EXHIBIT 136



Planet Depos®
We Make It Happen™

HIGHLY CONFIDENTIAL
ATTORNEYS' EYES ONLY

Transcript of Dawn Drain

Date: November 12, 2025
Case: Concord Music Group -v- Anthropic

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

1

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

SAN JOSE DIVISION

CONCORD MUSIC GROUP, INC., ET AL., Plaintiffs, vs. ANTHROPIC PBC, Defendant.) Case No. 24-cv-03811-EKL-SVK

***

HIGHLY CONFIDENTIAL - ATTORNEYS' EYES ONLY

***

REPORTER'S TRANSCRIPT

VIDEOTAPED DEPOSITION OF

DAWN DRAIN

Wednesday, November 12, 2025

Stenographically Reported by:

Shannon Borge, RPR,
CSR No. 9436

Job No. 606064
Pages 1-231

2

A P P E A R A N C E S

---oOo---

ON BEHALF OF THE PLAINTIFFS:

OPPENHEIM + ZEBRAK, LLP
BY: BRET MATERA, ESQUIRE
COREY MILLER, ESQUIRE
461 5th Avenue
19th Floor
New York, New York 10017
212.951.1833

ON BEHALF OF THE DEFENDANT AND THE WITNESS:

WILMER CUTLER PICKERING HALE AND DORR LLP
By: TAYLOR GOOCH, ESQUIRE
ALYSSA MILSTEAD, ESQUIRE (Via Zoom)
50 California Street
Suite 3600
San Francisco, California 94111
628.235.1026

ALSO PRESENT:

Gary Garcia, Planet Depos Videographer

Ian Chen, Anthropic Counsel

---oOo---

3

DEPOSITION SUPPORT INDEX

QUESTIONS INSTRUCTED NOT TO ANSWER:

PAGE	LINE
211	8

REQUEST FOR PRODUCTION OF DOCUMENTS

PAGE	LINE
16	22
191	22

STIPULATIONS

PAGE	LINE
(None)	

QUESTIONS MARKED

PAGE	LINE
(None)	

REPORTER'S NOTE:

QUOTATION MARKS ARE USED FOR CLARITY AND DO NOT NECESSARILY REFLECT A DIRECT QUOTE.

4

I N D E X

---oOo---

	PAGE
EXAMINATION BY MR. MATERA.....	7
EXAMINATION BY MR. GOOCH.....	228

---oOo---

E X H I B I T S

EXHIBIT	DESCRIPTION	PAGE
Exhibit 135	LinkedIn profile	19
Exhibit 136	Blog post (Anthropic_0000013558)	58
Exhibit 137	"Studying Data Repetition Through Scaling Laws and Mechanistic Interpretability (Anthropic_0000520123)	59
Exhibit 138	11/09/23 PowerPoint (Anthropic_0000500599)	85
Exhibit 139	"Exact Dedupe for [redacted]" (Anthropic_0000510628)	131
Exhibit 140	08/15/23 "Deduplication Algorithms" (Anthropic_0000507750)	146
Exhibit 141	"Dawn's Thoughts On Data Priorities" (Anthropic_0000216875)	165
Exhibit 142	"pranay dawn 1:1" (Anthropic_00052912)	195
Exhibit 143	Slack thread (Anthropic_0000024220)	197
Exhibit 144	Memorization Prompt (Anthropic_0000242229)	199

Pages 57-72

57

working in the best interest of humanity and will follow instructions and generally be helpful and harmless.

Q. Do you know if Anthropic has any policies related to intellectual property rights?

A. Um, yeah, that's not my area of expertise. I don't -- I don't know what Anthropic's policies are.

Q. You don't know what Anthropic's policies are related to intellectual property rights?

A. No. I'm not -- I'm not a lawyer. I don't know what Anthropic's policies are related to IP rights.

Q. Okay. Broadly speaking, how does Anthropic train its AI models?

A. At a high level, there -- there is pretraining and fine-tuning, and pretraining consists of training from scratch on a very large amount of data, um, which is raw text that has been, um, written by people, and then after pretraining, there's fine-tuning, which mostly consists of reenforcement learning where we generate samples from the model in various environments and then give it feedback based on its samples.

Q. And the training on very large amounts of data that you referred to, that data's contained in the data dictionary you were describing earlier?

MR. GOOCH: Objection. Calls for speculation.

THE WITNESS: Um, yeah, I don't -- I don't know

58

that all the components are in the dictionary, but, um, um, yeah -- yeah, probably some of the components are.

BY MR. MATERA:

Q. Do you do any pretraining in your work?

A. Um, yeah. Like I said earlier, um, I've been working for the last few months on having Claude make these architectural changes, and then we train -- or, pretrain tiny models on -- on those modified codebases.

Q. And the modified codebases that you use, are those contained in the data dictionary?

A. No.

Q. Where are they located?

A. Um, um, so we'll save the diff to S3 or Google storage and typically also push it to GitHub.

Q. What do you mean "push it to GitHub"?

A. I mean upload to GitHub.

MR. MATERA: Okay. Okay. I'm going to introduce what will be Plaintiff's Exhibit 136. It bears the Anthropic Bates stamp 0000013558.

