1 Joseph R. Saveri (State Bar No. 130064)
   Steven N. Williams (State Bar No. 175489)
2 Cadio Zirpoli (State Bar No. 179108)
   Elissa A. Buchanan (State Bar No. 249996)
3 Travis Manfredi (State Bar No. 281779)
4 **JOSEPH SAVERI LAW FIRM, LLP**
   601 California Street, Suite 1000
5 San Francisco, California 94108
   Telephone: (415) 500-6800
6 Facsimile: (415) 395-9940
   Email: jsaveri@saverilawfirm.com
7   swilliams@saverilawfirm.com
                  czirpoli@saverilawfirm.com
8   eabuchanan@saverilawfirm.com
                  tmanfredi@saverilawfirm.com
9
10 Matthew Butterick (State Bar No. 250953)
11 1920 Hillhurst Avenue, #406
   Los Angeles, CA 90027
12 Telephone: (323) 968-2632
   Facsimile: (415) 395-9940
13 Email: mb@butticklaw.com

14 *Counsel for Individual and Representative*
   Plaintiffs and the Proposed Class
15
16   **UNITED STATES DISTRICT COURT**
17   **NORTHERN DISTRICT OF CALIFORNIA**
                   SAN FRANCISCO DIVISION
18
   J. DOE 3 and J. DOE 4, individually and on behalf of   Case No.
19 all others similarly situated,
                                                             COMPLAINT
20   Individual and Representative Plaintiffs,
                                                             CLASS ACTION
21   v.

22 GITHUB, INC., a Delaware corporation;   **DEMAND FOR JURY TRIAL**
   MICROSOFT CORPORATION, a Washington
23 corporation; OPENAI, INC., a Delaware nonprofit
   corporation; OPENAI, L.P., a Delaware limited
24 partnership; OPENAI GP, L.L.C., a Delaware limited
   liability company; OPENAI STARTUP FUND GP I,
25 L.L.C., a Delaware limited liability company;
   OPENAI STARTUP FUND I, L.P., a Delaware
26 limited partnership; OPENAI STARTUP FUND
   MANAGEMENT, LLC, a Delaware limited liability
27 company,

28   Defendants.

───

TABLE OF CONTENTS

CLASS ACTION COMPLAINT

Plaintiffs J. Doe 3 and J. Doe 4 ("Plaintiffs"), on behalf of themselves and all others similarly situated, bring this Class Action Complaint (the "Complaint") against Defendants GitHub, Inc.; Microsoft Corporation; OpenAI, Inc.; OpenAI, L.P.; OpenAI GP, L.L.C.; OpenAI Startup Fund GP I, L.L.C.; OpenAI Startup Fund I, L.P.; and OpenAI Startup Fund Management, LLC[1] for violation of the Digital Millennium Copyright Act, 17 U.S.C. §§ 1201–1205 (the "DMCA"); violation of the Lanham Act, 15 U.S.C. § 1125; violation of Unfair Competition law, *Cal. Bus. & Prof. Code* §§ 17200, *et seq.*; violation of the California Consumer Privacy Act, *Cal. Civ. Code* § 1798.150 (the "CCPA"); and Breach of Contract regarding the Suggested Licenses, GitHub's Privacy Statement, and GitHub's Terms of Service, *Cal. Bus. & Prof. Code* §§ 22575–22579, *Cal. Civ. Code* § 1798.150. Plaintiffs and the Class also bring this Complaint against Defendants for their Tortious Interference in Plaintiffs' Contractual Relationships; Fraud, and Negligence regarding handling of sensitive data.

I.   OVERVIEW: A BRAVE NEW WORLD OF SOFTWARE PIRACY

1.   Plaintiffs and the Class are owners of copyright interests in materials made available publicly on GitHub that are subject to various licenses containing conditions for use of those works (the "Licensed Materials."). All the licenses at issue here (the "Licenses") contain certain common terms (the "License Terms").

2.   "Artificial Intelligence" is referred to herein as "AI." AI is defined for the purposes of this Complaint as a computer program that algorithmically simulates human reasoning or inference, often using statistical methods. Machine Learning ("ML") is a subset of AI in which the behavior of the program is derived from studying a corpus of material called training data.

---

[1] GitHub, Inc. is referred to as "GitHub." Microsoft Corporation is referred to as "Microsoft." OpenAI, Inc.; OpenAI, L.P.; OpenAI GP, L.L.C.; OpenAI Startup Fund GP I, L.L.C.; OpenAI Startup Fund I, L.P.; and OpenAI Startup Fund Management, LLC are referred to collectively herein as "OpenAI." Collectively, GitHub, Inc., Microsoft Corporation, OpenAI, Inc.; OpenAI, L.P.; OpenAI GP, L.L.C.; OpenAI Startup Fund GP I, L.L.C.; OpenAI Startup Fund I, L.P.; and OpenAI Startup Fund Management, LLC are referred to herein as "Defendants."

3.   GitHub is a company founded in 2008 by a team of open-source enthusiasts. At the time, GitHub's stated goal was to support open-source development, especially by hosting open-source source code on the website github.com. Over the next 10 years, GitHub, based on these representations succeeded wildly, attracting nearly 25 million developers.

4.   Developers published Licensed Materials on GitHub pursuant to written Licenses. In particular, the most popular ones share a common term: use of the Licensed Materials requires some form of *attribution*, usually by, among other things, including a copy of the license along with the name and copyright notice of the original author.

5.   On October 26, 2018, Microsoft acquired GitHub for $7.5 billion. Though some members of the open-source community were skeptical of this union, Microsoft repeated one mantra throughout: "Microsoft Loves Open Source". For the first few years, Microsoft's representations seemed credible.

6.   Microsoft invested $1 billion in OpenAI LP in July 2019 at a $20 billion valuation. In 2020, Microsoft became exclusive licensee of OpenAI's GPT-3 language model—despite OpenAI's continued claims its products are meant to benefit "humanity" at large. In 2021, Microsoft began offering GPT-3 through its Azure cloud-computing platform. On October 20, 2022, it was reported that OpenAI "is in advanced talks to raise more funding from Microsoft" at that same $20 billion valuation. Copilot runs on Microsoft's Azure platform. Microsoft has used Copilot to promote Azure's processing power, particularly regarding AI.

7.   On information and belief, Microsoft obtained a partial ownership interest in OpenAI in exchange for its $1 billion investment. As OpenAI's largest investor and largest service provider—specifically in connection with Microsoft's Azure product—Microsoft exerts considerable control over OpenAI.

8.   In June 2021, GitHub and OpenAI launched Copilot, an AI-based product that promises to assist software coders by providing or filling in blocks of code using AI. GitHub charges Copilot users $10 per month or $100 per year for this service. Copilot ignores, violates, and removes the Licenses offered by thousands—possibly millions—of software developers, thereby accomplishing software piracy on an unprecedented scale. Copilot outputs text derived

from Plaintiffs' and the Class's Licensed Materials without adhering to the applicable License

Terms and applicable laws. Copilot's output is referred herein as "Output."

9.   On August 10, 2021, OpenAI debuted its Codex product, which converts natural

language into code and is integrated into Copilot. (Copilot and Codex can be called either AIs or

MLs. Herein they will be referred to as AIs unless a distinction is required.)

10.   Though Defendants have been cagey about what data was used to train the AI,[2]

they have conceded that the training data includes data in vast numbers of publicly accessible

repositories on GitHub,[3] which include and are limited by Licenses.

11.   Among other things, Defendants stripped Plaintiffs' and the Class's attribution,

copyright notice, and license terms from their code in violation of the Licenses and Plaintiffs' and

the Class's rights. Defendants used Copilot to distribute the now-anonymized code to Copilot

users as if it were created by Copilot.

12.   Copilot is run entirely on Microsoft's Azure cloud-computing platform.

13.   Copilot often simply reproduces code that can be traced back to open-source

repositories or open-source licensees. Contrary to and in violation of the Licenses, code

reproduced by Copilot *never* includes attributions to the underlying authors.

14.   GitHub and OpenAI have offered shifting accounts of the source and amount of

the code or other data used to train and operate Copilot. They have also offered shifting

justifications for why a commercial AI product like Copilot should be exempt from these license

requirements, often citing "fair use."

15.   It is not fair, permitted, or justified. On the contrary, Copilot's goal is to replace a

huge swath of open source by taking it and keeping it inside a GitHub-controlled paywall. It

violates the licenses that open-source programmers chose and monetizes their code despite

GitHub's pledge never to do so.

---

[2] "Training" an AI, as described in greater detail below, means feeding it large amounts of data
that it interprets using given criteria. Feedback is then given to it to fine-tune its Output until it
can provide Output with minimal errors.

[3] Repositories are containers for individual coding projects. They are where GitHub users upload
their code and where other users can find it. Most GitHub users have multiple repositories.

II.   JURISDICTION AND VENUE

16.   Plaintiffs bring this action on their own behalf as well as representatives of a Class of similarly situated individuals and entities. They seek to recover injunctive relief and damages as a result and consequence of Defendants' unlawful conduct.

17.   Jurisdiction and venue are proper in this judicial district pursuant to Defendants' violation of the Digital Millennium Copyright Act, 17 U.S.C. §§ 1201–1205 (the "DMCA"); Reverse Passing Off, Unjust Enrichment, and Unfair Competition under the Lanham Act, 15 U.S.C. § 1125; and because a substantial part of the events giving rise to Plaintiff's claims occurred in this District, a substantial portion of the affected interstate trade and commerce was carried out in this District, and three or more of the Defendants reside in this District and/or are licensed to do business in this District. Each Defendant has transacted business, maintained substantial contacts, and/or committed overt acts in furtherance of the illegal scheme and conspiracy throughout the United States, including in this District. Defendants' conduct has had the intended and foreseeable effect of causing injury to persons residing in, located in, or doing business throughout the United States, including in this District.

III.   INTRADISTRICT ASSIGNMENT

18.   Pursuant to Civil Local Rule 3.2 (c) and (e), assignment of this case to the San Francisco Division of the United States District Court for the Northern District of California is proper because a substantial amount of the development of the Copilot product as well as of the interstate trade and commerce involved and affected by Defendants' conduct giving rise to the claims herein occurred in this Division. Furthermore, Defendants GitHub and all the OpenAI entities are headquartered within this Division.

IV.   PARTIES

PLAINTIFFS

19.   Plaintiff J. Doe 3 is a resident of the State of Idaho. Plaintiff Doe 3 published Licensed Materials they owned a copyright interest in to at least one GitHub repository under one of the Suggested Licenses. Specifically, Doe 3 has published Licensed Materials they claim a copyright interest in under the following Suggested Licenses: GNU Affero General Public

License version 3. Plaintiff was, and continues to be, injured during the Class Period as a result of Defendants' unlawful conduct alleged herein.

20.   Plaintiff J. Doe 4 is a resident of the State of South Carolina. Plaintiff Doe 4 published Licensed Materials they owned a copyright interest in to at least one GitHub repository under one of the Suggested Licenses. Specifically, Doe 4 has published Licensed Materials they claim a copyright interest in under the following Suggested Licenses: GNU General Public License v2.0 and GNU General Public License v3.0. Plaintiff was, and continues to be, injured during the Class Period as a result of Defendants' unlawful conduct alleged herein.

DEFENDANTS

21.   Defendant GitHub, Inc. is a Delaware corporation with its principal place of business located at 88 Colin P Kelly Jr Street, San Francisco, CA 94107. GitHub sells, markets, and distributes Copilot throughout the internet and other sales channels throughout the United States, including in this District. GitHub released Copilot on a limited "technical preview" basis on June 29, 2021. On June 21, 2022, Copilot was released to the public as a subscription-based service for individual developers. GitHub is a party to the unlawful conduct alleged herein.

22.   Defendant Microsoft Corporation is a Washington corporation with its principal place of business located at One Microsoft Way, Redmond, Washington 98052. Microsoft announced its acquisition of Defendant GitHub, Inc. on June 4, 2018. On October 26, 2018, Microsoft finalized its acquisition of GitHub. Microsoft owns and operates GitHub. Through its corporate ownership, control of the GitHub Board of Directors, active management, and other means, Microsoft sells, markets, and distributes Copilot. Microsoft is a party to the unlawful conduct alleged herein.

