EXHIBIT F

SONAL N. MEHTA (SBN 222086)
sonal.mehta@wilmerhale.com
ALLISON BINGXUE QUE (SBN 324044)
allison.que@wilmerhale.com
WILMER CUTLER PICKERING HALE AND DORR LLP
2600 El Camino Real, Suite 400
Palo Alto, CA 94306
Telephone: (650) 858-6000

JOSEPH R. WETZEL (SBN 238008)
joe.wetzel@lw.com
LATHAM & WATKINS LLP
505 Montgomery Street, Suite 2000
San Francisco, California 94111
Telephone: (415) 391-0600

Attorneys for Defendant
ANTHROPIC PBC
(Additional Counsel on Next Page)

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

SAN JOSE DIVISION

CONCORD MUSIC GROUP, INC., ET AL.,

Plaintiffs,

v.

ANTHROPIC PBC,

Defendant.

Case No. 5:24-cv-03811-EKL-SVK

DEFENDANT ANTHROPIC PBC'S AMENDED AND SUPPLEMENTAL RESPONSES AND OBJECTIONS TO PLAINTIFFS' FIRST SET OF INTERROGATORIES (NOS. 9, 11, 12)
Judge Eumi K. Lee
Magistrate Judge Susan van Keulen

HIGHLY CONFIDENTIAL – ATTORNEYS' EYES ONLY PORTIONS HIGHLIGHTED IN GREY

LOUIS W. TOMPROS (*Pro Hac Vice*)
louis.tompros@wilmerhale.com
STEPHANIE LIN (*Pro Hac Vice*)
stephanie.lin@wilmerhale.com
DISHA PATEL (*Pro Hac Vice*)
disha.patel@wilmerhale.com
WILMER CUTLER PICKERING HALE AND DORR LLP
60 State Street
Boston, MA 02109
Telephone: (617) 526-6000

ARI HOLTZBLATT (SBN 354361)
ari.holtzblatt@wilmerhale.com
ROBIN C. BURRELL (*Pro Hac Vice*)
robin.burrell@wilmerhale.com
WILMER CUTLER PICKERING HALE AND DORR LLP
2100 Pennsylvania Ave, NW
Washington, DC 20006
Telephone: (202) 663-6000

TAYLOR GOOCH (SBN 294282)
taylor.gooch@wilmerhale.com
KYLE EDWARDS HAUGH (SBN 323952)
kyle.haugh@wilmerhale.com
WILMER CUTLER PICKERING HALE AND DORR LLP
50 California St.
San Francisco, CA 94111
Telephone: (628) 235-1000

ANDREW M. GASS (SBN 259694)
andrew.gass@lw.com
BRITTANY N. LOVEJOY (SBN 286813)
brittany.lovejoy@lw.com
IVANA DUKANOVIC (SBN 312937)
ivana.dukanovic@lw.com
LATHAM & WATKINS LLP
505 Montgomery Street, Suite 2000
San Francisco, California 94111
Telephone: (415) 391-0600

SARANG V. DAMLE (*Pro Hac Vice*)
sy.damle@lw.com
LATHAM & WATKINS LLP
555 Eleventh Street NW, Suite 1000
Washington, DC 20004
Telephone: (202) 637-2200

ALLISON L. STILLMAN (*Pro Hac Vice*)
alli.stillman@lw.com
LATHAM & WATKINS LLP
1271 Avenue of the Americas
New York, New York 10020
Telephone: (212) 906-1747

Pursuant to Rules 26 and 33 of the Federal Rules of Civil Procedure 33 and the Local Rules of this Court, Defendant Anthropic PBC ("Anthropic") hereby supplements its responses to Plaintiffs' Interrogatory No. 9, Interrogatory No. 11, and Interrogatory 12. These supplemental responses are made subject to and without waiving the responses and objections in Anthropic's prior responses. Anthropic has not completed its own investigation and discovery. Accordingly, the following responses are based upon Anthropic's knowledge, information, and belief as of the date hereof. Anthropic reserves the right to make any further response if it appears that any omission or error has been made in connection with these responses or that future or more accurate information is available. These responses are made without prejudice to Anthropic's right to present in any proceeding in this litigation such additional evidence as may be hereafter discovered or evaluated.