(Exhibit 136 was marked for identification.)

MR. GOOCH: Thank you.

MR. MATERA: Yeah.

BY MR. MATERA:

Q. Miss Drain, do you recognize this document?

A. Um, let me -- let me read it.

59

Q. Uh-huh. While you read that, I will also introduce and mark exhibit -- Plaintiff's Exhibit 137, which bears the Anthropic Bates stamp 0000520123.

(Exhibit 137 was marked for identification.)

BY MR. MATERA:

Q. Miss Drain, do you recognize Exhibit 136?

A. Yeah.

Q. What is it?

A. Um, this looks like a blog post linking to a paper primarily written by Danny on repeated data.

Q. Who's Danny?

A. Danny Hernandez is a coworker of mine.

Q. And does this blog post appear on Anthropic's website?

A. Um, yeah.

Q. Does it appear to be a true and correct copy of what it purports to be?

A. Yeah.

Q. What's the paper it's linking to?

A. Um, the paper titled "Scaling Laws and Interpretability of Learning from Repeated Data."

Q. You were a coauthor of that paper; correct?

A. Yeah.

Q. Do you know if this paper is available on Anthropic's website?

60

A. Um, I mean, it -- it's on arXiv and you can access it through Anthropic's website.

Q. And if you could look at Exhibit 137, please. You can put 136 away for now.

A. Um, this is going to take a while to read.

Q. Sure.

A. Sorry. (Reviewing document.)

Q. Have you had enough time to review, Miss Drain?

A. I'm on the last page. All right. I've -- I've reviewed the document.

Q. Do you recognize Exhibit 137, Miss Drain?

A. Yeah.

Q. What is it?

A. This is Danny's paper "Studying Data Repetition Through Scaling Laws and Mechanistic Interpretability."

Q. It's the paper referenced in Exhibit 136?

A. Yeah.

Q. And your name is listed among the authors; correct?

A. Yeah.

Q. Does this paper appear to be a true and accurate copy of that paper?

A. Um, yeah. Some of the -- the graphs are a little bit distorted, but yeah.

Q. Understood. Do you know when the paper was

61

published?
A. Um, I think it was published in 2021.
Q. Okay.
A. The date was on the -- on the blog post. Oh, sorry, sorry, sorry. On the blog post, it says May 1st, 2022. Sorry.
Q. No problem. You're referring to Exhibit 136; correct?
A. Yeah.
Q. What is this paper about?
A. Um, this paper, um, trains several models on datasets that have introduced synthetic, um, duplication, um, where a subset of the data is repeated many times, and this causes the models to overfit to that sub-distribution and, um, memorize the data in a way which is harmful for training. It makes the models, um, a lot dumber so it's pointing this out as something to avoid, um, so that we don't accidentally train on repeated data, um, so that you can avoid memorizing it and degrading model performance.
Q. Has Anthropic attempted to implement the conclusions of this paper into practice?
MR. GOOCH: Objection. Calls for speculation.
THE WITNESS: Um, which -- which conclusions are you -- you referring to?

62

BY MR. MATERA:
Q. The ones you were just describing.
MR. GOOCH: Same objection.
THE WITNESS: Um, we worked on trying to deduplicate our datasets, yeah.
BY MR. MATERA:
Q. And why do you work on deduplicating datasets?
A. Um, yeah, as -- as shown in this paper, if you do a poor job deduplicating a dataset, you can have a subcomponent which is repeated many times, and then if you train the model, um, on this dataset, it will overfit to the sub-distribution and it will make the model a lot dumber.
Q. This phenomenon you're describing can occur in Anthropic's AI models? Is that what you're saying?
A. Um, I am saying that any language model, um, at least shown in this paper, um, for -- for these transformer models, um, yeah, if they're trained on highly, um, duplicated text, like I said, they could orient to the sub-distribution and become dumber.
Q. Okay. What does "overfit" mean?
A. Um, so overfitting refers to a machine learning phenomenon where if you train for too many epochs on a dataset, you can continue to improve on your objective for the training data, but, um, this doesn't actually

63

generalize to your test set and so you're not actually improving the model. You're actually making it worse.
Q. What's an "epoch"?
A. An epoch is one pass over your dataset.
Q. What do you mean by a pass over your dataset?
A. Um, I mean it's training the model on every data point in your dataset once.
Q. You also mentioned the term "memorize" or "memorization." Do you recall that?
A. Um, yeah.
Q. What is memorization in an AI training context?
A. Um, there are a few different technical definitions of memorization. Um, the one that I'm most partial to is when you have a large span of tokens, all which have pathologically low loss or a pathologically high probability of being produced by the model, which generally represents a superficial form of understanding which was actually damaging to the model and, in particular, hurts its performance of generalization.
Q. You said memorization is damaging to the AI model; is that right?
A. Um, um, yep. Like I said, if a model or if it's to a piece of text or a sub-distribution, by memorizing it according to its technical definition, then generally that makes the model dumber.