23.   Defendant OpenAI, Inc. is a Delaware nonprofit corporation with its principal place of business located at 3180 18th Street, San Francisco, CA 94110. OpenAI, Inc. is a party to the unlawful conduct alleged herein. It—along with OpenAI, L.P.—programed, trained, and maintains Codex, which infringes all the same rights at Copilot and is also an integral piece of Copilot. Copilot requires Codex to function. OpenAI, Inc. is a party to the unlawful conduct

1   alleged herein. OpenAI, Inc. founded, owns, and exercises control over all the other OpenAI

2   entities, including those set forth in Paragraphs 24–28.

3   24.   Defendant OpenAI, L.P. is a Delaware limited partnership with its principal place

4   of business located at 3180 18th Street, San Francisco, CA 94110. OpenAI, L.P. is a party to the

5   unlawful conduct alleged herein. Its primary activity is research and technology. OpenAI, L.P. is a

6   wholly owned subsidiary of OpenAI, Inc. that is operated for profit. OpenAI, L.P. is the OpenAI

7   entity that co-created Copilot and offers it jointly with GitHub. OpenAI's revenue, including

8   revenue from Copilot, is received by OpenAI, L.P. OpenAI, Inc. controls OpenAI, L.P. directly

9   and through the other OpenAI entities.

10   25.   Defendant OpenAI GP, L.L.C. ("OpenAI GP") is a Delaware limited liability

11   company with its principal place of business located at 3180 18th Street, San Francisco, CA

12   94110. OpenAI GP is the general partner of OpenAI, L.P. OpenAI GP manages and operates the

13   day-to-day business and affairs of OpenAI, L.P. OpenAI GP is liable for the debts, liabilities and

14   obligations of OpenAI, L.P., including litigation and judgments. OpenAI GP is a party to the

15   unlawful conduct alleged herein. Its primary activity is research and technology. OpenAI GP is

16   the general partner of OpenAI, L.P. OpenAI GP was aware of the unlawful conduct alleged herein

17   and exercised control over OpenAI, L.P. throughout the Class Period. OpenAI, Inc. directly

18   controls OpenAI GP.

19   26.   Defendant OpenAI Startup Fund I, L.P. ("OpenAI Startup Fund I") is a Delaware

20   limited partnership with its principal place of business located at 3180 18th Street, San Francisco,

21   CA 94110. OpenAI Startup Fund I was instrumental in the foundation of OpenAI, L.P., including

22   the creation of its business strategy and providing initial funding. Through participation in

23   OpenAI Startup Fund I, certain entities and individuals obtained an ownership interest in

24   OpenAI, L.P. Plaintiffs are informed and believed, and on that basis allege that OpenAI Startup

25   Fund I participated in the organization and operation of OpenAI, L.P. OpenAI Startup Fund I is a

26   party to the unlawful conduct alleged herein. OpenAI Startup Fund I was aware of the unlawful

27   conduct alleged herein and exercised control over OpenAI, L.P. throughout the Class Period.

28

27.   Defendant OpenAI Startup Fund GP I, L.L.C. ("OpenAI Startup Fund GP I") is a Delaware limited liability company with its principal place of business located at 3180 18th Street, San Francisco, CA 94110. OpenAI Startup Fund GP I is the general partner of OpenAI Startup Fund I. OpenAI Startup Fund GP I manages and operates the day-to-day business and affairs of OpenAI Startup Fund I. OpenAI Startup Fund GP I is liable for the debts, liabilities and obligations of OpenAI Startup Fund I, including litigation and judgments. OpenAI Startup Fund GP I was aware of the unlawful conduct alleged herein and exercised control over OpenAI, L.P. throughout the Class Period. OpenAI Startup Fund GP I is a party to the unlawful conduct alleged herein. Sam Altman, co-founder, CEO, and Board member of OpenAI, Inc. is the Manager of OpenAI Startup Fund GP I. OpenAI Startup Fund GP I is the General Partner of OpenAI Startup Fund I, L.P.

28.   Defendant OpenAI Startup Fund Management, LLC ("OpenAI Startup Fund Management") is a Delaware limited liability company with its principal place of business located at 3180 18th Street, San Francisco, CA 94110. OpenAI Startup Fund Management is a party to the unlawful conduct alleged herein. OpenAI Startup Fund Management was aware of the unlawful conduct alleged herein and exercised control over OpenAI, L.P. throughout the Class Period.

V.   AGENTS AND CO-CONSPIRATORS

29.   The unlawful acts alleged against the Defendants in this class action complaint were authorized, ordered, or performed by the Defendants' respective officers, agents, employees, representatives, or shareholders while actively engaged in the management, direction, or control of the Defendants' businesses or affairs.

30.   The Defendants' agents operated under the explicit and apparent authority of their principals.

31.   Each Defendant, and its subsidiaries, affiliates and agents operated as a single unified entity.

32.   Various persons and/or firms not named as Defendants herein may have

participated as coconspirators in the violations alleged herein and may have performed acts and

made statements in furtherance thereof.

33.   Each acted as the principal, agent, or joint venture of, or for other Defendants with

respect to the acts, violations, and common course of conduct alleged herein.

VI.   CLASS ALLEGATIONS

A.   Class Definitions

34.   Plaintiffs bring this action for damages and injunctive relief on behalf of

themselves and all others similarly situated as a class action pursuant to Rules 23(a), 23(b)(2), and

23(b)(3) of the Federal Rules of Civil Procedure, on behalf of the following Classes:

"Injunctive Relief Class" under Rule 23(b)(2):

> All persons or entities domiciled in the United States that, (1)
> owned an interest in at least one US copyright in any work; (2)
> offered that work under one of GitHub's Suggested Licenses[4]; and
> (3) stored Licensed Materials in any public GitHub repositories at
> any time between January 1, 2015 and the present (the "Class
> Period").

---

[4] When a GitHub user creates a new repository, they have the option of selecting one of thirteen licenses from a dropdown menu to apply to the contents of that repository. (They can also apply a different license later, or no license.) The Creative Commons Zero v1.0 Universal and the Unlicense donate the covered work to the public domain and/or otherwise waive all copyrights and related rights. Because they do not contain the necessary provisions nor do they even allow the owner to make copyright claims in most circumstances, they are not included in the Class Definition. We refer to the remaining eleven options as the "Suggested Licenses," which are: (1) Apache License 2.0 ("Apache 2.0"); (2) GNU General Public License version 3 ("GPL-3.0"); (3) MIT License ("MIT"); (4) The 2-Clause BSD License ("BSD 2"); (5) The 3-Clause BSD License ("BSD 3"); (6) Boost Software License ("BSL-1.0"); (7) Eclipse Public License 2.0 ("EPL-2.0"); (8) GNU Affero General Public License version 3 ("AGPL-3.0"); (9) GNU General Public License version 2 ("GPL-2.0"); (10) GNU Lesser General Public License version 2.1 ("LGPL-2.1"); and (11) Mozilla Public License 2.0 ("MPL-2.0"). These Suggested Licenses each contain at least three common requirements for use of the Licensed Materials in a derivative work or copy: attribution to the owner of the Licensed Materials ("Attribution"), inclusion of a copyright notice ("Copyright Notice"), and inclusion of the applicable Suggested License's text ("License Terms").

"Damages Class" under Rule 23(b)(3):

> All persons or entities domiciled in the United States that, (1) owned an interest in at least one US copyright in any work; (2) offered that work under one of GitHub's Suggested Licenses; and (3) stored Licensed Materials in any public GitHub repositories at any time during the Class Period.

These "Class Definitions" specifically exclude the following person or entities:

      a.   Any of the Defendants named herein;

      b.   Any of the Defendants' co-conspirators;

      c.   Any of Defendants' parent companies, subsidiaries, and affiliates;

      d.   Any of Defendants' officers, directors, management, employees, subsidiaries, affiliates, or agents;

      e.   All governmental entities; and

      f.   The judges and chambers staff in this case, as well as any members of their immediate families.

B.   Numerosity

35.   Plaintiffs do not know the exact number of Class members, because such information is in the exclusive control of Defendants. Plaintiffs are informed and believe that there are at least thousands of Class members geographically dispersed throughout the United States such that joinder of all Class members in the prosecution of this action is impracticable.

C.   Typicality

36.   Plaintiffs' claims are typical of the claims of their fellow Class members because Plaintiffs and Class members all own code published under a License. Plaintiffs and the Class published work subject to a License to GitHub later used by Copilot. Plaintiffs and absent Class members were damaged by this and other wrongful conduct of Defendants as alleged herein. Damages and the other relief sought herein is common to all members of the Class.

D.   Commonality & Predominance

37.   Numerous questions of law or fact common to the entire Class arise from Defendants' conduct—including, but not limited to those identified below:

1. **DMCA Violations**

- Whether Defendants' conduct violated the Class's rights under the DMCA when GitHub and OpenAI caused Codex and Copilot to ingest and distribute Licensed Materials without including any associated Attribution, Copyright Notice, or License Terms.

2. **Contract-Related Conduct**

- Whether Defendants violated the Licenses governing use of the Licensed Materials by using them to train Copilot and for republishing those materials without appending the required Attribution, Copyright Notice, or License Terms.

- Whether Defendants interfered in contractual relations between the Class and the public regarding the Licensed Materials by concealing the License Terms.

- Whether GitHub committed Fraud when it promised not to sell or distribute Licensed Materials outside GitHub in the GitHub Terms of Service and Privacy Statement.

3. **Unlawful-Competition Conduct**

- Whether Defendants passed-off the Licensed Materials as its own creation and/or Copilot's creation.

- Whether Defendants were unjustly enriched by the unlawful conduct alleged herein.

- Whether Defendants Copilot-related conduct constitutes Unfair Competition under California law.

4. **Privacy Violations**

- Whether GitHub violated the Class's rights under the California Consumer Privacy Act ("CCPA"), the GitHub Privacy Statement, and/or the California Constitution by, *inter alia*, sharing the Class's sensitive personal information (or, in the alternative, by not addressing an ongoing data breach of which it is aware); creating a product that contains personal data GitHub cannot delete,

alter, nor share with the applicable Class member; and selling the Class's personal data.

- Whether GitHub committed Negligence when it failed to stop a still-ongoing data breach it was and continues to be aware of.

5. Injunctive Relief

- Whether this Court should enjoin Defendants from engaging in the unlawful conduct alleged herein. And what the scope of that injunction would be.

6. Defenses

- Whether any affirmative defense excuses Defendants' conduct.
- Whether any statutes of limitation limit Plaintiffs' and the Class's potential for recovery.
- Whether any applicable statutes of limitation should be tolled as a result of Defendants' fraudulent concealment of their unlawful conduct.

38.   These and other questions of law and fact are common to the Class and predominate over any questions affecting the Class members individually.

E.   Adequacy

39.   Plaintiffs will fairly and adequately represent the interests of the Class because they have experienced the same harms as the Class and have no conflicts with any other members of the Class. Furthermore, Plaintiffs have retained sophisticated and competent counsel ("Class Counsel") who are experienced in prosecuting Federal and state class actions throughout the United States and other complex litigation and have extensive experience advising clients and litigating intellectual property, competition, contract, and privacy matters.

F.   Other Class Considerations

40.   Defendants have acted on grounds generally applicable to the Class, thereby making final injunctive relief appropriate with respect to the Class as a whole.

41.   This class action is superior to alternatives, if any, for the fair and efficient adjudication of this controversy. Prosecuting the claims pleaded herein as a class action will

eliminate the possibility of repetitive litigation. There will be no material difficulty in the management of this action as a class action.

42.   The prosecution of separate actions by individual Class members would create the risk of inconsistent or varying adjudications, establishing incompatible standards of conduct for Defendants.

VII.   FACTUAL ALLEGATIONS

A.   Introduction

43.   This class action against Defendants concerns an OpenAI product called Codex and a GitHub product called Copilot.

44.   OpenAI began development of Codex sometime after OpenAI was founded in December 2015 and released Codex on a limited basis in August 2021.

45.   GitHub began development of Copilot sometime in 2019, released it on a limited basis in June 2021, and released it as a paid subscription service in June 2022.