PRELIMINARY STATEMENT AND GENERAL OBJECTIONS

1. These responses are made solely for the purpose of this action. Each response is subject to all objections as to relevance, materiality, propriety, and admissibility, and any and all other objections and grounds that would require the exclusion of the responses if such responses were to be introduced at any hearing or trial of this matter, all of which objections and grounds are reserved and may be interposed at the time of such hearing or trial.

2. These responses are made solely on Anthropic's behalf, and do not include information from other parties that is not already in Anthropic's possession, custody, or control.

3. Anthropic has attempted in good faith to respond to these Interrogatories, and its responses are based on information currently available to it at the time of this response. Anthropic has not completed investigating the facts related to this action, and further discovery, investigation, research, and analysis may supply additional information that could alter the disclosures reflected in these responses. In light of the continuing discovery and investigation, Anthropic reserves the right to further amend and/or supplement its responses, and expressly reserves the right to later use and present at trial any evidence, facts, and/or testimony developed or analyzed at a later date that

is responsive to the Requests. Notwithstanding the foregoing, Anthropic assumes no obligation to amend and/or supplement its responses beyond that which is required by law.

4. Anthropic objects to the Interrogatories that purport to impose on Anthropic discovery obligations that exceed those provided for in the Federal Rules of Civil Procedure or in the Local Rules for the United States District Court for the Northern District of California, orders entered in this case, agreements among the parties, or other applicable law.

5. Anthropic objects to the Interrogatories that seek information protected from disclosure by the attorney-client privilege, the work-product doctrine, the common interest doctrine, Fed. R. Civ. P. 26(b)(4)(A), and/or any other applicable privilege, immunity, or protection. The identification or production of any such document or information shall not constitute a waiver of any applicable objection or privilege. The disclosure of any information or production or identification of documents or communications that are privileged or otherwise immune from discovery shall be handled in accordance with the M.D. Tenn. Protective Order and the ESI Protocol that the Parties have agreed to follow pending entry of a Protective Order and ESI Protocol in the Northern District of California. *See* ECF No. 63-1 ("M.D. Tenn. ESI Protocol"). For the purposes of responding to these Interrogatories, Anthropic will interpret each Interrogatory, definition, and instruction as excluding information subject to privilege or immunity.

6. Anthropic objects to the Interrogatories that seek third-party confidential and/or proprietary information and/or information that is subject to or precluded by restrictions of confidentiality imposed by, or pursuant to, agreements between Anthropic and third parties. Anthropic will cooperate with Plaintiffs to identify what confidential third party information, if any, is relevant and reasonably necessary for investigation of the claims and defenses herein, subject to notice obligations.

7. Anthropic objects to the Interrogatories that seek information not currently in Anthropic's possession, custody, or control.

8. Anthropic objects to the Interrogatories to the extent they seek trade secrets or confidential or proprietary information of Anthropic. Anthropic will produce confidential documents and information as set forth herein, subject to the terms of the M.D. Tenn. Protective Order and any other relevant agreements, including the forthcoming protective order to be entered in the Northern District of California.

9. Anthropic objects to the Interrogatories to the extent they seek information that (i) does not currently exist; (ii) is unavailable; (iii) is only available to be compiled with the expenditure of extraordinary efforts; and/or (iv) may be impractical or impossible to create and/or reasonably capture, resulting in a significant, punitive expense to Anthropic.

10. Anthropic objects to the Interrogatories that prematurely seek to require Anthropic to disclose individuals or information inconsistent with the Federal Rules of Civil Procedure, including Federal Rule of Civil Procedure 26(a)(2) regarding expert testimony, as well as the rules and orders of this Court.

11. Anthropic objects to the Interrogatories to the extent to which any terms used therein are vague, ambiguous, or overbroad.

12. Anthropic objects to the Interrogatories to the extent they are cumulative and duplicative of other discovery requests propounded by Plaintiffs.

13. Anthropic objects to the Interrogatories to the extent they call for legal conclusions. Any response by Anthropic shall not be construed as providing a legal conclusion regarding the meaning or application of any terms or phrases used in Plaintiffs' Interrogatories, definitions, or instructions.

14. Anthropic objects to each Interrogatory to the extent it contains multiple subparts or the definition or instruction renders an Interrogatory compound. Such Interrogatories constitute separate interrogatories that count towards the limit on interrogatories that Plaintiffs are permitted to propound. Anthropic objects to each and every Interrogatory that would exceed such limit.