64

Q. What do you mean by "makes the model dumber"?
A. Um, so the way that we normally define this is we -- because there's a notion of the compute efficient frontier. We're given a fixed budget of compute. There's an optimal way to train your model and an optimal loss you'll achieve on your objective function. Um, so we can use this as a measuring stick comparing various models. So, for example, um, if you made the model much dumber, it might shift it down so as if you only had half as much compute to train that model.
Q. And could you briefly explain what "compute" is?
A. Um, "compute" is normally expressed in terms of training FLOPs for the number of, um, FLOPs meant using your GPUs or TPUs or other accelerators.
(Court reporter requested clarification.)
THE WITNESS: GPUs or TPUs or other accelerators.
BY MR. MATERA:
Q. So "compute" relates to the model's computing capacity; is that fair to say?
A. No. In -- in this context, um, I'm referring to training compute, how much compute was used to train the model.
Q. And what's a FLOP?

65

A. Um, a FLOP is a floating point operation, like an add or a multiply.

Q. And so when a model memorizes something, that memorization has a tendency to impact the model's compute; is that correct?

A. Um, so the way that -- that I was trying to -- to frame things is, um, if -- if the model memorizes some sub-distribution of the text, according to the previous technical definition, then, um, the model's overall performance will be degraded as if you had used less compute to train it.

Q. And as a user of the model, how is that degradation you just described seen?

A. Um, so we have -- um, we have -- Anthropic has three different offerings of model sizes that vary by how much compute has gone into them, and if a model is made dumber, then it will be as if it was performing more similarly -- similarly to a model from a smaller class.

Q. And you mentioned there's a compute budget; is that correct?

A. Um, when I -- when I use the phrased "compute budget," I was saying the definition of the compute efficient frontier for -- for -- if you were given a fixed compute budget, there's a -- an optimal loss you

66

can achieve given -- like given a codebase or a training codebase.

Q. And so in your role, your objective is to find the optimal loss; is that correct?

A. I -- think it's more that we're trying to shift the compute efficient frontier to make training more efficient so you can achieve a better loss with the same amount of compute.

Q. I see. So, broadly speaking, in your work you want to reduce memorization; is that correct?

A. Um, that isn't something that I'm working on now, but, um, especially in 2023, I was looking into ways to, um, reduce memorization because memorization generally makes the models dumber and we want to reduce it.

Q. You said you worked on that in 2023?

A. Yeah.

Q. And did you succeed in your work?

A. Um, yeah, we -- we developed techniques to reduce duplication and, um, make it so that there was less text that the, um, models were memorizing so that they were, therefore, not being harmed as much or made dumber by memorizing that text.

Q. I'd like to come back to the techniques you developed to reduce duplication, but in the meantime, I

67

do want to ask does Anthropic decide what its AI models memorize?

MR. GOOCH: Objection. Calls for speculation.

THE WITNESS: Um, um, yeah. So -- so, in practice, um, based on the training data, um, if some piece of text is repeated many times and the model memorizes it, which is harmful to the model's performance, it makes the model dumber. Um, so by improving deduplication, um, we can reduce the number of things that are memorized.

BY MR. MATERA:

Q. Conversely, if you were to add repeated data, would you be able to increase the probability that a model memorizes that repeated data?

A. Yeah, you -- you wouldn't want to do --

MR. GOOCH: Sorry. Just give me --

THE WITNESS: Sorry.

MR. GOOCH: -- five seconds to object. No worries.

Same objection and also, um, asked and answered.

Go ahead.

THE WITNESS: Um, yeah, we wouldn't want to do that. Um, you could train a model on a heavily repeated dataset and have it memorize the text in that dataset,

68

but you wouldn't want to do that because it would -- it would make the model dumber.

BY MR. MATERA:

Q. Is there anything you would want the model to memorize?

A. Um --

MR. GOOCH: Same objections.

THE WITNESS: Um, um, yeah. So there's this, um, definition of memorization where the model has this pathologically high probability of generating some piece of text that it's been trained to memorize in this superficial way that, um, makes the model dumber, um, and then there's a more colloquial sense in which we want the model to -- um, there's a more colloquial sense or like another -- another definition where the model's simply capable of generating a sequence. Um, and, yeah, we want -- yeah. I think, um, it's good for the model to have memorized things like famous quotes that like you or I might memorize.

BY MR. MATERA:

Q. So are you saying there's a incidental memorization that comes with training and then a more deliberate memorization that you described as colloquial?

MR. GOOCH: Objection. Misstates prior

69

testimony and vague.
THE WITNESS: No, I wouldn't -- I wouldn't say that. I wouldn't say that there's intentional memorization, no.
BY MR. MATERA:
Q. Fair enough. I'm just trying to nail down these two types of --
A. Yeah.
Q. -- memorizations you mentioned. There's one -- if I'm understanding you correctly, you're saying there's one memorization that occurs through training and then there's another kind of memorization. Is that correct?
MR. GOOCH: Same objections.
THE WITNESS: Um, um, yeah, I'm saying that there's -- um, um, there's pathological memorization which can happen during training where the model overfits to a piece of text --
(Court reporter requested clarification.)
THE WITNESS: -- overfits to a piece of text by memorizing it in a superficial damaging way, um, and then separately there's a, um, capability where the model could, um, generate a -- a known piece of text like a famous quote or like a -- a famous excerpt from the Bible.