46.   Codex and Copilot are assistive AI-based systems offered to software programmers. These AI systems are each trained on a large corpus of publicly accessible software code and other materials, including all the Licensed Materials. Defendants describe these products as making coding easier by accepting a code "prompt" from a programmer and emitting a possible completion of that code, which is referred to herein as Output, as set forth above. Defendants claim Codex and Copilot do not retain copies of the materials they are trained on. In practice, however, the Output is often a near-identical reproduction of code from the training data.

47.   Codex and Copilot are related. Copilot is a joint venture between GitHub and OpenAI. Codex is a standalone product released by OpenAI that also "powers GitHub Copilot, which [OpenAI] built and launched in partnership with GitHub."[5] "GitHub Copilot uses the OpenAI Codex to suggest code and entire functions in real-time, right from your editor."[6] The

---

[5] *See* https://openai.com/blog/openai-codex/.

[6] *See* https://github.com/features/copilot/.

1   version of Codex used to power Copilot is modified from OpenAI's standalone version.

2   Defendants have kept secret the details of Codex's modifications and its integration into or

3   interaction with Copilot.

4

B.   Codex Outputs Copyrighted Materials Without Following the Terms of the
Applicable Licenses

5

6   48.   Below is an explanation of how Codex functions. When Codex is prompted with:

7   ```
function isEven(n) {
```

8   it assumes this is the beginning of a function written in the JavaScript language that will test

9   whether a number is even.

10  ```
///
```

11 ```
///
```

12  ```
///
```

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

49.   Based on this assumption, Codex will then provide Output meant to complete the rest of the function. Based on the given prompt, it produced the following response:[7]

```
function isEven(n) {
  if (n == 0)
    return true;
  else if (n == 1)
    return false;
  else if (n < 0)
    return isEven(-n);
  else
    return isEven(n - 2);
}
console.log(isEven(50));
// → true
console.log(isEven(75));
// → false
console.log(isEven(-1));
// → ??
```

50.   The function itself occupies the first ten lines. Six additional lines follow the function, beginning with "`console.log(isEven(50))`". On possible explanation for Codex's inclusion of these lines is to test the "`isEven`" function. Though not part of the function itself, the lines will confirm the function works for certain values. In this case, the code implies that "`isEven(50)`" should return the value "`true`", and "`isEven(75)`" should return "`false`". Those answers are correct.

---

[7] Due to the nature of Codex, Copilot, and AI in general, Plaintiffs cannot be certain these examples would produce the same results if attempted following additional trainings of Codex and/or Copilot. However, these examples are representative of Codex and Copilot's Output at the time just prior to the filing of this Complaint.

1   51.   The penultimate line indicates "`isEven(-1)`" should return "`??`". This is an

2   error, as "`isEven(-1)`" should return "`false`".

3   52.   Codex cannot and does not understand the meaning of software code or any other

4   Licensed Materials. But in training, what became Codex was exposed to an enormous amount of

5   existing software code (its "Training Data") and—with input from its trainers and its own

6   internal processes—inferred certain statistical patterns governing the structure of code and other

7   Licensed Materials. The finished version of Codex, once trained, is known as a "Model."

8   53.   When given a prompt, such as the initial prompt discussed above—"`function`

9   `isEven(n) {`"—Codex identifies the most statistically likely completion, based on the

10   examples it reviewed in training. Every instance of Output from Codex is derived from material in

11   its Training Data. Most of its Training Data consisted of Licensed Materials.

12   54.   Codex does not "write" code the way a human would, because it does not

13   understand the meaning of code. Codex's lack of understanding of code is evidenced when it

14   emits extra code that is not relevant under the circumstances. Here, Codex was only prompted to

15   produce a function called "`isEven`". To produce its answer, Codex relied on Training Data that

16   also appended the extra testing lines. Having encountered this function and the follow-up lines

17   together frequently, Codex extrapolates they are all part of one function. A human with even a

18   basic understanding of how JavaScript works would know the extra lines aren't part of the

19   function itself.

20   55.   Beyond the superfluous and inaccurate extra lines, this "`isEven`" function also

21   contains two major defects. First, it assumes the variable "`n`" holds an integer. It could contain

22   some other kind of value, like a decimal number or text string, which would cause an error.

23   Second, even if "`n`" does hold an integer, the function will trigger a memory error called a "stack

24   overflow" for sufficiently large integers. For these reasons, experienced programmers would not

25   use Codex's Output.

26   56.   Codex does not identify the owner of the copyright to this Output, nor any

27   other—it has not been trained to provide Attribution. Nor does it include a Copyright Notice nor

28   any License Terms attached to the Output. This is by design—Codex was not coded or trained to

1 track or reproduce such data. The Output in the example above is taken from *Eloquent Javascript*

2 by Marijn Haverbeke.[8]

3   57.   Here is the exercise from *Eloquent Javascript*:

4   `// Your code here.`

5

6   `console.log(isEven(50));`

7   `// → true`

8   `console.log(isEven(75));`

9   `// → false`

10   `console.log(isEven(-1));`

11   `// → ??`

12   58.   The exercise includes the " ?? " error. However, for Haverbeke's purposes, this is

13 not an error but a placeholder value for the reader to fill in. Codex—as a mere probabilistic

14 model—fails to recognize this nuance. The inclusion of the double question marks confirms

15 unequivocally that Codex took this code directly from a copyrighted source without following any

16 of the attendant License Terms.

17   59.   Haverbeke provides the following solution to the function discussed above:

18   `function isEven(n) {`

19   `if (n == 0) return true;`

20   `else if (n == 1) return false;`

21   `else if (n < 0) return isEven(-n);`

22   `else return isEven(n - 2);`

23   `}`

24

25   `console.log(isEven(50));`

26

27

28

[8] https://eloquentjavascript.net/code/#3.2. *Eloquent Javascript* is "Licensed under a Creative Commons [A]ttribution-[N]oncommercial license. All code in this book may also be considered licensed under an MIT license." *See* https://eloquentjavascript.net/. Thus, having also been posted on GitHub, the code Codex relied on meets the definition of Licensed Materials.

    CLASS ACTION COMPLAINT

```
// → true
console.log(isEven(75));
// → false
console.log(isEven(-1));
// → false
```

60.   Aside from different line breaks—which are not semantically meaningful in JavaScript—this code for the function "`isEven`" is the same as what Codex produced. The tests are also the same, though in this case Haverbeke provides the right answer for "`isEven(-1)`", which is "`false`". Codex has reproduced Haverbeke's Licensed Material almost verbatim, with the only difference being drawn from a different portion of those same Licensed Materials.

61.   There are many copies of Haverbeke's code stored in public repositories on GitHub, where programmers who are working through Haverbeke's book store their answers.

62.   The MIT license provides that "The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software."[9] Any person taking this code directly from *Eloquent JavaScript* would have direct access to these License Terms and know to follow them if incorporating the Licensed Materials into a derivative work and/or copying them. Codex does not provide these License Terms.

63.   OpenAI Codex's Output would frequently, perhaps even constantly, contain Licensed Materials, i.e., it would have conditions associated with it through its associated license. In its 2021 research paper about Codex called "Evaluating Large Language Models Trained on Code," OpenAI stated Codex's Output is "often incorrect" and can contain security vulnerabilities and other "misalignments" (meaning, departures from what the user requested).

64.   Most open-source licenses require attribution of the author, notice of their copyright, and a copy of the license specifically to ensure that future coders can easily credit all previous authors and ensure they adhere to all applicable licenses. All the Suggested Licenses include these requirements.

---

[9] *See* Appendix A for full text of the MIT License.

Class Action Complaint

65.   Ultimately, Codex derives its value primarily from its ability to locate and output potentially useful Licensed Materials. And from its obfuscation of any rights associated with those materials.

C.   Copilot Outputs Copyrighted Materials Without Following the Terms of the Applicable Licenses

66.   GitHub Copilot works in a similar way to OpenAI Codex. As mentioned above, a modified version of Codex is used as the engine that powers Copilot.

67.   Copilot is installed by the end user as an extension to various code editors, including Microsoft's Visual Studio and VS Code. As the user types into the editor, their code is uploaded in real time to Microsoft's Azure cloud platform, where they become prompts for Copilot.

68.   When we give Copilot the same prompt discussed above in Paragraph 48, "`function isEven(n) {`", it interprets the prompt as the beginning of a function written in the JavaScript language that will test whether a number is even, same as Codex.

69.   However, the Output of Copilot in response to the prompt is different than Codex, namely:

```
function isEven(n) {
    return n % 2 === 0;
}
```

70.   This function is much closer to what a human programmer might use as compared to Codex's older, inaccurate offering. It handles all values and types of "n" correctly. It does not cause a stack overflow for larger values of "n" like the Codex Output.

71.   Copilot's Output, like Codex's, is derived from existing code. Namely, sample code that appears in the online book *Mastering JS*, written by Valeri Karpov.[10] Like Codex's Output, Copilot's is also based upon copyrighted educational material. *Mastering JS* is a set of educational exercises for programmers. Like *Eloquent Javascript*, there are many copies of

---

[10] https://masteringjs.io/tutorials/fundamentals/modulus/.

1 Karpov's exercise stored in public repositories on GitHub. Programmers working through

2 *Mastering JS* store their answers there.

3   72.   If Copilot is prompted with the name of a function that will test whether a number

4 is prime (that is, a number that can only be evenly divided by 1 and itself), namely "`function`

5 `isPrime(n) {`", it returns:

```
function isPrime(n) {
    if (n < 2) {
        return false;
    }
    for (let i = 2; i < n; i++) {
        if (n % i === 0) {
            return false;
        }
    }
    return true;
}
```

17   73.   Though this function will work, it contains an error often made by beginner

18 programmers that makes it much slower than it could be. Namely, the loop in the middle, which

19 checks possible divisors, does not need to check every divisor smaller than "n," only the divisors

20 smaller than the square root of "n". As with Codex, Copilot has no understanding of how the

21 code works. It knows that more functions called "`isPrime`" contain the portion that checks for

22 all divisors smaller than "n", so that is what it offers. It does not return what it "thinks" is best, it

23 returns what it has seen *the most*. It is not writing, it is reproducing (i.e., copying).

24   74.   Like the other examples above—and most of Copilot's Output—this output is

25 nearly a verbatim copy of copyrighted code. In this case, it is substantially similar to the

26 "`isPrime`" function in the book *Think JavaScript* by Matthew X. Curinga et al,[11] which is:

27

---

28 [11] https://matt.curinga.com/think-js/#solving-problems-with-for-loops.

      CLASS ACTION COMPLAINT

```
function isPrime(n) {
    if (n < 2) {
        return false;
    }
    for (let i = 2; i < n; i++) {
        if (n % i === 0) {
            return false;
        }
    }
    return true;
}
```

75.   As with the other examples above, the source of Copilot's Output is a programming textbook. Also like the books the other examples were taken from, there are many copies of Curinga's code stored in public repositories on GitHub where programmers who are working through Curinga's book keep copies of their answers.

76.   The material in Curinga's book is made available under the GNU Free Documentation License. Although this is not one of the Suggested Licenses, it contains similar attribution provisions, namely that "You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License."[12]

77.   As with Codex, Copilot does not provide the end user any attribution of the original author of the code, nor anything about their license requirements. There is no way for the Copilot user to know that they must provide attribution, copyright notice, nor a copy of the license's text. And with regard to the GNU Free Documentation License, Copilot users would not be aware that they are limited in what conditions they can place on the use of derivative works

---

[12] https://matt.curinga.com/think-js/#gnu-free-documentation-license.

they make using this copyrighted code. Had the Copilot user found this code in a public GitHub

repository or a copy of the book it was originally published in, they would find the GNU Free

Documentation License at the same time and be aware of its terms. Copilot finds that code for the

user but excises the license terms, copyright notice, and attribution. This practice allows its users

to assume that the code can be used without restriction. It cannot.

D.   Codex and Copilot Were Trained on Copyrighted Materials Offered Under Licenses

78.   Codex is an AI system. Another way to describe it is a "model." Without Codex,

Copilot, or another AI-code-lookup-tool, code is written both by originating code from the

writer's own knowledge of how to write code as well as by finding pre-written portions of code

that—under the terms of the applicable license—may be incorporated into the coding project.

