? Subset of the?
A Oh. GitHub repository where this change was recorded.
Q (BY ATTORNEY MILLER) And who is Kevin Kelley?

71

A Kevin Kelley is a member of technical staff at Anthropic.
Q Do you know what his specific role is at Anthropic?
ATTORNEY GOOCH: Objection, calls for speculation.
A It's a specific question. I think it's best to check.
He was the person that I spoke to about the times and dates of the -- when these changes were being made as part of preparation of -- for various defenses that Anthropic had put up.
Q (BY ATTORNEY MILLER) Have you, yourself, reviewed the -- the source code cited in line 20 on page 10?
A This particular source code, no.
Q Prior to your preparation for this deposition in your conversation with Mr. Kelley, were you aware that this particular artifact cited on line 20 existed?
A So I'll go back to what I had answered before we got into this document.
I had -- I have -- I had read the constitution as part of my interview preparation, so I had a broad understanding. Do a line-by-line and

72

word-by-word precision, I don't remember.
Q Are you familiar with something at Anthropic called the "content compare system"?
A Yes.
Q Are you familiar with something at Anthropic called CopyFairy?
A Yes.
Q What is the relationship between the content compare system and CopyFairy?
A So they are -- the outcome of what it sought to do, they are the same.
The idea is to look at completions as they're being streamed and detect if there is a regurgitation of copyrighted material so that way they are similar in terms of how they -- what they do.
In terms of how they are -- how the engineering implementation is, the cease cont -- cont -- content completion system is an updated version of CopyFairy.
Q Why did Anthropic develop the content compare system?
A Just to clarify, we are talking about -- specifically about CCS not CopyFairy.
Q Thanks for that clarification.

Pages 81-84

81

that a match threshold or is there some other term by which that -- that requirement known?

A So from a purely engineering perspective, there are -- there are two sources that need to be compared.

One is what is -- what the model is generating in terms of tokens, and the other one is this CCS database, which is what I called the reference corpus.

And for every X number of tokens, you run a check against this -- all of this if those tokens exist in that order in the corpus.

[redacted]

Q Just to clarify, in your answer just now you referenced "tokens."

Is the -- was this version of CCS matching [redacted]?

A Let's -- let's stick with words. I misspoke. Let's use words, yeah.

Q [redacted]

82

[redacted]

A Substantially. The reason is there is a -- there is a bunch of checks -- or changes that we to -- to normalize it, so -- and it's step -- set above what the normalization is in line 20.

So we do remove a bunch of characters that could potentially cause it to not match because of -- it could become highly particular for that, so it generalizes a little bit more.

[redacted]

The -- other than that, yes, it is those words and matching them against the reference corpus, which is also similarly normalized.

Q What would be the result of the operation of the CCS system if the potential output contained [redacted]?

A So the entire -- like, just clarification. The entire output of CLAUDE at that time was less [redacted]?

Q Yes.

A Then it would not match it.

Q So for any copyrighted work that contained

83

[redacted]

excuse me. Withdrawn.

For any copyrighted work that contained [redacted], this version of the copyright -- excuse me -- the content compare system would never block an output that contained that copyrighted work; is that correct?

A So in the context of the overall set of defenses that are built, CCS is a part of it. And when the match is set to be less than -- set to be [redacted] But that doesn't mean there are no other guardrails to -- to protect.

Q What happens -- what is the result under this version of the content compare system if the [redacted]?

A Total output?

Q Do you understand what I mean?

A I'm trying to understand. Can you give me an example? Yeah.

Q Sure.

Let's say somebody prompted CLAUDE: Explain to me the meaning of the lyrics to a

84

particular song, and CLAUDE wrote a -- and the lyrics to the song are only 150 words long, and CLAUDE wrote a 500-word essay that included -- so [redacted] included the -- all 150 of the words in the song lyrics; [redacted].

Would this version of the content compare system have blocked that output, assuming that the song lyrics were in the reference corpus?

A Right. Okay.