70

BY MR. MATERA:
Q. That capability you're describing where a model could generate a known piece of text like a famous quote or like a famous excerpt from the Bible, how does that capability work?
MR. GOOCH: Same objections.
THE WITNESS: Um, um, I mean, mechanically, the model, um, will -- um, is capable of, um, generating that piece of text, um...
BY MR. MATERA:
Q. Does it do that on its own or do you tell the model what to memorize specifically?
A. Um, so -- so, again, in general, we're trying to reduce memorization, um, because it's harmful for the model, um, because it makes the model dumber and it works at generalizing, um. Yeah. Sorry. I don't know of any -- any text that we've intentionally had the model memorize.
Q. I understand. I'm -- I'm asking more about the process. You gave some specific examples of "famous quotes and excerpts from the Bible," and I'm wondering how the model learns to memorize those specific examples that you listed?
A. Um, yeah, it --
MR. GOOCH: And same objection. Calls for

71

speculation and vague.
THE WITNESS: Um, um, yeah, so I don't -- I don't know that the model is -- is memorizing those things, um, yeah. But, in general, if a model is trained on a piece of text many times, then it will memorize the text, which in a typical case makes the model -- model dumber. It's a form of overfitting.
BY MR. MATERA:
Q. I understand that part. I'm talking about the specific example you gave of a model being capable of reciting, for example, a -- a bible excerpt. And I want to know how it does that and if it does that in a way that's different than the other memorization that you have discussed?
A. Um --
MR. GOOCH: Objection. Vague and asked and answered.
THE WITNESS: Yeah, I think -- I think I've answered this question.
BY MR. MATERA:
Q. I don't think you have. I'm asking specifically about this second form of memorization that you've been describing, and I just want to know how that works, who decides in a hypothetical situation what to specifically memorize?

72

MR. GOOCH: Objection. Vague. Compound. Calls for a hypothetical and speculation.
THE WITNESS: Could you ask a -- a narrower question?
BY MR. MATERA:
Q. If you wanted to train an AI model to memorize a specific bible excerpt, how would you do that?
MR. GOOCH: Objection. Calls for speculation and an expert opinion.
THE WITNESS: Um, again, in general, um, you could use, um, either pretraining or supervised learning or reenforcement learning.
(Court reporter requested clarification.)
THE WITNESS: Supervised learning.
But I don't -- I don't know of a way to do that which wouldn't, um, be damaging to the model. It's -- it's a hard problem because generally if you, um, just train on a piece of text, then it will result in this very superficial form of memorization, which is harmful to the model.
BY MR. MATERA:
Q. Earlier you said, quote, "I think it's good for the model to have memorized things like famous quotes that you or I might memorize." Do you recall that testimony?

Pages 77-80

77

THE WITNESS: Yeah, I haven't -- I haven't personally tested this, um, um...
BY MR. MATERA:
Q. Based on your knowledge of how Claude works, is that something you expect could happen?
A. Um --
MR. GOOCH: Calls for speculation.
THE WITNESS: -- um, um, yeah, I haven't -- I haven't personally tested or seen this, um, but, um, it probably -- there are things like famous quotes that the model can complete.
BY MR. MATERA:
Q. You said probably?
A. Yeah. I -- I haven't -- I haven't tested it so I don't know.
Q. You don't know for certain?
A. I haven't seen, yeah.
Q. Okay. We've been going about an hour. Would you like to take a break?
A. Yeah.
THE VIDEOGRAPHER: Okay. We are going off the record. The time is 11:12 a.m.
(Whereupon, a recess was taken.)
THE VIDEOGRAPHER: We are back on the record. The time is 11:26 a.m.

78

BY MR. MATERA:
Q. Miss Drain, do you recall before the break, we were talking about different types of memorization?
A. Yeah.
Q. You mentioned one type of memorization that you described as pathological and harmful to a model. Do you recall that?
A. Um, yeah. That's during when a model trains on a piece of text many times and memorizes it. Um, then it will overfit to that sub-distribution and become dumber, yeah.
Q. And you mentioned a second type of memorization that you referred to as colloquial memorization. Do you recall that testimony?
A. Um, um, yeah.
Q. Is there a technical term for that second type of memorization?
A. No. That wouldn't -- yeah, I think that was something that I was just saying in the moment. You know, I don't work with a technical definition of colloquial def- -- of memorization. This is just like -- um, I was more referring to things like you and I would have memorized.
Q. You weren't describing Claude's ability to memorize specific text in a manner that you and I would refer to colloquially as memorization?
MR. GOOCH: Objection. Asked and answered.
THE WITNESS: Um, I've -- I've already answered that.
BY MR. MATERA:
Q. I'm asking you to confirm.
MR. GOOCH: To the extent that's a -- are -- is -- "asking you to confirm," is that a question or are you re-asking? What are you doing?
Okay. Same objections.
THE WITNESS: Can you -- can you repeat the question?
BY MR. MATERA:
Q. You mentioned a second kind of memorization which you referred to as colloquial memorization, and I'm asking are you retracting that testimony now?
A. Um --
MR. GOOCH: Misstates the record and misstates the witness's answer.
You can answer.
THE WITNESS: I'm not making a retraction.
BY MR. MATERA:
Q. Okay. But to confirm, there is a capability of Claude to memorize specific excerpts of test -- text; is that correct.