15. Anthropic objects to each Interrogatory to the extent it seeks to require Anthropic to do more than conduct an examination of those files or sources that reasonably may be expected

to yield responsive information, or an inquiry of those persons who may be reasonably expected to possess responsive information.

16. Anthropic objects to the Interrogatories to the extent they purport to place an obligation on Anthropic to obtain or provide information that is publicly available or otherwise as readily available to Plaintiffs as it is to Anthropic.

17. Anthropic's response to part or all of any Interrogatory is not intended to be, and shall not be construed to be, a waiver of any part of any objection to any Interrogatory.

18. Anthropic's willingness to respond to any Interrogatory does not constitute an admission that Anthropic agrees with any characterization, definition, or assumption contained in the Interrogatory or an assumption or an acknowledgement by Anthropic that the Interrogatory is proper, that the information sought is within the proper bounds of discovery, or that demands for similar information will be treated in similar fashion.

OBJECTIONS TO DEFINITIONS

1. Anthropic objects to the Definitions of "Defendant," "Anthropic," "You," or "Your" on the basis that they are overbroad. The inclusion of "any other person or entity . . . previously acting or purporting to act" on behalf of Anthropic would unreasonably expand Anthropic's burden of responding to these Interrogatories. Anthropic will interpret "Defendant," "Anthropic," "You," or "Your" to mean the Anthropic PBC, as identified in the complaint.

2. Anthropic objects to the definition of "Claude" as overbroad, in that the definition includes general purpose generative-AI large language models that are not relevant to Plaintiffs' claims in this action, including models that were never made commercially available or used to develop commercial Claude, were not trained on datasets that potentially include Publishers' Works, and products hosted by third parties. Anthropic will interpret "Claude" as ordered by the Court in the November 26, 2024 Order (ECF No. 273), including with respect to negotiating an end date for relevant product development. *Id.* at n.1.

3. Anthropic objects to the Definition of "AI Company" as vague and ambiguous to the extent that "AI company" is used interchangeably with "companies that currently provide generative AI services."

4. Anthropic objects to the Definition of "Training Corpus" as vague and ambiguous to the extent it refers to any "collection of materials used to train a generative AI model."

5. Anthropic objects to the Definition of "Employee" as overbroad to the extent that it includes "subsidiar[ies]" and "servant[s]."

6. Anthropic objects to the definitions of "communication," "document," and "identify" on the grounds that they are overbroad to the extent they include documents or information beyond the requirements of the Federal Rules of Civil Procedure; documents or information that are not in Anthropic's possession, custody, or control; materials that are not proportionally relevant to the needs of this case; and materials from sources that are not reasonably accessible. Anthropic's response should not be construed as acquiescence to Plaintiffs' Definition of these terms or as a limitation of the scope of Anthropic's defenses or legal positions.

7. Anthropic objects to Plaintiffs' failure to define terms that are vague, ambiguous, and/or subject to varying interpretations. Where possible, Anthropic will make a reasonable effort to respond to each Interrogatory to the extent not otherwise objectionable, based upon its understanding and interpretation of the Interrogatory, including by interpreting each undefined term or phrase in accordance with its ordinary meaning. If Plaintiffs subsequently assert any interpretation of any Interrogatory that differs from Anthropic's interpretation, Anthropic reserves the right to supplement or amend its objections and responses.

OBJECTIONS TO INSTRUCTIONS

1. Anthropic objects to Instruction No. 8, and to each Interrogatory, as providing an overbroad timeframe.

2. Anthropic objects to Instruction No. 5, and each Interrogatory, to the extent it purports to impose procedures or requirements for identifying information over which Anthropic claims privilege.

3. Anthropic reserves the right to respond to the Interrogatories pursuant to Federal Rule of Civil Procedure 33(d), where appropriate.

* * *

Each of the foregoing Objections is incorporated in full to each response below as if restated fully therein.

RESPONSES TO INTERROGATORIES

INTERROGATORY NO. 9:

Describe Anthropic's methods and process for cleaning, filtering, editing, removing data, or otherwise modifying data or datasets used to train Claude, including but not limited to data or datasets used in pretraining and finetuning.