79.   Unlike a human programmer that has learned how code works and notices when

code it is copying has attached license terms, a copyright notice, and/or attribution, Codex and

Copilot were developed by feeding a corpus of material, called "training data," into them. These

AI programs ingest all the data and, through a complex probabilistic process, predict what the

most likely solution to a given prompt a user would input is. Though more complicated in

practice, essentially Copilot returns the solution it has found in the most projects when those

projects are somehow weighted to adjust for whatever variables Codex or Copilot have identified

as relevant.

80.   Codex and Copilot were not programmed to treat attribution, copyright notices,

and license terms as legally essential. Defendants made a deliberate choice to expedite the release

of Copilot rather than ensure it would not provide unlawful Output.

81.   The words "study" and "training" and "learning" in connection with AI describe

algorithmic processes that are not analogous to human reasoning. An AI models cannot "learn"

as humans do, nor can it "understand" semantics and context the way humans do. Rather, it

detects statistically significant patterns in its training data and provides Output derived from its

training data when statistically appropriate. A "brute force" approach like this would not be

efficient nor even possible for humans. A human could not memorize, statistically analyze, and

easily access thousands of gigabytes of existing code, a task now possible for powerful computers

1   like those that make up Microsoft's Azure cloud platform. To accomplish the same task, a human

2   may search for Licensed Materials that serve their purpose if they believe such materials exist.

3   And if that human finds such materials, they will probably abide by its License Terms rather than

4   risk infringing its owners' rights. At the very least, if they incorporate those Licensed Materials

5   into their own project without following its terms they will be doing so knowingly.

6   **E.   Copilot Was Launched Despite Its Propensity for Producing Unlawful Outputs**

7   82.   GitHub and OpenAI have not provided much detail regarding what data Codex

8   and OpenAI were trained on. Plaintiffs know for certain from GitHub and OpenAI's statements,

9   that both systems were trained on publicly available GitHub repositories, with Copilot having

10 been trained on all available public GitHub repositories. Thus, if Licensed Materials have been

11 posted to a GitHub public repository, Plaintiffs and the Class can be reasonably certain it was

12 ingested by Copilot and is sometimes returned to users as Output.

13   83.   According to OpenAI, Codex was trained on "billions of lines of source code from

14 publicly available sources, including code in public GitHub repositories". Similarly, GitHub has

15 described[13] Copilot's training material as "billions of lines of public code." GitHub researcher

16 Eddie Aftandilian confirmed in a recent podcast[14] that Copilot is "train[ed] on public repos on

17 GitHub."

18   84.   In a recent customer-support message, GitHub's support department clarified

19 certain facts about training Copilot. First, GitHub said that "training for Codex (the model used

20 by Copilot) is done by OpenAI, not GitHub." Second, in its support message, GitHub put

21 forward a more detailed justification for its use of copyrighted code as training data:

22

23

24

25

26

27 [13] https://github.blog/2021-06-30-github-copilot-research-recitation/.

28 [14] https://www.se-radio.net/2022/10/episode-533-eddie-aftandilian-on-github-copilot/.

> Training machine learning models on publicly available data is considered fair use across the machine learning community . . . OpenAI's training of Codex is done in accordance with global copyright laws which permit the use of publicly accessible materials for computational analysis and training of machine learning models, and do not require consent of the owner of such materials. Such laws are intended to benefit society by enabling machines to learn and understand using copyrighted works, much as humans have done throughout history, and to ensure public benefit, these rights cannot generally be restricted by owners who have chosen to make their materials publicly accessible.

The claim that training ML models on publicly available code is widely accepted as fair use is not true. And regardless of this concept's level of acceptance in "the machine learning community," under Federal law, it is illegal.

85.   Former GitHub CEO Nat Friedman said in June 2021—when Copilot was released to a limited number of customers—that "training ML systems on public data is fair use."[15] Friedman's statement is pure speculation; no Court has considered the question of whether "training ML systems on public data is fair use." The Fair Use affirmative defense is only applicable to Section 501 copyright infringement. It is not a defense to violations of the DMCA, Breach of Contract, nor any other claim alleged herein. It cannot be used to avoid liability here. At the same time Friedman asserted "the output [of Copilot] belongs to the operator."

86.   Other open-source stakeholders have made this point already. For example, in June 2021, Software Freedom Conservancy ("SFC"), a prominent open-source advocacy organization, asked Microsoft and GitHub to provide "legal references for GitHub's public legal positions." No references were provided by any of the Defendants.[16]

87.   Beyond the examples above, Copilot regularly Output's verbatim copies of Licensed Materials. For example, Copilot reproduced verbatim well-known code from the game Quake III, use of which is governed by one of the Suggested Licenses—GPL-2.[17]

---

[15] https://twitter.com/natfriedman/status/1409914420579344385/.

[16] https://sfconservancy.org/blog/2022/feb/03/github-copilot-copyleft-gpl/.

[17] https://twitter.com/stefankarpinski/status/1410971061181681674/.

1   88.   Copilot also reproduced code that had been released under a license that allowed

2   its use only for free games and required attribution by including a copy of the license. Copilot did

3   not mention nor include the underlying license when providing a copy of this code as Output.[18]

4   89.   Texas A&M computer-science professor Tim Davis has provided numerous

5   examples of Copilot reproducing code belonging to him without its license or attribution.[19]

6   90.   GitHub concedes that in ordinary use, Copilot will reproduce passages of code

7   verbatim: "Our latest internal research shows that about 1% of the time, a suggestion [Output]

8   may contain some code snippets longer than ~150 characters that matches" code from the

9   training data. This standard is more limited than is necessary for copyright infringement. But

10   even using GitHub's own metric and the most conservative possible criteria, Copilot has violated

11   the DMCA at least tens of thousands of times.

12   91.   In June 2022, Copilot had 1,200,000 users. If only 1% of users have ever received

13   Output based on Licensed Materials and only once each, Defendants have "only" breached

14   Plaintiffs' and the Class's Licenses 12,000 times. However, each time Copilot outputs Licensed

15   Materials without attribution, the copyright notice, or the License Terms it violates the DMCA

16   three times. Thus, even using this extreme underestimate, Copilot has "only" violated the

17   DMCA 36,000 times.[20] Because Copilot constantly Outputs code as a user writes, and because

18   nearly all of Copilot's training data was Licensed Material, this number is most likely

19   exponentially lower than the true number of breaches and DMCA violations.

20

21

22

23

24

25   [18] https://twitter.com/ChrisGr93091552/status/1539731632931803137/.

[19] https://twitter.com/DocSparse/status/1581461734665367554/.

26   [20] These violations of Section 1202 of the DMCA each incur statutory damages of "not less than

27   $2,500 or more than $25,000." 17 U.S.C. § 1203(c)(3)(B). This extremely conservative estimate
     of Defendants' number of direct violations translates to $90 million to $900 million in statutory

28   damages.

F.   Codex and Copilot Were Designed to Withhold Attribution, Copyright Notices, and License Terms from Their Users

92.   Codex and Copilot have no way to determine whether license text or other Copyright Management Information ("CMI")[21] is part of the code it appears immediately before or after. Unless instructed otherwise, it will assume that CMI that usually appears just before a given block of code is an important part of that code or otherwise necessary for it to function.

93.   It is a common practice to provide the applicable license text at the top of every source file in the codebase. The purpose of this practice is to avoid the code from being divorced from the license. This may occur via "vendoring," a method of creating a derivative work by including source files from a copyrighted project directly into another project without following the terms of the license or providing attribution or a copyright notice. Copilot circumvents this protective measure to mask the degree of vendoring it engages in.

94.   Early iterations of Copilot reproduced license text. For example, in a blog post, GitHub noted "In one instance, GitHub Copilot suggested starting an empty file with something it had even seen more than a whopping 700,000 different times during training–that was the GNU General Public License."[22] Copilot no longer suggests licenses in this way because it has been altered not to. As GitHub explains: "GitHub Copilot *has* changed to require a minimum file content. So some of the suggestions flagged here would not have been shown by the current version."

95.   In July 2021, near Copilot's launch, it would sometimes produce license text, attribution, and copyright notices. This CMI was not always accurate. Copilot no longer reproduces these types of CMI, incorrect or otherwise, on a regular basis. It has been altered not to.

---

[21] CMI is defined in detail below in Paragraph 144.

[22] https://github.blog/2021-06-30-github-copilot-research-recitation/.

G.   Open-Source Licenses Began to Appear in the Early 1990s

96.   In 1991, software engineer Linus Torvalds began a project to create a UNIX-like operating system that would run on common PC hardware. This project became known as Linux.

97.   To encourage adoption of his system, and persuade other programmers to contribute, he released Linux under what was then an unusual software license called the GNU General Public License, or GPL.

98.   The GPL is a software license. But whereas most software licenses required payment, software under the GPL is provided for free. Whereas most software licenses did not include source code, GPL software always included source code. And whereas most software licenses prohibited derivative works, the GPL not only allowed it, but encouraged it.

99.   In certain ways, however, the GPL still operated like a traditional software license. For example, consistent with copyright law, it depended on an assertion of copyright by the software author. Even though GPL software was available at no charge, the GPL contained conditions on its users as licensees.

100.   One license requirement was that a program derived from GPL software had to redistribute certain information about that software:

> You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this General Public License and to the absence of any warranty; and give any other recipients of the Program a copy of this General Public License along with the Program.[23]

Failure to adhere to these conditions constituted a violation of the license, triggering the possibility of legal action. Provisions of the GPL are enforceable, and many GPL licensors have sought to enforce GPL licenses though court proceedings and other litigation.

101.   The early years of Linux paralleled the early years of the World Wide Web. The fact that Linux was free and ran on common computer hardware made it a popular choice for web

---

[23] https://www.gnu.org/licenses/old-licenses/gpl-1.0.en.html.

servers. Because of its contrarian GPL licensing, Linux became hugely popular. A large ecosystem

of other programs and tools grew around it. This contributed to the explosive growth of the web

and other network services across the rest of the 1990s.

102.   In turn, the growth of the World Wide Web made it easier for developers in

different places to collaborate on software. The GPL, and licenses like it, were a natural fit for this

kind of collaborative work.

103.   Around 1998, a new name was coined as an umbrella term for these principles of

software licensing and development: *open source*.

H.   Microsoft Has a History of Flouting Open-Source License Requirements

104.   During the 1980s and 1990s, Microsoft was primarily a software company,

focusing largely on operating systems and related applications. These included its DOS operating

system and later, its Windows operating system. Windows generated billions of dollars in revenue

from its sale and licensing as proprietary software for desktop computers and servers. Microsoft

derived substantial income from sale of licensed products and devotes substantial resources to

protecting and enforcing such licenses.

105.   Windows is a graphical operating system. It allows users to view and store files,

run software and games, play videos, and provides a way to connect to the internet.

106.   Linux represented a competitive threat to Windows. It ran on the same hardware.

It performed many of the same functions. It was free. Many programmers at the time considered

Linux to be functionally superior to Windows.

107.   Microsoft has engaged in a problematic practice known as "vaporware," where

products are announced but are in fact late, never manufactured, or canceled. Typically the

company promising vaporware never has any intention of providing it. The term vaporware was

coined by Microsoft in 1982 in reference to the development of its Xenix operating system.

108.   Microsoft described its anti-Linux strategy as "FUD," standing for fear,

uncertainty, and doubt. Microsoft focused extra attention to Linux's open-source aspects.

109.   In 1998, a source at Microsoft leaked what became known as the "Halloween

Documents", revealing Microsoft's thinking on how to counter the competitive threat from

1 Linux. Among other things, the documents emphasized the importance of countering the "long

2 term developer mindshare threat", and concluded that to defeat open source, "[Microsoft] must

3 target a process rather than a company".[24]

4   110.   In 2001, Microsoft CEO Steve Ballmer said "The way the [GPL] is written, if you

5 use any open-source software, you must make the rest of your software open source. . . . Linux is

6 a cancer that attaches itself in an intellectual property sense to everything it touches."[25]

7 Ballmer's summary of GPL licensing was not accurate. In 2001, Linux was being used by

8 corporations of every size. The growth of open source up to that point, and since, has been made

9 possible by the open-source community's respect for and compliance with applicable licenses.