79

80

MR. GOOCH: Objection. Misstates the witness's prior testimony.
THE WITNESS: Um, yeah. I mean, I would say that when you train a -- a model on a piece of text many times, then, um, it will get a very low loss on that piece of text, which is --
(Court reporter requested clarification.)
THE WITNESS: It will get a very low loss on that piece of text and be able to generate it verbatim, but that's a very superficial understanding. That's harmful for the model.
BY MR. MATERA:
Q. Do you recall your testimony earlier about the Gettysburg address?
A. Yeah.
Q. And you said that the Gettysburg address is something that Claude -- that would be good for Claude to memorize?
A. Um, yeah, that's an example, yeah, of a important piece of our culture, which was public domain, that like a reasonable adult would expected to have memorize and it would be reasonable for Claude to be able to generate that text.
Q. Is there a technical term for Claude's ability to memorize the Gettysburg address?

Pages 105-112

105

memorizing The Bee Movie." Have you seen that happen before?
A. Um, yeah, I haven't personally seen it happen. Um, yeah. Yeah.
Q. Are you aware of it happening before?
A. Um, um, I mean, I remember jokes about it from a few years ago, but I didn't -- yeah. Yeah, I haven't like actually seen it happen.
Q. Jokes from your coworkers?
A. Yeah.
Q. And by including The Bee Movie in this list of pillars of human culture, did you think your coworkers would get the joke when you presented it?
A. Um, yeah.
Q. That's because your coworkers have seen The Bee Movie replicated?
A. Um, no. Just because there's this -- well, I mean, first of all, obviously The Bee Movie's not in the same tier as like the Constitution. So even if they weren't familiar with -- with this, they could probably just like infer that it was a joke. It's also very long. Um, yeah. I don't -- I don't -- I don't know how many people are actually like familiar with the joke, um, but I think they could still infer it based on the context.

106

Q. You said The Bee Movie's very long?
A. Um, yeah. It's a whole movie.

107

[redacted]
Q. I got you. Okay. You listed famous poems here; correct?
A. Yeah.
Q. What's a famous poem that you consider to be a pillar of human culture?
A. Um, the -- there's like a Robert Frost about like two roads diverging in a -- in a forest. I mean, I don't -- I don't have it memorized, but, yeah, something like that.
Q. Do you know if that poem's in the public domain?
A. Um, yeah, like I'm not aware of whether the poem's on public domain. I don't know.
Q. You ended this list with "et cetera"; correct?
A. Yeah.
Q. Suggesting there are other pillars of human culture?
A. Yeah.
Q. What are some other pillars of human culture that you would want Claude to memorize verbatim?

108

A. Um, I mean, it's a pretty short list since, in general, verbatim memorization is bad and makes them all dumber, um, but, you know, one -- one example we had earlier was like famous bible passages that people have memorized. Um, but it would be a pretty short list.
Q. Are there any songs that you would consider to be pillars of human culture?
A. Um, maybe like some of The Beatles songs, yeah.
Q. Like which ones?
A. Um, like "Let It Be" or something. Um...
Q. Are you familiar with "American Pie" by Don McLean?
A. Yeah.
Q. Would you consider "American Pie" to be a pillar of cul- -- of human culture?
MR. GOOCH: Objection. Form.
THE WITNESS: Could you ask a different question or clarify the question?
BY MR. MATERA:
Q. You said you were familiar with the song "American Pie" by Don McLean; correct?
A. Yeah.
Q. Would you consider that song to be a pillar of human culture?
MR. GOOCH: Objection. Form.

109

THE WITNESS: Um, I -- I mean, I don't know if I would literally use the phrase "pillar of human culture." Yeah.
BY MR. MATERA:
Q. Is that a "no"?
A. I mean, it's hard -- I mean, it's not like the US Constitution.
Q. I'm just asking for your opinion on whether you think it's a pillar of human culture?
A. Um, yeah, maybe more like American culture. Um...
(Cross-talk.)
BY MR. MATERA:
Q. Are you familiar with the song "Blowing in the Wind" by Bob Dylan?
MR. GOOCH: You can answer this question, but then let's take a lunch.
MR. MATERA: I have a few more questions on this line I'd like to finish before we take a break, but I am happy to do a break soon after that. We are not yet at an hour.
MR. GOOCH: Miss Drain, would you like to take a lunch?
THE WITNESS: Yeah. Let's, you know, take a lunch.