DECEMBER 20, 2024 AMENDED SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 9 (HIGHLY CONFIDENTIAL — ATTORNEYS' EYES ONLY):

Anthropic incorporates its General Objections, Objections to Definitions, and Objections to Instructions as if fully asserted herein. Anthropic objects to this Interrogatory as vague and ambiguous insofar as it includes the undefined terms "cleaning," "filtering," "editing" and "modifying." Anthropic further objects to this Interrogatory as overbroad, unduly burdensome, and disproportionate to the needs of this case, because it asks for descriptions of technically complex methods and processes that are highly proprietary and sensitive and that are neither relevant to either party's claims or defenses nor proportional to the needs of this case. Anthropic further objects to this Request as overbroad, unduly burdensome, and not proportional to the needs of this case to the extent it seeks information relating to AI models other than Claude, as defined by the Court in the November 26, 2024 Order, or future versions of Claude released beyond the product development end date to be negotiated by the Parties. Anthropic further objects to this Interrogatory because it is compound and consists of multiple subparts, each of which constitutes a separate interrogatory that will count toward the limit of interrogatories that Plaintiffs are permitted to propound. Anthropic objects to this Interrogatory to the extent it seeks information protected by the attorney-client privilege, work product doctrine, or any other applicable privilege

or immunity. Anthropic also objects to this Interrogatory to the extent that it is cumulative and/or duplicative of other discovery requests by calling for materials or information that Plaintiffs have sought elsewhere, for example in Request for Production No. 27.

Subject to the foregoing General and Specific Objections, Anthropic provides the following overview of Anthropic's process for filtering, cleaning, and otherwise modifying pre-training data: At a high level of generality, Anthropic's process of training Claude is described as follows. Anthropic requires a massive quantity and diversity of datasets on which to train Claude. These datasets include data derived from many diverse sources, such as the Internet, including but not limited to websites and blogs, books, research papers, source code, patents, news media, and advertising. Anthropic must then engage in a series of activities to transform the raw data it collects into a format that can be used for training. After Anthropic collects the data, it cleans the data and processes the data for training.

Data Cleaning. Anthropic needs to "clean" the data it collects to reduce the amount of noise. "Noise" includes junk data and low quality data that would detract from the goal of teaching Claude to recognize true patterns in human language. Junk data might include boilerplate HTML text (e.g., "Lorum Ipsum") and web spam (e.g., duplicitous and inaccurate keyword stuffing). Data at the lower end of the quality spectrum might include data that is highly repetitive (i.e., the same or similar words over and over again) and websites that are just headers and footers without any meaningful content.

Anthropic filters out junk and low quality data using classifiers. Classifiers are supplemental tools, trained on a pre-categorized dataset so that they can predict the categories for new, unseen data, that are used to categorize, label, or score input data. The classifiers used at this stage are trained to detect when data looks like junk or low quality data, and to filter that data out. Filtering using classifiers is time and resource intensive, so for certain categories, Anthropic only filters on samples of the data, not the entire corpora of data that it collects. As a result, the training corpus still contains certain low quality data and remains relatively "noisy." This means that not every piece of data on which Claude is trained is high quality.

Deduplication. Duplicated or skewed datasets can ultimately result in less intelligent outputs. Anthropic engages in deduplication efforts to prevent data points from inadvertently being weighted too heavily by the algorithm during the training process because of their frequency. A data point that is overweighted might result in an LLM becoming biased towards that piece of data and overfitting its outputs to conform with the patterns it learned from that data, a phenomenon that degrades model performance.

Additionally, if a company reaches out to remove certain discrete corpora from Anthropic's training datasets, Anthropic will consider that request and may work with the company to exclude that data from new training runs that have not commenced. Requests are more likely to be feasible if the company points to specific URLs for exclusion.

Data Processing. Under this category, Anthropic normalizes and tokenizes the data. Normalization, for example, may include removing extra white space and other efforts to make the data easier for the model to understand.

A key part of data processing is tokenization. Tokenization refers to the process of converting human-readable text into binary computer code and then breaking down this code into smaller numerical units called tokens that an LLM like Claude can understand. Tokens are subsets of words in numerical form. On average, 1.4 tokens are required to represent an English word. Once the data is tokenized, sequences of tokens are scrambled and may be stored in the training corpus in a different order than their human-readable text equivalents appear in the original text of the material.

Only at this point will the model parse the data, analyzing each token within the context of the overall text to glean language principles and the relationship between the words and phrases in the corpus. The language patterns that the model discovers and records in the form of statistical information are then organized in a dataset of values representing weights and biases, or parameters. When combined with software, these parameters—not the training datasets themselves—allow Claude to generate textual outputs in response to user prompts.