10   111.   In 2001, Microsoft was the defendant in a major software-related antitrust case,

11 *United States v. Microsoft Corporation.*[26] In this case, the U.S. Department of Justice accused

12 Microsoft of maintaining a software monopoly by illegally imposing technical restrictions on

13 manufacturers of personal computers, including "tying" violations related to the Internet

14 Explorer web browser. Judge Thomas Penfield Jackson, who presided over the antitrust trial,

15 opined that Microsoft is "a company with an institutional disdain for both the truth and for rules

16 of law that lesser entities must respect. It is also a company whose 'senior management' is not

17 averse to offering specious testimony to support spurious defenses to claims of its wrongdoing."[27]

18   112.   In 2007, Microsoft admitted that it tried to influence the vote of an ISO open-

19 standards committee by offering money to certain business partners in Sweden to vote for

20 Microsoft's preferred outcome.[28]

21   113.   After observing the rapid growth of Amazon's original cloud computing products,

22 Microsoft has expanded its business into cloud computing, which it has branded Microsoft Azure

23

24 [24] http://www.catb.org/esr/halloween/halloween1.html.

25 [25] https://lwn.net/2001/0607/a/esr-big-lie.php3.

26 [26] No. Civ.A. 00–1457 TPJ.

27 [27] *Jackson v. Microsoft Corp.*, 135 F. Supp. 2d 38 (D.D.C. 2001).

28 [28] https://learn.microsoft.com/en-us/archive/blogs/jasonmatusow/open-xml-the-vote-in-sweden/.

1   or simply Azure. Microsoft announced Azure to developers in 2008. It was formally released in

2   2010. Azure uses large-scale virtualization at Microsoft data centers and offers many hundreds of

3   services, including infrastructure as a service ("IaaS"), platform as a service ("PaaS"), compute

4   services, Azure Active Directory, mobile services, storage services, communication services, data

5   management, messaging, developer services, Azure AI, blockchain, and others.

6   **I.   GitHub Was Designed to Cater to Open-Source Projects**

7   114.   By 2002, Linux had become immensely popular. But the project itself had become

8   unwieldy and had outgrown its reliance on informal systems of managing software source code

9   (also known as *source-control systems*). The Linux community needed something better.

10   115.   Linus Torvalds set about writing a new source-control system. He named his new

11   system Git. He released it under the GPL. It quickly became the source-control system of choice

12   for open-source programmers.

13   116.   A single software project stored in Git is called a *source repository*, commonly

14   shortened to *repository* or just *repo*. A Git source repository would typically be stored on a

15   networked server accessible to a group of programmers.

16   117.   This became less convenient, however, when programmers were distributed

17   among multiple locations, rather than being in a single location. A Git repository could be stored

18   on an internet-accessible server. But setting up that server hardware and being responsible for it

19   was inconvenient and expensive.

20   118.   In 2008, a group of open-source developers in San Francisco, California founded

21   GitHub. GitHub managed internet servers that hosted Git source repositories. With an account at

22   GitHub, an open-source developer could easily set up a Git project accessible to collaborators

23   anywhere in the world. From early on, GitHub's core market was open-source developers, whom

24   it attracted by making many of its hosting services free.

25   119.   Most open-source programmers used GitHub to create "public" repositories,

26   meaning that anyone could view them & access them. GitHub also allowed programmers and

27   organizations to create "private" repositories, which were not accessible from the public GitHub

28   website, and required password access.

120.   Open-source licensing was integral to GitHub. GitHub encouraged open-source developers to understand and use open-source licenses for their work. Many—though not all—public repositories on GitHub carry an open-source license. By convention, this license is stored at the top level of each repository in a file called LICENSE. GitHub's interface also includes a button on the front pages of most repositories users can click to see details of the applicable license. A human user could easily find the license in either of these locations—as could an AI anywhere near as powerful as Codex or Copilot.

121.   Though the GPL is one of the early open-source licenses and remains common, it's not the only open-source license. Examples of other common open-source licenses include the MIT License, the Apache License, and the Berkeley Software Distribution License (all of which are included in the Suggested Licenses).

122.   Though these licenses differ in their wording and their details, most of them share a requirement that a copy of the license be included with any copy, derivative, or redistribution of the software, and that the author's name and copyright notice remains intact. This is not a controversial requirement of open-source licenses—indeed, it has been an integral part of the GPL for over 30 years.

123.   There are also many public repositories on GitHub that have no license. Though GitHub has encouraged awareness of licenses among its users, it has never imposed a default license on public repositories. A public repository without a license is subject to ordinary rules of U.S. copyright.

124.   Open-source developers flocked to GitHub. By 2018, GitHub had become the largest and most successful Git hosting service, hosting millions of users and projects.

125.   In October 2018, Microsoft acquired GitHub for $7.5 billion. It was important to Microsoft that programmers use GitHub. Microsoft had developed a well-deserved poor reputation because of its documented vaporware, FUD, and other business practices, including those targeted at open-source programs and programming, and open-source licensing specifically. Microsoft made false and misleading statements and omissions to assuage such concerns,

1 including its primary mantra intended to win over the open-source community: "Microsoft Loves

2 Open Source."

3 **J.   OpenAI Is Intertwined with Microsoft and GitHub**

4   126.   OpenAI, Inc. is a nonprofit corporation founded in December 2015 by a group that

5 included Greg Brockman, Ilya Sutskever, and other AI researchers; Elon Musk, CEO of Tesla;

6 and Sam Altman, president of Y Combinator, a tech-startup incubator with hundreds of

7 companies in its portfolio. Musk and Altman served as co-chairs of OpenAI, Inc. One of OpenAI,

8 Inc.'s current board members is Reid Hoffman, founder of LinkedIn, which is now a Microsoft

9 subsidiary. Mr. Hoffman is also a member of the Microsoft Board of Directors.

10   127.   Less than a year later, in November 2016, it first partnered with Microsoft. It

11 described the partnership as follows: "We're working with Microsoft to start running most of our

12 large-scale experiments on Azure. This will make Azure the primary cloud platform that OpenAI

13 is using for deep learning and AI, and will let us conduct more research and share the results with

14 the world."

15   128.   Initially, OpenAI, Inc. held itself out as a "non-profit artificial intelligence research

16 company" that sought to shape AI "in the way that is most likely to benefit humanity as a whole."

17   129.   OpenAI, Inc. reportedly secured $1 billion in initial funding, from sources that

18 were largely not disclosed, but included at least most of its founders.

19   130.   OpenAI, Inc. obtained its initial source of training data from its founders'

20 companies. According to reporting at the time, Musk and Altman planned to "pool[] online data

21 from their respective companies" to serve as training data for OpenAI, Inc. projects. Musk

22 planned to contribute data from Tesla; Altman planned to have Y Combinator companies "share

23 their data with OpenAI."[29]

24   131.   In February 2019, Altman created OpenAI, LP, a for-profit subsidiary of the

25 nonprofit entity OpenAI, Inc. The new OpenAI, LP entity would serve as a vessel for accepting

26 traditional outside investment in exchange for equity and distributing profits.

27 ──────────────────────────────

[29] https://www.wired.com/2015/12/elon-musks-billion-dollar-ai-plan-is-about-far-more-than-
28 saving-the-world/.

132.   In July 2019, OpenAI, L.P. accepted a $1 billion investment from Microsoft. In addition to cash, Microsoft would become the exclusive licensor of certain OpenAI, LP products (including GPT-3, described below in Paragraph 135). Also, as part of this alliance, OpenAI, LP would use Microsoft's cloud-computing platform, Azure, exclusively to develop and host its products. Some portion of Microsoft's investment was paid in credits for use of Azure rather than cash. Finally, Microsoft and OpenAI agreed to "jointly build new Azure AI supercomputing technologies."

133.   Azure is a major growth area for Microsoft. In its most recent earnings report on October 25, 2022, "Azure and other cloud services" grew by 35% from the previous quarter, more than any other product.[30] Azure has grown rapidly since Microsoft began its partnership with OpenAI in 2016. Its revenue grew by 50% or more every quarter from 2016 through the first three quarters of 2020.

134.   In May 2020, Microsoft and OpenAI announced they had jointly built a supercomputer in Azure that would be used exclusively by OpenAI to train its AI models. Microsoft's influence over and frequent collaboration with OpenAI has led some to describe Microsoft as "the unofficial owner of OpenAI."[31]

135.   One of OpenAI's projects is GPT-3, a so-called "large language model" designed to emit naturalistic text. When researchers noticed that GPT-3 could also generate software code, they started studying whether they could make a new AI model specifically trained for this purpose. This project became known as Codex.

136.   Sometime after July 2019, OpenAI and Microsoft began collaborating on a code-completion product for GitHub that would use Codex as its underlying model. This product became known as Copilot.

137.   On September 28, 2022, OpenAI released an image-generation AI called DALL-E-2. Much like Copilot, DALL-E-2 removes any attribution and/or copyright notice from the

---

[30] https://www.microsoft.com/en-us/Investor/earnings/FY-2023-Q1/press-release-webcast/.

[31] https://venturebeat.com/ai/what-to-expect-from-openais-codex-api/.

1   images it uses to create derivative works. Like with Codex, here, OpenAI ignores the rights of the

2   owners of copyrights to images it has ingested.

3   138.   In another joint project, Microsoft and OpenAI recently launched a preview of a

4   product called "Azure OpenAI Service."[32] This service will "Leverage large-scale, generative AI

5   models with deep understandings of language and code to enable new reasoning and

6   comprehension capabilities for building cutting-edge applications. Apply these coding and

7   language models to a variety of use cases, such as writing assistance, code generation, and

8   reasoning over data. Detect and mitigate harmful use with built-in responsible AI and access

9   enterprise-grade Azure security."

10 **K.**   **Conclusion of Factual Allegations**

11   139.   Future AI products may represent a bold and innovative step forward. GitHub

12 Copilot and OpenAI Codex, however, do not. Defendants should not have released these

13 products until they could ensure that they did not constantly violate Plaintiffs' and the Class's

14 intellectual-property rights, licenses, and other rights.

15   140.   Defendants have made no attempt to comply with the open-source licenses that

16 are attached to much of their training data. Instead, they have pretended those licenses do not

17 exist, and trained Codex and Copilot to do the same. By simultaneously violating the open-source

18 licenses of tens-of-thousands—possibly millions—of software developers, Defendants have

19 accomplished software piracy on an unprecedented scale. As Microsoft's Co-Founder Bill Gates

20 once said regarding software piracy: "the thing you do is theft."[33]

21   141.   There is no inherent limitation or constraint of AI systems that made any of this

22 necessary. Defendants chose to build AI systems designed to enhance their own profit at the

23 expense of a global open-source community that they had once sought to foster and protect.

24 GitHub and OpenAI are profiting at the expense of Plaintiffs' and the Class's rights.

25

26

---

27 [32] https://azure.microsoft.com/en-us/products/cognitive-services/openai-service/.

28 [33] https://www.digibarn.com/collections/newsletters/homebrew/V2_01/gatesletter.html

VIII. CLAIMS FOR RELIEF

COUNT I
VIOLATION OF THE DIGITAL MILLENIUM COPYRIGHT ACT
17 U.S.C. §§ 1201–1205
(Direct, Vicarious, and Contributory)
(Against All Defendants)

142.   Plaintiffs and the Class hereby repeat and incorporate by reference each preceding and succeeding paragraph as though fully set forth herein.

143.   Plaintiffs and members of the Class own the copyrights to Licensed Materials used to train Codex and Copilot. Copilot was trained on millions—possibly billions—of lines of code publicly available on GitHub. Copilot runs on Microsoft's Azure cloud platform exclusively and Microsoft had input in the creation of Copilot. Microsoft is aware that Copilot ignores License Terms and that it was trained almost exclusively on Licensed Materials.

144.   Plaintiffs and members of the Class included the following Copyright Management Information (as defined in Section 1202(c) of the DMCA) ("CMI") in the Licensed Materials:

      a. copyright notices;

      b. the title and other information identifying the Licensed Materials;

      c. the name of, and other identifying information about, the authors of the Licensed Materials;

      d. the name of, and other identifying information about, the copyright owners of the Licensed Materials;

      e. terms and conditions for use of the Licensed Materials, specifically the Suggested Licenses; and

      f. identifying numbers or symbols referring to CMI or links to CMI.

145.   Defendants did not contact Plaintiffs and the Class to obtain authority to remove or alter CMI from the Licensed Materials within the meaning of the DMCA.