110

MR. GOOCH: Okay.
THE VIDEOGRAPHER: Okay. We are going off the record. The time is 12:24 p.m.
(Whereupon, a lunch recess was taken.)
THE VIDEOGRAPHER: We are back on the record. The time is 1:03 p.m.
BY MR. MATERA:
Q. Miss Drain, we were just looking at Exhibit 138 before the lunch break. Do you recall that?
A. Yeah.
Q. And if you could please turn to page ending in 636. Do you see it says, "Excision hurts 'good' memorization evals"?
A. I see that.
Q. What does that mean?
A. Let me look at this. Yeah, so the -- the font's, um, very small so it's hard for me to read the labels for the graphs.
Q. Do you need the graphs to interpret the --
A. No. No. I -- I think I can say. Um, um, so this was -- um, Jamie had this, um, small eval set or test set which contained, um, some things that we thought it would be reasonable for the model to be able to generate that, um, were like touchstones of civilization like quotes from the bible that were in the public domain.
(Court reporter requested clarification.)
[redacted]
BY MR. MATERA:
Q. Okay. Who is Jamie?
A. Jamie's a coworker. He works on pretraining.
Q. Jamie's last name?
A. Jamie Sully.
Q. Okay. And you said Jamie conducted this eval; is that correct?
A. Yeah.
Q. And that eval contained certain touchstones of civilization quotes? Is that what you said?
A. Um, yeah. I said, yeah, touchstones of civilization like, um, especially from the public domain like important quotes from the bible.
Q. Did you help Jamie in this exercise?
A. Um, no, not really.
Q. Do you know what touchstones of civilization Jamie used in these evals?
A. Um, I mean, one of them says, "Performance of Bible Accuracy," but I don't -- I don't know any -- I don't know any of the details of the eval.
Q. You're looking at the third graph on this page --
A. Yeah.
Q. -- is that correct?
The top graph says, "Famous books." Do you see that?
A. Yeah.
Q. What does that mean?
A. Um --
MR. GOOCH: Objection. Calls for speculation.
THE WITNESS: Yeah. I -- I didn't make this -- this test set. Um, yeah. I don't -- I don't know what's in it.
BY MR. MATERA:
Q. You made this document though; correct?
A. Yeah. This was copying Jamie's graphs.
Q. So you just copied in the graphs?
A. Um, yes.
Q. Did you ask Jamie what any of these evals meant?
A. Um -- um, yeah, I don't -- I don't remember, but I don't think so.
Q. Okay. And the second graph here says, "Quotations." Do you see that?

Pages 117-124

117

[redacted]

Q. I want to skip ahead to the page ending in 641.

118

[redacted]

Q. And you say here, "Generic commoncrawl has all these garbage headers and footers." Do you see that?

A. Yeah.

Q. What did you mean by that?

A. Um, so by "garbage" here, I mean, um, if you train on it, it makes the model dumber, um, and you can look at the example here, there's this header, um, with all these like one- or two- or three-word lines, um, where pretty much every page on this website would have, um, a similar header and it would be very repetitive. And it's not like coherent sentences or anything so you wouldn't want to train a model on this.

Q. Okay. So can you help me out with headers here? You're saying there's a portion of a web page that has a header? Is -- is that right?

A. Um, yeah. On this web page, the header starts with "Jackson Booth Tattoos" and then probably goes down to, um, I guess the final occurrence of "Jackson Booth Tattoos." That's what -- that's what we call the header.

Q. I see. And you said the header repeats on every page from the website?

A. Um, yeah.

Q. Okay. And a footer?

A. Um, so the footer in this example would start with "Jackson's Latest Works" and go to the bottom, and it's similar where a lot of web pages have, um, all these links at the top and the bottom. And then they end up getting reproduced on -- on every page and this is highly repetitive text that you don't want to train on.

Q. Okay. And so do you remove these headers and footers during preprocessing?

A. Um, yeah. I don't know the details of how our current preprocessing works. Yeah. I mean, I think it would be a good intervention to remove them, but, yeah, I don't know what we're currently doing.

Q. But this exercise of removing headers and footers has been done in the past?

MR. GOOCH: Objection. Calls for speculation and vague.

THE WITNESS: Um, yeah. In -- in the past there were at least some models we -- we had similar heuristics to remove the headers and footers, but --

(Court reporter requested clarification.)

THE WITNESS: -- heuristics to remove headers and footers, but, um, um, yeah, I don't know what we're currently doing, um...

BY MR. MATERA:

Q. Have you ever removed headers and footers from a dataset that you're working on?

A. Um, um, yeah. Yeah, I don't -- I don't remember doing that for a dataset that I trained on.

(Court reporter requested clarification.)

BY MR. MATERA:

Q. But in this exhibit and on this page, are you suggesting that Anthropic could preprocess Common Crawl better if it removed these, quote, "garbage headers and footers"?