November 12, 2025 AMENDED SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 9:

Subject to the foregoing General and Specific Objections, Anthropic incorporates herein its prior responses to Interrogatory No. 9. In addition, Anthropic further responds as follows:

Certain types of low-quality data that may be present on websites that may need to be filtered out can include signatures, user information, or page navigation. *See*, *e.g.*, Anthropic_0000218687. In addition, the deduplication process prevents repetition of low-quality data and memorization or regurgitation. *See*, *e.g.*, Anthropic_0000000759; Anthropic_0000347404; Anthropic_0000241415; Anthropic_0000347360; Anthropic_0000343002. Anthropic further states that information responsive to this interrogatory may be found in the deposition testimony or exhibits of the following witnesses: Ben Mann, Daniela Amodei, Dawn Drain, Deep Ganguli, Devin Whittle, Jane Leibrock, Jared Kaplan, Pranay Sangani, Tom Brown, Tom Turvey, and Vinay Rao.

INTERROGATORY NO. 11:

Describe any and all efforts to eliminate copyrighted material from data or datasets used to train Claude, including without limitation data or datasets used in pretraining and finetuning.

DECEMBER 20, 2024 AMENDED RESPONSE TO INTERROGATORY NO. 11:

Anthropic incorporates its General Objections, Objections to Definitions, and Objections to Instructions as if fully asserted herein. Anthropic objects to this Interrogatory as vague and ambiguous insofar as it includes the undefined term "copyrighted material," which does not offer appropriate parameters for providing a response. Anthropic further objects to this Request as overbroad, unduly burdensome, and not proportional to the needs of this case to the extent it seeks information relating to AI models other than Claude, as defined by the Court in the November 26, 2024 Order, or future versions of Claude released beyond the product development end date to be negotiated by the Parties. Anthropic also objects to this Interrogatory because it seeks information that is not relevant to either party's claims or defenses in this case. Anthropic objects to the Interrogatory to the extent that it seeks legal conclusions about the copyright eligibility or status

of the enormous quantity of materials used to train Claude. Anthropic objects to this Interrogatory to the extent that it is cumulative and/or duplicative of other discovery requests by calling for materials or information that Plaintiffs have sought elsewhere, for example in Request for Production No. 40. Anthropic further objects to this Interrogatory to the extent it seeks information protected by the attorney-client privilege, work product doctrine, or any other applicable privilege or immunity.

Subject to the foregoing General and Specific Objections, Anthropic responds as follows: Anthropic engages in efforts to ensure that certain content is not contained with its training corpus. First, Anthropic respects industry practices with respect to robots.txt instructions and other signals that website operators use to indicate whether they permit crawling of the content on their sites. Second, Anthropic operates its crawling system transparently, which means website operators can easily identify Anthropic visits and signal their preferences to Anthropic directly, including whether they want to block Anthropic's access to content on their sites. Third, as a general matter, Anthropic implements deduplication efforts in connection with its training data to prevent inadvertent "memorization" of specific texts. The process is designed to minimize regurgitation, thereby mitigating unwanted duplication of material in the training corpus.

Anthropic's investigation and discovery is ongoing, and Anthropic will supplement its response to this Interrogatory as appropriate.

NOVEMBER 12, 2025 AMENDED SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 11:

Subject to the foregoing General and Specific Objections, Anthropic incorporates herein its prior responses to Interrogatory No. 11. In addition, Anthropic responds as follows:

Anthropic is transparent about the ways in which Anthropic accesses, collects, and uses web content and details how website operators can signal their crawling preferences to Anthropic directly. *See*, *e.g.*, Anthropic_0000344022. Anthropic also has a reasonable and easily accessible take down form that anyone can submit if they believe someone is infringing their intellectual property rights. *See* Anthropic_0000000586. The terms of service also state that users should

report any potential intellectual property infringement by a third-party. *See* Anthropic_0000000550. Anthropic further states that information responsive to this interrogatory may be found in the deposition testimony or exhibits of the following witnesses: Ben Mann, Daniela Amodei, Dawn Drain, Deep Ganguli, Devin Whittle, Jane Leibrock, Jared Kaplan, Pranay Sangani, Tom Brown, Tom Turvey, and Vinay Rao.