146.   Defendants knew that they did not contact Plaintiffs and the Class to obtain authority to remove or alter CMI from the Licensed Materials within the meaning of the DMCA.

1   147.   As part of the scheme, Defendants did not attempt to contact Plaintiffs to obtain

2 authority to remove or alter CMI from the Licensed Materials within the meaning of the DMCA.

3 In fact, Defendants' removal of CMI made it difficult or impossible to contact Plaintiffs and the

4 Class to obtain authority to remove or alter CMI from the Licensed Materials within the meaning

5 of the DMCA. Rather, Defendants removed or altered CMI from open-source code that is owned

6 by Plaintiffs and the Class after the code was uploaded to a GitHub repository by incorporating it

7 into Copilot with its CMI removed.

8   148.   Without the authority of Plaintiffs and the Class, Defendants intentionally

9 removed or altered CMI from the Licensed Materials after they were uploaded to one or more

10 GitHub repositories.

11   149.   Defendants had access to but were not licensed by Plaintiffs nor the Class to train

12 any machine learning, AI, or other pseudo-intelligent computer program, algorithm, or other

13 functional prediction engine using the Licensed Materials.

14   150.   Defendants had access to but were not licensed by Plaintiffs nor the Class to

15 incorporate the Licensed Materials into Copilot.

16   151.   Defendants had access to but were not licensed by Plaintiffs nor the Class to create

17 Derivative Works[34] based upon the Licensed Materials.

18   152.   Defendants had access to but were not licensed by Plaintiffs nor the Class to

19 distribute the Licensed Materials as they do through Copilot.

20   153.   Without the authority of Plaintiffs and the Class, Defendants distributed CMI

21 knowing that the CMI had been removed or altered without authority of the copyright owner or

22 the law with respect to the Licensed Materials.

23   154.   Defendants distributed copies of the Licensed Materials knowing and intending

24 that CMI had been removed or altered without authority of the copyright owner or the law, with

25 respect to the Licensed Materials.

26

27 [34] "Derivative Works" as used herein refers to Copilot's Output to the extent they are derived
from Licensed Materials. The definition also includes the Copilot product itself, which is a
28 Derivative Work based upon a large corpus of Licensed Materials.

1   155.   Defendants removed or altered CMI from the Licensed Materials knowing and

2   intending that it would induce, enable, facilitate, or conceal infringement of copyright.

3   156.   Without the CMI associated with the Licensed Materials, Copilot users are

4   induced or enabled to copy the Licensed Materials. Because CMI has been removed, Copilot

5   users do not know whether Output is owned by someone else and subject to restrictions on use.

6   Without the CMI, copyright infringement is facilitated or concealed, because Plaintiffs and the

7   Class are prevented from knowing or learning that the Output is based upon one or more of the

8   Licensed Materials. Use of the Licensed Materials is not infringement when the terms of the

9   applicable Suggested License are followed. Had the CMI not been removed, Copilot users would

10   be aware of the Licenses and their obligations under them. The terms of the applicable Suggested

11   License would have allowed those users to use the Licensed Materials without infringement. By

12   withholding and concealing license information and other CMI, Defendants prevented Copilot

13   users from making non-infringing use of the Licensed Materials. This contradicts the express

14   wishes of Plaintiffs and the Class, which are set forth explicitly in the Suggested Licenses under

15   which the Licensed Materials are offered.

16   157.   Defendants removed or altered CMI from Licensed Materials owned by Plaintiffs

17   and the Class while possessing reasonable grounds to know that it would induce, enable, facilitate,

18   and/or conceal infringement of copyright in violation of the DMCA. By omitting and concealing

19   CMI from Copilot's Output, Defendants have reasonable grounds to know that innocent

20   infringers are induced or enabled to copy the Licensed Materials, because CMI has been

21   removed. Without the CMI, Defendants have reasonable grounds to know copyright infringement

22   is facilitated or concealed, because Plaintiffs and the Class have the difficult or impossible task of

23   proving the Licensed Materials belong to them.

24   158.   Defendants knowingly provided CMI that is false with respect to the Licensed

25   Materials. Defendants have a business practice of asserting and/or implying that Copilot is the

26   author of the Licensed Materials. Defendants knowingly distributed CMI that is false, with

27   respect to the Licensed Materials. Defendants have a business practice of asserting and/or

28   implying that Copilot is the author of the Licensed Materials.

159.   Defendants provided or distributed false CMI from the Licensed Materials with respect to Copilot's Output with the intent and foreseeable result to induce, enable, facilitate, or conceal infringement. Defendants have a business practice of asserting and/or implying that Copilot is the author of the Licensed Materials. This false CMI induces or enables Defendants or Copilot users to copy the Licensed Materials. Defendants' false description of the source of Copilot's Output facilitated or concealed infringement by Defendants and Copilot users because Plaintiffs and the Class have the difficult or impossible task of proving that the copyrights to the suggested portions of their Licensed Materials belong to them once those Licensed Materials have been delinked from all identifying information and all license terms governing their use.

160.   The profits attributable to Defendants' violation of the DMCA include the revenue from: Copilot subscription fees, sales of or subscriptions to Defendants' Copilot-related products and/or services that are used to run Copilot, hosting Copilot on Azure, and any other of Defendants' products that contain copies of the Licensed Materials without all the original CMI. The Licensed Materials add nearly all value to the Copilot product because the purpose of Copilot is to provide code and the source of that code is the Licensed Materials. Without the Licensed Materials, Copilot would not be functional.

161.   On information and belief, Defendants could have trained Copilot to include attribution, copyright notices, and license terms when it provides Output covered by a License.

162.   Defendants did not request or obtain permission from Plaintiffs and the Class to use the Licensed Materials for Defendants' Copilot product.

163.   Defendants use of the Licensed Materials does not follow the requirements of the Suggested Licenses associated with the Licensed Materials. In particular, Copilot fails to provide attribution for the creator nor the owner of the Work. Copilot fails to include the required copyright notice included in the License. Copilot fails to include the applicable Suggested License's text.

164.   Defendants are sophisticated with respect to intellectual property matters related to open-source code. Microsoft in particular has extensive experience granting licenses, obtaining licenses, and enforcing license terms. Its most recent Annual Report states:

We protect our intellectual property investments in a variety of ways. We work actively in the U.S. and internationally to ensure the enforcement of copyright, trademark, trade secret, and other protections that apply to our software and hardware products, services, business plans, and branding. We are a leader among technology companies in pursuing patents and currently have a portfolio of over 69,000 U.S. and international patents issued and over 19,000 pending worldwide. While we employ much of our internally-developed intellectual property exclusively in our products and services, we also engage in outbound licensing of specific patented technologies that are incorporated into licensees' products. From time to time, we enter into broader cross-license agreements with other technology companies covering entire groups of patents. We may also purchase or license technology that we incorporate into our products and services. At times, we make select intellectual property broadly available at no or low cost to achieve a strategic objective, such as promoting industry standards, advancing interoperability, supporting societal and/or environmental efforts, or attracting and enabling our external development community. **Our increasing engagement with open source software will also cause us to license our intellectual property rights broadly in certain situations.**

Microsoft Corporation Annual Report, Form 10-K at 27 (July 28, 2022) (emphasis added).[35]

165. GitHub, which offers the Copilot product jointly with OpenAI, also has extensive experience with the DMCA. GitHub knows or reasonably should know that the Licensed Materials it hosts are subject to copyright. It provides the language of the Suggested Licenses to users, all of which include copyright notices. Its 2022 Transparency Report—January to June[36] states: "Copyright-related takedowns (which we often refer to as DMCA takedowns) are particularly relevant to GitHub because so much of our users' content is software code and can be eligible for copyright protection."[37] In the first six months of 2022, GitHub processed 1220 DMCA takedown requests. Its DMCA Takedown Policy[38] notes "GitHub probably never would have existed without the DMCA."

---

[35] https://microsoft.gcs-web.com/static-files/07cf3c30-cfc3-4567-b20f-f4b0f0bd5087/.

[36] https://github.blog/2022-08-16-2022-transparency-report-january-to-june/.

[37] https://github.blog/2022-08-16-2022-transparency-report-january-to-june/.

[38] https://docs.github.com/en/site-policy/content-removal-policies/dmca-takedown-policy#what-is-the-dmca/.

CLASS ACTION COMPLAINT

166.   GitHub also knows or reasonably should know the portions of the DMCA giving rise to Plaintiffs' claim. In its 2021 Transparency Report, "Before removing content based on alleged circumvention of copyright controls (under Section 1201 of the US DMCA or similar laws in other countries), we carefully review both the legal and technical claims, and we sponsor a Developer Defense Fund to provide developers with meaningful access to legal resources."[39]

167.   GitHub is aware that Copilot's removal of CMI is illegal. For example, it states that "publishing or sharing tools that enable circumvention are not [permitted]"[40] and "Distributing tools that enable circumvention is prohibited, even if their use by developers falls under the exemption [for security research]."[41] GitHub has also frequently published articles discussing the DMCA, its application, and the Copyright Office's guidance on its scope and exceptions.[42]

168.   Unless Defendants are enjoined from violating the DMCA, Plaintiffs and the Class will suffer great and irreparable harm by depriving them of the right to identify and control the reproduction and/or distribution of their copyrighted works, to have the terms of their open-source licenses followed, and to pursue copyright-infringement remedies. Defendants will not be damaged if they are required to comply with the DMCA. Plaintiffs and the Class members are therefore entitled to an injunction barring Defendants from violating the DMCA and impounding any device or product that is in the custody or control of Defendants and that the court has reasonable cause to believe was involved in a violation of the DMCA.

169.   Plaintiffs and the Class are further entitled to recover from Defendants the actual or statutory damages Plaintiffs and the Class sustained pursuant to 17 U.S.C. § 1203(c) and for Plaintiffs' and the Class's costs and attorneys' fees in enforcing the Licenses. Plaintiffs and the Class are also entitled to recover as restitution from Defendants for any unjust enrichment,

---

[39] https://github.blog/2022-01-27-2021-transparency-report/.

[40] https://github.blog/2020-11-19-take-action-dmca-anti-circumvention-and-developer-innovation/#what-dmca-exemptions-do-not-do/.

[41] https://github.blog/2021-11-23-copyright-office-expands-security-research-rights/.

[42] *See, e.g.*, Footnotes 36–41.

1   including gains, profits, and advantages that Defendants have obtained as a result of their breach

2   of the Licenses.

3   170.   Defendants conspired together and acted jointly and in concert pursuant to their

4   scheme to commit the acts that violated the DMCA alleged herein.

5   171.   Defendants induced Copilot users to unknowingly violate the DMCA by

6   withholding attribution, licensing, and other information as described herein.

<div align="center">

COUNT II
BREACH OF CONTRACT—OPEN-SOURCE LICENSE VIOLATIONS
Common Law
(Against All Defendants)

</div>

10   172.   Plaintiffs and the Class hereby repeat and incorporate by reference each preceding

11   and succeeding paragraph as though fully set forth herein.

12   173.   Plaintiffs and the Class offer code under various Licenses, the most common of

13   which are set forth in Appendix A. Use of each of the Licensed Materials is allowed only pursuant

14   to the terms of the applicable Suggested License.

15   174.   Plaintiffs and the Class granted Defendants a license to copy, distribute, and/or

16   create Derivative Works under the Suggested Licenses. Each of the Suggested Licenses requires

17   at least (1) that attribution be given to the owner of the Licensed Materials used, (2) inclusion of a

18   copyright notice for the Licensed Materials used, and (3) inclusion of the terms of the applicable

19   Suggested License. When providing Output, Copilot does not comply with any of these terms.

20   175.   Defendants accepted the terms of Plaintiffs' and the Class's Licenses when it used

21   the licensed code to create Copilot and when it incorporated the licensed code into Copilot. They

22   have accepted and continue to accept the applicable Licenses every time Copilot Output's

23   Plaintiffs' or the Class's copyrighted code. As such, contracts have been formed between

24   Defendants on the one hand and Plaintiffs and the Class on the other.

25   176.   Plaintiffs and the Class have performed each of the conditions, covenants, and

26   obligations imposed on them by the terms of the License associated with their Licensed

27   Materials.

28

1   177.   Plaintiffs and members of the Class hold the copyright in the contents of one or

2   more code repositories that have been hosted on GitHub's platform.