So assuming that the prompt got through the prompt classifier, assuming that the model decided to include the contents -- because this is a part of the overall system -- and if the number of words set to match -- that's how it is configured is greater than the number of words in the lyrics -- then yes, it would miss.

Q For this version of the content compare system, [redacted]?

A I -- I will have to check to get the exact answer why it was set.

What I can say is this -- this is a fairly complex engineering system. It needs to be highly performant to be able to do this matching for every

Pages 93-100

93

whether there is copyrighted content or not.
And in a way that -- it is built in a way that it works on all the surfaces that Anthropic's models are -- are available. So the engineering work is -- I mean, it has to be a more performant system. Engineering is -- is a lift there.
ATTORNEY MILLER: Tab 18.
Q (BY ATTORNEY MILLER) If you'll turn to your Zoom, you should see what I am marking as Exhibit 86, which for the record is a spreadsheet bearing the Bates Number Anthropic_0000347810.
We made the native version available, which -- in the Zoom.
(Exhibit 86 was marked for identification and is attached to the transcript.)
Q (BY ATTORNEY MILLER) And you can -- let me know when you have that open.
Do you have the document open?
A Yes.
Q Okay. Do you recognize this document, sir?
A Yes.
Q What is this document?
A I mean, as it says, it is a description of
[redacted]

94

[redacted]
Q All right. Do you see there's also a Column D?
A Yes.
[redacted]
A Yes.
Q -- you see that?
[redacted]
Q Do you see the first row of content, which is Row 2 -- step back a second.
I'll represent to you that the metadata for this document indicates it was created on May 23, 2024.

95

Do you have any reason to doubt that?
A May 23? Doubt what? The date?
Q Yeah.
Does anything about this document indicate to you that it -- that -- whether that date is -- that that date is inaccurate or accurate?
A I don't understand the question very well. Like, this is --
Q Let --
A -- sorry. Go ahead.
Q -- yeah, let me ask a better question.
So the metadata indicates this document was created on May 23, 2024.
Does this document represent the contents of the CCS corpus as of that date?
A I will have to check to confirm if it is exactly that, and that would be really best done by looking at the code, yeah.
Q Do you know who created this document?
ATTORNEY GOOCH: Objection, scope.
A I will have to go back and check.
Q (BY ATTORNEY MILLER) Did you create this document?
A No.
Q And do you know for what purpose this

96

document was created?
ATTORNEY GOOCH: Objection, scope.
A I will have to check.
Q (BY ATTORNEY MILLER) Row 2, which is the first row after the title row, the first entry is
[redacted]
A Yes.
[redacted]
A I -- the second part of the question, I think, is privileged. It -- what it refers to is
[redacted]
ATTORNEY GOOCH: We don't assert privilege over that fact so --
ATTORNEY MILLER:
A Yeah.
ATTORNEY GOOCH: -- you can testify regarding that.
A Yes, it is.
Q (BY ATTORNEY MILLER) Okay. Let me ask it again just so we have a clear record of a question and answer because that got a little bit garbled.
[redacted]

97

[redacted]
A Yes.
[redacted]
on S3, it appears?
A Correct.
Q What is that a link to?
ATTORNEY GOOCH: Objection, calls for speculation; and outside the scope.
A I have to confirm. From an engineering perspective, it could be -- but it needs confirmation -- is S3 is a kind of storage system available on Cloud platforms. It needs to be confirmed.
Q (BY ATTORNEY MILLER) Okay. And do you see in the link there's a reference to "October 21 prompts"?
Do you see that?
A Yes.
Q Do you know what that refers to?
A I will have to check.
ATTORNEY GOOCH: I'm sorry. I should have said it out loud -- not that I thought it -- objection, scope. I realize that your instruction