NOVEMBER 12, 2025 AMENDED SUPPLEMENTAL RESPONSE TO INTERROGATORY NO. 12 (HIGHLY CONFIDENTIAL - ATTORNEYS' EYES ONLY):

Subject to and without waiving the foregoing General and Specific Objections, Anthropic amends and supplements its response to Interrogatory No. 12 as follows:

Anthropic intends its commercially released Claude models to serve humans in a helpful, honest, and harmless way. Anthropic has developed policies to prohibit users from misusing Claude in a way that is inconsistent with its mission. With respect to intellectual property specifically, Anthropic's Consumer Terms of Service expressly prohibit users from "access[ing] or us[ing]" Anthropic's services to "infringe, misappropriate, or violate intellectual property or other legal rights." Anthropic_0000437178. Both Anthropic's Consumer and Commercial Terms of Service incorporate Anthropic's Usage Policy, also referred to as an "Acceptable Use Policy," Anthropic_0000437176; Anthropic_0000437192, which expressly states that users are not allowed to "[v]iolate [a]pplicable [l]aws or [e]ngage in [i]llegal [a]ctivity," including "using [Anthropic's] products or services to: [i]nfringe, misappropriate, or violate the intellectual property rights of a third party," Anthropic_0000437165; Anthropic_0000440731. Anthropic's Safeguards Team enforces the Usage Policy through monitoring and detection, and may "throttle, suspend, or terminate" access to Claude for violations of the policy. Anthropic_0000437165.

In addition to implementing these user-facing policies, Anthropic restricts its commercially released Claude models from inadvertently allowing users to generate outputs that infringe copyrights to further Anthropic's goal of providing a harmless, general purpose, generative-AI model. Prior to releasing Claude to commercial customers in March 2023, Anthropic implemented certain technological systems intended to prevent the generation of outputs that may infringe

copyrights. Since then, Anthropic has augmented and iterated upon those systems and continues to make improvements.

Constitutional Copyright Principle. Commercially released Claude models incorporate "Constitutional AI" as a part of their finetuning process. "Constitutional AI" is an approach that incorporates ethical principles and behavioral guidelines directly into the model's training objectives. The goal of Constitutional AI is to create AI systems that reliably follow intended behaviors and values, even in new situations they have not encountered before. Constitutional AI means that ethical behavior is fundamental to how Claude thinks and operates. One of the Constitutional AI principles that Claude is finetuned on is to avoid substantially reproducing copyrighted content. The principle asks Claude to "understand and respect intellectual property rights" and to avoid "reproduc[ing] portions of copyrighted text in a way that could infringe on copyright." This means that Claude can leverage its own intelligence to prevent users from generating offending outputs and police itself before turning to any auxiliary guardrail systems that sit on top of the model.[I] When Claude detects that a user is requesting an output that would violate the principle against substantially reproducing copyrighted content, Claude is finetuned to avoid responding to the user's prompt and to generate a refusal instead. This method is particularly effective at preventing users from writing creative prompts to get around more traditional auxiliary systems. Anthropic first added the principle against substantially reproducing copyrighted content to its Constitution on or around [redacted]. Constitution.r1 artifact, Commit [redacted].[II]

Prompt Shield. In addition to Claude's constitutional AI principle, which acts as the first line of defense, Anthropic uses two auxiliary systems to further detect and prevent the generation of copyrighted content. The first auxiliary system that Anthropic uses is an intervention tool called "Prompt Shield." Prompt Shield uses a specially trained classifier model to detect when a user attempts to generate an output that would violate Anthropic's Terms of Use or Acceptable Use

[I] An auxiliary system is one that is separate from the model's training or finetuning itself. Two auxiliary guardrail systems, Prompt Shield and Content Compare System, are discussed in more detail below.

[II] Citations to "artifacts" refer to portions of the source code, which Anthropic has produced in the inspection environment.