3   178.   Plaintiffs and the Class have appended one of the Suggested Licenses to each of

4   the Licensed Materials.

5   179.   Plaintiffs and the Class did not know about, authorize, approve, or license the

6   Defendants' use of the Licensed Materials in the matter at issue in this Complaint before they

7   were used by Defendants.

8   180.   Defendants have substantially and materially breached the applicable Licenses by

9   failing to provide the source code of Copilot nor a written offer to provide the source code upon

10   the request of each licensee.

11   181.   Defendants have substantially and materially breached the applicable Licenses by

12   failing to provide attribution to the creator and/or owner of the Licensed Materials.

13   182.   Defendants have substantially and materially breached the applicable Licenses by

14   failing to include copyright notices when Copilot Outputs copyrighted OS code.

15   183.   Defendants have substantially and materially breached the applicable Licenses by

16   failing to identify the License applicable to the Work and/or including its text when Copilot

17   Outputs code including a portion of a Work.

18   184.   Plaintiffs and the Class have suffered monetary damages as a result of Defendants'

19   conduct.

20   185.   The conduct of Defendants is causing and, unless enjoined and restrained by this

21   Court, will continue to cause Plaintiffs and the Class great and irreparable injury that cannot fully

22   be compensated or measured in money.

23   186.   As a direct and proximate result of these material breaches by Defendants,

24   Plaintiffs and the Class are entitled to an injunction requiring Defendants to comply with all the

25   terms of any License governing use of code that was used to train Copilot, otherwise incorporated

26   into Copilot, and/or reproduced as Output by Copilot.

27   187.   Plaintiffs and the Class are further entitled to recover from Defendants the

28   damages Plaintiffs and the Class sustained—including consequential damages—for Plaintiffs' and

the Class's costs in enforcing their contractual rights. Plaintiffs and the Class are also entitled to recover as restitution from Defendants for any unjust enrichment, including gains, profits, and advantages that Defendants have obtained as a result of their breach of contract.

COUNT III
TORTIOUS INTERFERENCE IN A CONTRACTUAL RELATIONSHIP
Common Law
(Against All Defendants)

188.   Plaintiffs and the Class hereby repeat and incorporate by reference each preceding and succeeding paragraph as though fully set forth herein.

189.   Defendants have wrongfully interfered with the business interests and expectations of Plaintiffs and the Class by improperly using Copilot to create Derivative Works that compete against OSC.

190.   At GitHub's upcoming yearly conference, GitHub Universe 2022, it will host a presentation called "How to compete with open source—and win."

191.   Plaintiffs and the Class have suffered monetary, reputational, and other damages as a result of Defendants' conduct.

192.   The harm was the actual, proximate, intentional, direct, and foreseeable consequence of Defendant's conduct.

193.   The conduct of Defendants is causing and, unless enjoined and restrained by this Court, will continue to cause Plaintiffs and the Class great and irreparable injury that cannot fully be compensated or measured in money.

COUNT IV
FRAUD
Common Law
(Against GitHub)

194.   Plaintiffs and the Class hereby repeat and incorporate by reference each preceding and succeeding paragraph as though fully set forth herein.

195.   GitHub made certain representations to Plaintiffs and the Class to induce them to publicly post their code on GitHub. Specifically, in both its Terms of Service and its Privacy Statement, GitHub promises not to sell Licensed Materials or anything else uploaded to or shared

1   with GitHub. It also promises not to distribute Licensed Materials outside GitHub. As explained

2   above, Copilot operates on an individual's computer as an extension to their editor as well as on

3   Microsoft's Azure cloud platform. Neither are part of GitHub. It Outputs in the user's editor,

4   which is not part of GitHub.

5   196.   Plaintiffs and the Class relied upon those representations in choosing to upload

6   Licensed Materials to GitHub. GitHub has long held itself out as the best place to host open-

7   source code repositories. It has courted the business of users it expects will include Licenses with

8   their code. It facilitates this by allowing users to easily select the name of a license, including the

9   Suggested Licenses, when creating a repository rather than finding the text of the license and

10   adding it themselves. GitHub provides the terms, it can hardly claim to be unaware of what they

11   are or what they mean. If it didn't understand the requirements of a given Suggested License, it

12   would not have provided it as an option to its users.

13   197.   GitHub failed to honor its representations in creating and operating Copilot. It

14   sells Plaintiffs' and the Class's Licensed Materials as part of Copilot. It also distributes them. It

15   does so without following any of the License Terms.

16   198.   As such, GitHub failed to honor its representations in operating Copilot.

17   199.   The conduct of GitHub is causing and, unless enjoined and restrained by this

18   Court, will continue to cause Plaintiffs and the Class great and irreparable injury that cannot fully

19   be compensated or measured in money. Namely, it will continue the proliferation of copies of

20   Licensed Materials divorced from their licenses and identifying information until infringement is

21   so prevalent no amount of enforcement by Plaintiffs and the Class could stop its spread.

COUNT V
FALSE DESIGNATION OF ORIGIN—REVERSE PASSING OFF
15 U.S.C. § 1125
(GitHub and OpenAI)

25   200.   Plaintiffs and the Class hereby repeat and incorporate by reference each preceding

26   and succeeding paragraph as though fully set forth herein.

27   201.   GitHub and OpenAI have used or made, and will continue to use or make, in

28   commerce throughout the United States, including in California, one or more words, terms,

names, symbols, or devices, or any combination thereof, or any false and/or misleading

designation of origin, false and/or misleading description of fact, or false and/or misleading

representation of fact that is likely to cause consumer confusion, or to cause mistake, or to deceive

as to the affiliation, connection, or association of Plaintiffs' and the Class's Licensed Materials

and Copilot, or as to the origin, sponsorship, or approval of Plaintiffs' and the Class's Licensed

Materials and Copilot.

202.   As a result, GitHub and OpenAI have intentionally violated 15 U.S.C. §

1125(a)(1)(A).

203.   As an actual and proximate result of GitHub's and OpenAI's acts, Plaintiffs and

the Class have suffered and continue to suffer harm.

COUNT VI
UNJUST ENRICHMENT
Cal. Bus. & Prof. Code §§ 17200, *et seq.* and Common Law
(GitHub and OpenAI)

204.   Plaintiffs and the Class hereby repeat and incorporate by reference each preceding

and succeeding paragraph as though fully set forth herein.

205.   Plaintiffs and the Class have invested substantial time and energy in creating the

Licensed Materials.

206.   GitHub and OpenAI have unjustly utilized access to Licensed Materials hosted on

GitHub. This code is used to create Derivative Works that are licensed to third parties in

exchange for, *inter alia*, compliance with applicable License terms.

207.   GitHub and OpenAI derive profit or other benefits from removal of attribution,

copyright notices, and license terms from Licensed Materials and reselling it as Output through

Copilot.

208.   It would be unjust for GitHub and OpenAI to retain those benefits.

209.   Plaintiffs and the Class have suffered monetary damages as a result of GitHub's

and OpenAI's conduct.

210.   The conduct of GitHub and OpenAI is causing and, unless enjoined and restrained by this Court, will continue to cause Plaintiffs and the Class great and irreparable injury that cannot fully be compensated or measured in money.

COUNT VII
UNFAIR COMPETITION
15 U.S.C. § 1125; *Cal. Bus. & Prof. Code* §§ 17200, *et seq.*; and Common Law
(GitHub and OpenAI)

211.   Plaintiffs and the Class hereby repeat and incorporate by reference each preceding and succeeding paragraph as though fully set forth herein.

212.   GitHub and OpenAI have engaged in unlawful business practices, including:

    a.   Violations of Plaintiffs' and the Class's rights under the DMCA;

    b.   Tortious interference in Plaintiffs' and the Class's contractual relations with users of their code;

    c.   Passing off Copilot's Output as originating from Copilot, GitHub, and/or OpenAI;

    d.   Privacy violations, namely violation of the CCPA and negligent handling of personal information.

The details of the unlawful business practices are set forth herein.

213.   Plaintiffs and the Class have suffered monetary damages as a result of GitHub's and OpenAI's conduct.

214.   The conduct of GitHub and OpenAI is causing and, unless enjoined and restrained by this Court, will continue to cause Plaintiffs and the Class great and irreparable injury that cannot fully be compensated or measured in money.

COUNT VIII
BREACH OF CONTRACT
VIOLATION OF GITHUB PRIVACY POLICY and TERMS OF SERVICE
Cal. Bus. & Prof. Code § 22575–22579; *Cal. Civ. Code* § 1798.150; and Common Law
(Against GitHub)

215.   Plaintiffs and the Class hereby repeat and incorporate by reference each preceding and succeeding paragraph as though fully set forth herein.

216.   Plaintiffs and the Class are GitHub users who have accepted GitHub's Terms of Service. As a result, Plaintiffs and the Class have formed a contract, the terms of which are set

1   forth in GitHub's Terms of Service—including the additional GitHub Copilot Terms from

2   GitHub Terms for Additional Products and Features.

3   217.   Plaintiffs and the Class are GitHub users who have accepted GitHub's Privacy

4   Statement. As a result, Plaintiffs and the Class have formed a contract.

5   218.   GitHub's Privacy Statement, Terms of Service, and GitHub Copilot Terms share

6   definitions and refer to each other. As such, they are collectively referred to herein as "GitHub's

7   Policies" unless a distinction is necessary and are attached as Exhibit 1.

8   219.   Plaintiffs and the Class have performed each of the conditions, covenants, and

9   obligations imposed on them by the terms of GitHub's Policies.

10   220.   GitHub has substantially and materially breached GitHub's Policies in the

11   following ways:

12   a.   Sharing Plaintiffs' and the Class's personal data with unauthorized third parties in

13   violation of the GitHub Privacy Statement;

14   b.   Selling and distributing Plaintiffs' and the Class's personal data in contravention

15   of the GitHub Policies;

16   c.   Use of Plaintiffs' and the Class's personal data after the GitHub Privacy Statement

17   explicitly claims it will be deleted;

18   d.   Use and distribution of Plaintiffs' and the Class's personal data outside the

19   limitations set forth in the GitHub Privacy Statement.

20   221.   Plaintiffs and the Class have suffered monetary damages as a result of GitHub's

21   conduct.

22   222.   GitHub's conduct is causing and, unless enjoined and restrained by this Court,

23   will continue to cause Plaintiffs and the Class great and irreparable injury that cannot fully be

24   compensated or measured in money.

25   223.   As a direct and proximate result of these material breaches by GitHub, Plaintiffs

26   and the Class are entitled to an injunction requiring GitHub to comply with all the terms of the

27   GitHub Policies.

28

224.   Plaintiffs and the Class are further entitled to recover from GitHub the damages Plaintiffs and the Class sustained—including consequential damages—for Plaintiffs' and the Class's costs in enforcing GitHub's Policies. Plaintiffs and the Class are also entitled to recover as restitution from GitHub for any unjust enrichment, including gains, profits, and advantages that it has obtained as a result of its breaches of the GitHub Policies.

COUNT IX
VIOLATION OF THE CALIFORNIA CONSUMER PRIVACY ACT
Cal. Civ. Code § 1798.150
(GitHub and OpenAI)

225.   Plaintiffs and the Class hereby repeat and incorporate by reference each preceding and succeeding paragraph as though fully set forth herein.

226.   At all relevant times, GitHub and OpenAI were "businesses" under the terms of Cal. Civ. Code § 1798.140(g) of the California Consumer Privacy Act ("CCPA") as sole proprietorships, partnerships, limited liability companies, corporations, associations, or other legal entities operating in the State of California that collect consumers' personal information. GitHub and OpenAI either have annual operating revenue above $25 million, collect the personal information of 50,000 or more California residents annually, or derive at least fifty percent of their annual revenue from the sale of personal information of California residents.

227.   At all relevant times, Plaintiffs and the Class were "consumers" under the terms of the CCPA as natural persons as defined in Section 17014 of Title 18 of the California Code of Regulations.

228.   "Personal Information" is defined in Section 1798.140(o)(1) of the CCPA, which protects consumers' personal information from collection, use, or sale by businesses without consumers' notice and consent.