98

to the witness to say things out loud should probably also apply to me.
ATTORNEY MILLER: I think the court reporter would appreciate that. She can't promise to transcribe thoughts.
ATTORNEY GOOCH: It would make my life easier.
ATTORNEY MILLER: I don't think you want her to.
Q (BY ATTORNEY MILLER) If you turn to Row 7,
[redacted]
A Yes.
[redacted]
A Oh, God. Wait a sec.
Yes.
Q Are these the 20,000-- withdrawn.
[redacted]
A Yes.
Q And then in Cell B7, there's again a link to an S3 storage location, correct?
A Yes.
Q And in part of that link there's a

99

[redacted]
Do you see that?
A Yes.
Q What does that mean?
ATTORNEY GOOCH: Objection, calls for speculation.
A We'll have to check with the person that created it, yeah.
Q (BY ATTORNEY MILLER) Okay. Do you know
[redacted]
ATTORNEY GOOCH: Objection, that calls for privileged --
A Privileged.
ATTORNEY GOOCH: -- information, and I instruct the witness not to answer.
(Instruction not to answer.)
ATTORNEY MILLER: Counsel, the question was: Do you know how?
ATTORNEY GOOCH: Oh.
ATTORNEY MILLER: See. let's see first whether he knows and then we can --
ATTORNEY GOOCH: I'm fine with that -- that plan. So I'm going to -- why don't you re-ask

100

the question.
Q (BY ATTORNEY MILLER) Sir, do you know how
[redacted]
ATTORNEY GOOCH: So, Doctor, only answer that question "Yes" or "No."
A No.
Q (BY ATTORNEY MILLER) Do you know who
[redacted]
ATTORNEY GOOCH: You can answer that "Yes" or "No."
A No.
Q (BY ATTORNEY MILLER) Look at Row 10, please. Do you see it says in Column A
[redacted]
A Yes.
[redacted]
A Yes.
[redacted]
A Yes.
Q And then in Cell B10, again, there's a

Pages 121-124

121

Resources" there's links for two documents?
A Yes.
Q And the first one -- the first link is to a document that appears to be titled "Steps for Training T&S Classifier"?
A Yes.
Q Are you familiar with that document?
A Without looking at it, I don't know exactly what that document is. But at a high level, that would be how to train the prompt classifier for trust and safety.
Q And the second document is titled "Problematic and Acceptable Prompts."
Do you see that?
A Yes.
Q Are you familiar with that document?
A Again, I don't know exactly what is in that document without seeing it.
[redacted]
Q Do you know who created the "Steps for

122

Training T&S Classifier" document?
ATTORNEY GOOCH: Objection, calls for speculation.
A Again, without seeing the document metadata, it would be hard to answer.
Q (BY ATTORNEY MILLER) Do you know who created the "Problematic and Acceptable Prompts" document?
ATTORNEY GOOCH: Objection, calls for speculation.
A I will have to check the metadata.
Q (BY ATTORNEY MILLER) And in the middle of page 1, where -- do you see it says "last updated December 4, 2023"?
A Yes.
Q And there it says:
[redacted]

123

Do you see that?
A Yes.
Q Why did Anthropic stop doing further iteration development in the context of [redacted] at this time?
A Okay. So this is where I was going to say the context. Around this time there were a bunch of different projects that were happening.
My team was short on staffing, so we had three engineers come over from another team, the three at the top of this document, to speed up the process of getting this classifier up and running.
So they developed the -- the classifier that is used to detect the likelihood of a prompt leading to copyrighted content in the output, so they built, tested; and at this point they wrote this document, and that's why this document is called "Transition," to transition off from what they were doing back to my team.
So they stopped working. That doesn't mean Anthropic or the -- at that time what was called the trust and safety to stop working.
Q And then there's a section of the document titled "Current State," correct?
A Yes.

124

Q And in the next paragraph, the document states:
[redacted]
Do you see that?
A Yes.
Q Did Anthropic do testing of the classifier to determine how effective it was at detecting direct requests for lyrics and mildly obfuscated requests for lyrics?
A I would have to go back and check in detail, but that's what is written in this document.
And that means you have both, like, simple prompts that say: Hey, give me this, versus some amount of obfuscation in the prompt.
Q And in the next paragraph, the document states that:
[redacted]