Policy, including by eliciting outputs that may infringe copyrights. When Prompt Shield detects a prompt that may seek to elicit copyrighted content, Prompt Shield will append a request, invisible to the end user, at the end of the prompt asking Claude to be particularly considerate and careful not to substantially reproduce copyrighted content when responding. If Claude determines that an output would violate the appended request not to substantially reproduce copyrighted content, it will not provide the requested output. Anthropic first implemented Prompt Shield for Claude.ai and its first-party API product on or around [redacted]. Promptshield.r2 artifact, Commit [redacted]. Anthropic enabled Prompt Shield on its third-party API product through Bedrock on or around [redacted]. Promptshield.r2 artifact, Commit [redacted]. On or around [redacted] product through Vertex with Prompt Shield enabled. Promptshield.r2 artifact, Commit [redacted]

Content Compare System. The second auxiliary system Anthropic uses to further detect and prevent the generation of outputs including substantial portions of copyrighted content is a recitation checker that sits on top of the model called "Content Compare System" or "CCS." A recitation checker is a tool that checks new content against a separate dataset, external to the model, to determine whether the new content is the same or similar to the external dataset's contents. CCS's external dataset is called the hash corpus. At the time that Claude generates an output, CCS checks the output against the hash corpus before the output is sent to the end user. When CCS performs its check, the texts are "normalized"—i.e., stripped of all spacing, casing, punctuation, and accents—to ensure that CCS catches matches in semantic content, not form. If CCS detects that the output would reproduce an amount of a work in the hash corpus above a certain threshold, CCS will block the output so that it does not reach the user. When the first CCS system was implemented on or around [redacted] that matched against the contents of its corpus. CCS artifact, Commit [redacted] with CCS enabled. CCS.r1b artifact, Commit [redacted] Anthropic enabled CCS on its third-party API product through Bedrock on or around [redacted] CCS.r1b artifact, Commit

[redacted] Anthropic enabled CCS on its third-party API product through Vertex on or around [redacted] [redacted] CCS.r1b artifact, Commit [redacted].

Separately, on or around [redacted] [redacted] CCS's hash corpus. CCS artifact, Commit [redacted]. In the same Commit, Anthropic added [redacted] that were added to CCS's hash corpus. *Id.* Anthropic [redacted] [redacted]. *Id.* A "context window" refers to the span of text that Claude considers when determining whether there has been a match with text in the hash corpus. Then, on or around [redacted] [redacted] to CCS's hash corpus. CCS artifact, Commit [redacted]. The guardrails that Anthropic has implemented prevent Claude from including copyrighted material, like song lyrics, in outputs. *See, e.g.*, Doc. ID 00454102 (from sample produced pursuant to Dkt. 377) ("Human: what are the lyrics of: Mississippi Goddamn and Four Women Assistant: Unfortunately I am unable to provide the full lyrics to the songs 'Mississippi Goddam' and 'Four Women' without potentially violating copyright law. However, I can give a brief overview of each song…."). Anthropic anticipates that it will make further refinements and improvements to CCS and its other guardrails as it continues its mission of building helpful, honest, and harmless AI models.

In addition, Anthropic responds as follows, pursuant to Fed. R. Civ. P. 33(d): Anthropic has produced the following document that explains how Claude is trained with Constitutional AI: Claude's Constitution (Anthropic_0000000485). Anthropic has additionally produced the following documents that explain Anthropic's implementation and use of Prompt Shield: Prompt Shield explainer (go/prompt-shield) (Anthropic_0000012694); [redacted] ML Transition Plan (Anthropic_0000241411). Finally, Anthropic has produced the following documents that explain the evolution of Anthropic's Content Compare System over time: CCS Changelog (Anthropic_0000440820); Execution: Copyfairy as a sampling Stage (Anthropic_0000232830); and Copyfairy Content Enums (Anthropic_0000347950). Anthropic further states that information

responsive to this interrogatory may be found in the deposition testimony or exhibits of the following witnesses: Ben Mann, Daniela Amodei, Dawn Drain, Deep Ganguli, Devin Whittle, Jane Leibrock, Jared Kaplan, Pranay Sangani, Tom Brown, Tom Turvey, and Vinay Rao.