229.   GitHub and OpenAI violated the CCPA by using customers' PII without providing the required notice under the CCPA. GitHub and OpenAI did not notify Plaintiffs nor the Class they were using, distributing, or selling their PII to unauthorized third parties, namely Copilot users.

230. GitHub and OpenAI also violated the CCPA by failing to provide notice to its customers of their right to opt-out of the disclosure of their PII to unauthorized third parties, namely Copilot users.

231. GitHub and OpenAI also violated the CCPA by incorporating Plaintiffs' and the Class's personal information into Copilot with no way to alter or delete. And also with no way to share that personal data with Plaintiffs or the Class upon request.

232. GitHub and OpenAI also violated the CCPA by failing to provide a clear and conspicuous link entitled "Do Not Sell My Personal Information" to a webpage that enables a consumer—or a person authorized by a consumer—to opt out of the sale of Plaintiffs' and the Class's personal data through Copilot.

233. By the acts described above, GitHub and OpenAI violated the CCPA by negligently, carelessly, and recklessly collecting, maintaining, and controlling their customers' sensitive personal information and by engineering, designing, maintaining, and controlling systems that exposed their customers' sensitive personal information of which GitHub and OpenAI had control and possession to the risk of exposure to unauthorized persons, thereby violating their duty to implement and maintain reasonable security procedures and practices appropriate to the nature of the information to protect the personal information. GitHub and OpenAI allowed unauthorized users to view, use, manipulate, exfiltrate, and steal the nonencrypted and nonredacted personal information of Plaintiffs and other customers, including their personal and financial information.

234. GitHub and OpenAI are aware of Copilot's propensity for revealing PII. They have altered Copilot to force it to provide mock PII in place of real PII. But due to the many varieties of PII and various forms it can take, there is no way GitHub and OpenAI can ensure that it will not, at times, reproduce unaltered PII. Unless PII is removed before Copilot is trained on a given repository, Copilot will incorporate that PII and could, under the right circumstances, reproduce that PII verbatim. Unless all PII is removed from Copilot, it will continue to sometimes provide unaltered PII in its Output. On information and belief, some Copilot users have devised methods to increase the chances Copilot will produce accurate PII.

COUNT X
NEGLIGENCE—NEGLIGENT HANDLING OF PERSONAL DATA
Common Law
(GitHub and OpenAI)

235.   Plaintiffs and the Class hereby repeat and incorporate by reference each preceding and succeeding paragraph as though fully set forth herein.

236.   GitHub and OpenAI owed a duty of reasonable care toward Plaintiffs and the Class based upon GitHub's and OpenAI's relationship to them. This duty is based upon GitHub's and OpenAI's contractual obligations, custom and practice, right to control information in its possession, exercise of control over the information in its possession, authority to control the information in its possession, and the commission of affirmative acts that resulted in said harms and losses. Additionally, this duty is based on the requirements of California Civil Code section 1714 requiring all "persons," including GitHub and OpenAI, to act in a reasonable manner toward others. This duty is also based on the specific statutory duties imposed on GitHub and OpenAI under California Civil Code sections 1798.100, *et seq.*, as businesses operating in the State of California that either have annual operating revenue above $25 million, collect the personal information of 50,000 or more California residents annually, or derive at least 50 percent of their annual revenue from the sale of personal information of California residents.

237.   GitHub and OpenAI breached their duties by negligently, carelessly, and recklessly collecting, maintaining, and controlling their customers' sensitive personal information and engineering, designing, maintaining, and controlling systems—including Copilot—that exposed and continue to expose their customers' sensitive personal information of which GitHub and OpenAI had control and possession to the risk of exposure to unauthorized persons.

238.   GitHub and OpenAI also committed per se breaches of said duty by negligently violating the dictates of California Civil Code sections 1798.82, *et seq.*, and 1798.100, *et seq.*, and the provisions of the California Constitution enshrining the right to privacy, by failing to inform Plaintiffs and the Class of the access to their sensitive personal information by unauthorized persons expeditiously and without delay and failing to adequately safeguard this information from unauthorized access even after GitHub and OpenAI became aware of multiple instances of release of this information by Copilot. The provisions of the California Civil Code and the

1   California Constitution that GitHub and OpenAI violated were enacted to protect the class of

2   Plaintiffs here involved from the type of injury here incurred, namely their right to privacy and

3   the protection of their personal data. Plaintiffs and the Class were within the class of persons and

4   consumers who were intended to be protected by California Civil Code sections 1798.82, *et seq.*,

5   and 1798.100, *et seq.*

6   239.   As a direct consequence of the actions described herein, and the breaches of

7   duties indicated thereby, unauthorized users gained access to, exfiltrated, stole, and gained

8   disclosure of the sensitive personal information of Plaintiffs and the Class, causing them harms

9   and losses including but not limited to economic loss, the loss of control over the use of their

10 identity, harm to their constitutional right to privacy, lost time dedicated to cure harm to their

11 privacy, the need for future expenses and time dedicated to the recovery and protection of further

12 loss, and privacy injuries associated with having their sensitive personal and financial information

13 disclosed.

14 <div align="center">

COUNT XI
CIVIL CONSPIRACY
Common Law
(Against All Defendants)
</div>

15

16

17   240.   Plaintiffs and the Class hereby repeat and incorporate by reference each preceding

18 and succeeding paragraph as though fully set forth herein.

19   241.   On information and belief, Microsoft, GitHub, OpenAI, and the Individual

20 Defendants have worked together to create Copilot. In creating Copilot, Defendants willfully

21 avoided determining whether and how Copilot's training and Output may violate the rights of

22 Plaintiffs and the Class and other stakeholders. This is because Defendants understood that

23 through Copilot they would be engaging in a variety of unlawful conduct. Defendants conduct

24 resulted in violations of Plaintiffs' and the Class's rights as set forth herein.

25   242.   On information and belief, OpenAI derives a financial or other valuable benefit

26 from the sale of Copilot. In exchange, OpenAI provided Microsoft an exclusive license to use its

27 GPT-3 language model.

28

1   243.   On information and belief, Microsoft derives a financial benefit from sales of

2 Copilot through payments or other form of compensation in exchange for GitHub's and

3 OpenAI's use of Azure to run Copilot.

4   244.   GitHub, Microsoft, and OpenAI agreed to a common plan or design to create, sell,

5 and run Copilot to commit and conceal the following tortious acts:

6   a. Violations of Plaintiffs' and the Class's rights under the DMCA Section 1202;

7   b. Tortious interference in Plaintiffs' and the Class's contractual relations with users

8   of their code;

9   c. Passing off Copilot's Output as originating either from Copilot itself or from

10   Defendants;

11   d. Unfair competition with Plaintiffs and the Class by releasing and operating

12   Copilot; and

13   e. Privacy violations, namely violation of the CCPA and negligent handling of

14   personal information.

COUNT XII
DECLARATORY RELIEF
28 U.S.C. § 2201(a) and Cal. Code Civ. Proc. § 1060
(Against All Defendants)

18   245.   Plaintiffs and the Class hereby repeat and incorporate by reference each preceding

19 and succeeding paragraph as though fully set forth herein.

20   246.   An actual controversy exists between the Class (including Plaintiffs) and

21 Defendants due to Defendants' operation of Copilot, which violates Plaintiffs' and the Class's

22 rights, including but not limited to their rights under DMCA Section 1202, the Lanham Act, the

23 CCPA, and Cal. Civ. Code § 17200 as alleged herein.

24   247.   Plaintiffs are entitled to a judgment declaring that Defendants' actions are

25 unlawful and, specifically, that Defendants violated DMCA Section 1202, the Lanham Act, the

26 CCPA, and Cal. Civ. Code § 17200.

1

IX.   DEMAND FOR JUDGMENT

2   **WHEREFORE**, Plaintiffs requests that the Court enter judgment on their behalf and on

3   behalf of the Class defined herein, by adjudging and decreeing that:

4   248.   This action may proceed as a class action, with Plaintiffs serving as Class

5   Representatives, and with Plaintiffs' counsel as Class Counsel;

6   a) Judgment in favor of Plaintiffs and the Class and against Defendants;

7   b) Permanent injunctive relief, including but not limited to making changes to its

8   Copilot product to ensure that all applicable information set forth in 17 U.S.C. §

9   1203(b)(1) is included in along with any Output including associated code;

10   c) An order of costs and allowable attorney's fees pursuant to 17 U.S.C. § 1203(b)(4)–

11   (5);

12   d) An award of statutory damages pursuant to 17 U.S.C. § 1203(b)(3) and 17 U.S.C. §

13   1203(c)(3),[43] or, in the alternative, an award of actual damages and any additional

14   profits pursuant to 17 U.S.C. § 1203(c)(2) (including tripling damages pursuant to

15   17 U.S.C. § 1203(c)(4) if applicable);

16   e) An award of damages for harms resulting from Defendants' breach of Licenses;

17   f) An award of damages, including punitive damages, for harms resulting from

18   Defendants' tortious interference in Plaintiffs' and the Class's prospective

19   contractual relations;

20   g) An award of damages for harms resulting from Defendants' false designation of

21   the origin of Copilot's Output;

22   h) An award of damages in the amount Defendants have been unjustly enriched

23
24
25
26
27
28

---

[43] Plaintiffs estimate that statutory damages for Defendants' direct violations of DMCA Section 1202 alone will exceed $9,000,000,000. That figure represents minimum statutory damages ($2,500) incurred three times for each of the 1.2 million Copilot users Microsoft reported in June 2022. Each time Copilot provides an unlawful Output it violates Section 1202 three times (distributing the Licensed Materials without: (1) attribution, (2) copyright notice, and (3) License Terms). So, if each user receives just one Output that violates Section 1202 throughout their time using Copilot (up to fifteen months for the earliest adopters), then GitHub and OpenAI have violated the DMCA 3,600,000 times. At minimum statutory damages of $2500 per violation, that translates to $9,000,000,000.

1   through their conduct as alleged herein as well as punitive damages in connection

2   with this conduct;

3   i)   An award of damages, including punitive damages, for harms resulting from

4   Defendants acts of unfair competition;

5   j)   Statutory damages and any other relief this Court deems proper for Defendants

6   violation of the CCPA;

7   k)   An award of damages for harms resulting from GitHub's breach of the GitHub

8   Policies; and

9   l)   An award of damages, including punitive damages, for harms resulting from

10   Defendants' negligent handling of Plaintiffs' and the Class's personal data.

11   249.   Injunctive relief sufficient to alleviate and stop Defendants' unlawful conduct

12   alleged herein.

13   250.   Plaintiffs and the Class are entitled to prejudgment and post-judgment interest on

14   the damages awarded them, and that such interest be awarded at the highest legal rate from and

15   after the date this class action complaint is first served on Defendants;

16   251.   Defendants are to be jointly and severally responsible financially for the costs and

17   expenses of a Court approved notice program through post and media designed to give immediate

18   notification to the Class.

19   252.   Plaintiffs and the Class receive such other or further relief as may be just and

20   proper.

21   / / /

22   / / /

23   / / /

24

25

26

27

28

X.   JURY TRIAL DEMANDED

Pursuant to Federal Rule of Civil Procedure 38(b), Plaintiffs demand a trial by jury of all the claims asserted in this Complaint so triable.

Dated: November 10, 2022

By: _/s/ Joseph R. Saveri_
              Joseph R. Saveri

Joseph R. Saveri (State Bar No. 130064)
Steven N. Williams (State Bar No. 175489)
Cadio Zirpoli (State Bar No. 179108)
Elissa A. Buchanan (State Bar No. 249996)
Travis Manfredi (State Bar No. 281779)
JOSEPH SAVERI LAW FIRM, LLP
601 California Street, Suite 1000
San Francisco, California 94108
Telephone: (415) 500-6800
Facsimile: (415) 395-9940
Email: jsaveri@saverilawfirm.com
                    swilliams@saverilawfirm.com
                    czirpoli@saverilawfirm.com
                    eabuchanan@saverilawfirm.com
                    tmanfredi@saverilawfirm.com

Matthew Butterick (State Bar No. 250953)
1920 Hillhurst Avenue, #406
Los Angeles, CA 90027
Telephone: (323) 968-2632
Facsimile: (415) 395-9940
Email: mb@butticklaw.com

Counsel for Plaintiffs and the Proposed Class