A. Um --

MR. GOOCH: Objection. Form.

THE WITNESS: Um, yeah. So I mean -- I mean, again, to -- to clarify, by "garbage," I mean if you

121

train on this highly repetitive text, which isn't even like complete sentences, then it generally makes the model dumber, um, and the other slide is saying that it could be beneficial for training to remove these repetitive headers and footers.
BY MR. MATERA:
Q. Can you pull back up in front of you Exhibit 135, please. It's the first one we looked at.
A. Okay.
Q. And -- oh, sorry. Maybe it's 136. If you could flip to the second page. Do you have a pen handy?
A. Nope.
Q. Would you like to use mine?
A. Okay.
Q. Can you please circle on Exhibit --
MR. GOOCH: I think you're on different exhibits.
MR. MATERA: No. We're on the right exhibit.
MR. GOOCH: You said 136?
BY MR. MATERA:
Q. We are on 136; correct?
A. Yeah.
MR. GOOCH: Okay.
MR. MATERA: I believe you're looking at 137.
MR. GOOCH: Great. I'll mark that on my little

122

sheet.
BY MR. MATERA:
Q. Miss Drain, Exhibit 136, with the pen can you please circle for me the header of this web page to the best of your ability.
MR. GOOCH: Objection. Form.
THE WITNESS: Um, yeah. So -- so previously I was referring to headers after html parsing, um, but what would presumably get turned into the header here would be the -- these tabs at the top.
BY MR. MATERA:
Q. Can you label that header for me, please?
A. Okay (marking).
Q. And is there a footer on this first page?
A. Um, yes. So, again, the -- the headers and footer are after html parsing.
Q. Understood.
A. But presumably, um, all these things at the bottom would get turned into the footer.
Q. And, for the record, you just circled the portion of Exhibit 136 that you understand to be the footer of the web page; is that correct?
MR. GOOCH: Objection. Misstates the witness's testimony.
THE WITNESS: Um, yeah. I circled the -- the

123

part that, um, after html parsing I imagine would turn into the footer.
BY MR. MATERA:
Q. Could you label that "footer" for me, please.
A. Okay (marking).
Q. Miss Drain, do you see in the portion that you just circled on the far right there appears a small symbol followed by the words "2025 Anthropic PBC"?
A. Yeah, I see that.
Q. Do you recognize that symbol?
A. I think that's the copyright symbol.
Q. Would that copyright symbol be removed --
MR. GOOCH: Objection. Vague --
BY MR. MATERA:
Q. -- as --
MR. GOOCH: -- and calls for speculation.
MR. MATERA: I wasn't finished answering -- asking my question.
MR. GOOCH: No worries. You had a good five-second pause so I thought you were. Go ahead and reask it.
MR. MATERA: I appreciate it.
BY MR. MATERA:
Q. Would this symbol be contained -- strike that.
Would you -- well, yeah.

124

Would this symbol be removed as part of the footer in the exercise we were just discussing?
MR. GOOCH: Objection. Calls for speculation and an expert opinion. Um, incomplete hypothetical.
THE WITNESS: Um, it would depend on what function you were actually using to remove the headers and footers, um...
BY MR. MATERA:
Q. Your answer's complete?
A. Um, yeah. It depends on what parcer you use and then what function you use to remove headers and footers.
(Court reporter requested clarification.)
THE WITNESS: What parcer you use and what function you use.
BY MR. MATERA:
Q. Okay. I appreciate your help with that. You can put it aside and I'll ask for my pen back if you don't mind. Thank you.
MR. GOOCH: Um, just so that we have it for the record, um, I'm going to move to strike the testimony from earlier today regarding Exhibit 126 now that the exhibit has been modified and is no longer representative of the exhibit that was testified about.
MR. MATERA: 126 is not --

Pages 149-160

149

[redacted]

(Court reporter requested clarification.)

[redacted]

Do you see that?

A. Yeah.

Q. What's a "GNU license"?

A. Um, a GNU license is a certain type of open source license that, um, is often placed at the top of a programming file.

Q. And that's the kind of license you described earlier as garbage; is that correct?

MR. GOOCH: Objection. Misstates prior testimony.

THE WITNESS: Um, okay. So -- so, again, to -- to contextualize, um, what the word "garbage" there meant was that there's, um, some piece of text which is, um, repeated many, many, many times, um, and if you were to train the model on it, the model would be overfit to that narrow sub-distribution of text and become dumber. Um, so that's -- that's what the word "garbage" meant there. Um, so, yeah, this is referring to, um, that type of license, which is often repeated at the top of every single source code file in a repo, which is very repetitive.

BY MR. MATERA:

Q. Understood. If you could please flip to page 10 of this document with a Bates stamp 759 at the end. It says, "Optimal dedupe parameters." Do you see that?

A. On -- on page 759?

Q. 507759.

A. Um, oh, yeah. Yeah. I see it, yeah.

Q. You see where it says, "Optimal dedupe parameters"?

A. Yeah.

151

[redacted]

MR. GOOCH: Objection. Calls for speculation.