Date: November 12, 2025

/s/ *Taylor Gooch*

JOSEPH R. WETZEL (SBN 238008)
joe.wetzel@lw.com
ANDREW M. GASS (SBN 259694)
andrew.gass@lw.com
BRITTANY N. LOVEJOY (SBN 286813)
brittany.lovejoy@lw.com
IVANA DUKANOVIC (SBN 312937)
ivana.dukanovic@lw.com
LATHAM & WATKINS LLP
505 Montgomery Street, Suite 2000
San Francisco, California 94111
Telephone: (415) 391-0600

SARANG V. DAMLE (*Pro Hac Vice*)
sy.damle@lw.com
LATHAM & WATKINS LLP
555 Eleventh Street NW, Suite 1000
Washington, DC 20004
Telephone: (202) 637-2200

ALLISON L. STILLMAN (*Pro Hac Vice*)
alli.stillman@lw.com
LATHAM & WATKINS LLP
1271 Avenue of the Americas
New York, New York 10020
Telephone: (212) 906-1747

SONAL N. MEHTA (SBN 222086)
sonal.mehta@wilmerhale.com
ALLISON BINGXUE QUE (SBN 324044)
allison.que@wilmerhale.com
WILMER CUTLER PICKERING HALE AND DORR LLP
2600 El Camino Real, Suite 400
Palo Alto, CA 94306
Telephone: (650) 858-6000

LOUIS W. TOMPROS (*Pro Hac Vice*)
louis.tompros@wilmerhale.com
STEPHANIE LIN (*Pro Hac Vice*)
stephanie.lin@wilmerhale.com
DISHA PATEL (*Pro Hac Vice*)
disha.patel@wilmerhale.com
WILMER CUTLER PICKERING HALE AND DORR LLP
60 State Street
Boston, MA 02109
Telephone: (617) 526-6220

ARI HOLTZBLATT (SBN 354361)
ari.holtzblatt@wilmerhale.com
ROBIN C. BURRELL (*Pro Hac Vice*)
robin.burrell@wilmerhale.com
WILMER CUTLER PICKERING HALE AND DORR LLP
2100 Pennsylvania Ave, NW
Washington, DC 20006
Telephone: (202) 663-6000

TAYLOR GOOCH (SBN 294282)
taylor.gooch@wilmerhale.com
KYLE EDWARDS HAUGH (SBN 323952)
kyle.haugh@wilmerhale.com
WILMER CUTLER PICKERING HALE AND DORR LLP
50 California St.
San Francisco, CA 94111

Telephone: (628) 235-1000

Attorneys for Defendant
ANTHROPIC PBC

CERTIFICATE OF SERVICE

I hereby certify that on November 12, 2025, I caused to be served a true and correct copy of the foregoing document to the following counsel of record:

OPPENHEIM + ZEBRAK, LLP
Matthew J. Oppenheim (*Pro Hac Vice*)
Nicholas C. Hailey (*Pro Hac Vice*)
Audrey Adu-Appiah (*Pro Hac Vice*)
Corey Miller (*Pro Hac Vice*)
Jeffrey M. Gould (*Pro Hac Vice*)
4530 Wisconsin Avenue, NW, 5th Floor
Washington, DC 20016
Telephone: (202) 480-2999
matt@oandzlaw.com
nick@oandzlaw.com
aadu-appiah@oandzlaw.com
corey@oandzlaw.com
jeff@oandzlaw.com

COBLENTZ PATCH DUFFY & BASS LLP
Jeffrey G. Knowles
Christopher J. Weiner
Bina Patel
Thomas A. Harvey
One Montgomery Street, Suite 3000
San Francisco, CA 94104
Telephone: (415) 772-5795
Ef-jgk@cpdb.com
cwiener@coblentzlaw.com
bpatel@coblentzlaw.com
Ef-Tah@cpdb.com

Jennifer Pariser (*Pro Hac Vice*)
Andrew Guerra (*Pro Hac Vice*)
Timothy Chung (*Pro Hac Vice*)
Michelle Gomez-Reichman (*Pro Hac Vice*)
Bret Matera (*Pro Hac Vice*)
Alexander Kaplan (*Pro Hac Vice*)
461 5th Avenue, 19th Floor
New York, NY 10017
Telephone: (212) 951-1156
jpariser@oandzlaw.com
andrew@oandzlaw.com
tchung@oandzlaw.com
mgomez-reichman@oandzlaw.com
bmatera@oandzlaw.com
alex@oandzlaw.com

COWN, LEIBOWITZ & LATMAN, P.C.
Richard S. Mandel (*Pro Hac Vice*)
Jonathan Z. King (*Pro Hac Vice*)
Richard Dannay (*Pro Hac Vice*)
114 West 47th Street
New York, NY 10036-1525
Telephone: (212) 790-9200
rsm@cll.com
jzk@cll.com
rxd@cll.com

Dated: November 12, 2025

By: */s/ Taylor Gooch*
Taylor Gooch