[redacted]

152

[redacted]

Q. Okay. Let's skip ahead to page 47, which ends in 796.

A. Okay.

Q. You see where it says, "Memorization," in large font?

A. Yep.

Q. Below that it says, "Large LOL" -- sorry. Excuse me.

Below that it says, "Large LMs memorize a LOT," in all caps, "like a LOT," in all caps.

A. Yeah.

Q. What did you mean by that?

A. Um, this is looking at, um, a couple technical definitions of memorization. Sample 1 could be that you have a large span of tokens that all have pathologically low losses, um, and if you do a bad job deduplicating your dataset, then you can have the same piece of text repeated many times and the model could memorize it in its pathological sense. And, um, before we invested a

153

lot more in improving our deduplication algorithms, this was a bigger problem.

Q. I'm sorry. You said this is looking at a couple definitions of what?

A. Of memorization.

Q. What are those definitions?

A. Um, yeah. So that -- um, the -- the main two definitions here are, um, if you have, um, a single token with an extremely low loss or if you have --

(Court reporter requested clarification.)

THE WITNESS: An extremely low loss.

Or if you have a large span of tokens, um, with a pathologically low loss. Yeah, those are the two main definitions here.

BY MR. MATERA:

Q. Are there other definitions of memorization?

A. Um, yeah, we -- we talked earlier about the greedy sampling definition.

Q. Any others?

A. Um, um, um, and you can define -- you could set the -- the threshold for the loss to be based on, um, the typical loss for the model. You could renormalize it that way. I've used that definition before.

Q. Any others?

A. Um, um, yeah. It's pretty much just those.

154

Yeah. I mean, there are probably some others in the literature, but I can't -- I can't think of any.

Q. None other that you're aware of at this moment?

A. Yeah.

[redacted]

Q. "cc.p8" is a Common Crawl band; is that correct?

A. Yeah.

[redacted]

155

[redacted]

MR. GOOCH: Objection. Vague.

THE WITNESS: Um, I mean, I -- ideally you don't want the model to memorize a lot of text because it makes the model dumber.

BY MR. MATERA:

Q. I'm just referring to the language you wrote here.

[redacted]

Do you see that?

A. Yeah.

Q. Can you explain that in more layman's terms?

[redacted]

156

[redacted]

(Court reporter requested clarification.)

[redacted]

157

[redacted]

Q. Miss Drain, can you turn to page 69, which ends in 818.

A. Okay.

Q. You say, "What's getting memorized here?" Or, sorry. You just say, "What's getting memorized?"

Do you see that?

A. Yep.

Q. And then you see the paragraph that says, "The first cc.p8 stuff that gets removed is lists, lists, and lists"?

A. Yep.

Q. "(Mostly website headers and footers and lots of semi-expanded links to other articles. Also boring copyright stuff, exhortations to share and subscribe or whatever. Also webspam that has somehow perfused the internet eh."

Did I read that correctly?

158

A. Yep.

Q. What are you describing in this paragraph?

[redacted]

Q. You said, "Boring copyright stuff." What did you mean by that?

A. Um, um, it was mostly in contrast to content that would make the model smarter, like something from a book or like the actual code from a source code file. Um, yeah. I don't -- I don't have examples here so I'm not sure what the boring copyright stuff is that's being removed, but by definition, it's something that's repeated many, many times and is, um, reasonably long.

Q. And so if you remove the boring copyright stuff, the model will become smarter?

159

[redacted]

Q. So if you remove the boring copyright stuff, does that mean Claude will not train on it?

A. Okay. So, again, to -- to contextualize the -- the boring copyright stuff here, this is -- this is saying if you were to remove all of the high repetitive text, which includes things like repetitive website headers and links and ads and then like repetitive GitHub licenses, um, and if you don't -- if you don't train Claude on that highly repetitive text and, instead, train it on, um, higher quality text like books and stuff, then, yeah, the model will be smarter and it won't be training on the high repetitive text.

Q. And if Claude does not train on the boring copyright stuff, Claude will not memorize it? Is that correct?

160

MR. GOOCH: Objection. Asked and answered.

THE WITNESS: I feel these are all kind of like the same question. Um, yeah, I'm just going to give the -- the same answer.

BY MR. MATERA:

Q. What is that answer?

A. Um, so if you, um, remove the high repetitive text, like repeated website headers or -- or ads or other webspam, then the model will get smarter and one of the main mechanisms is that, um, the model doesn't, um, waste compute on training on this high repetitive text that it could then memorize.

Q. I just had a yes-or-no question there. Claude would not memorize that data?

MR. GOOCH: Counsel, you know, she's free to answer any question with the truthful answer. Asked and answered.

THE WITNESS: Um, I guess we looked at, um, like some statistics previously, um, where as a function of how many times the model -- how -- how many times the piece of text was repeated, um, there was a higher risk of the model memorizing it and overfitting to that subdistribution and becoming dumber. So if the text is repeated fewer times --

THE COURT REPORTER: I'm not getting you. I