

1 **IRELL & MANELLA LLP**
Morgan Chu (SBN 70446)
2 Alan Heinrich (SBN 212782)
3 Iian Jablon (SBN 205458)
4 Abigail Sellers (SBN 342380)
Justin Koo (SBN 351547)
Henry White (SBN 351549)
5 1800 Avenue of the Stars, Suite 900
Los Angeles, CA 90067
6 Telephone: (310) 277-1010
7 Fax: (310) 203-7199
mchu@irell.com
8 aheinrich@irell.com
9 ijablon@irell.com
asellers@irell.com
10 jkoo@irell.com
hwhite@irell.com

11 Attorneys for Plaintiff Elon Musk

12 **SUPERIOR COURT OF CALIFORNIA**
13 **CGC-24-612746**
14 **IN AND FOR THE COUNTY OF SAN FRANCISCO**

15 ELON MUSK, an individual,
16 Plaintiff,
17 vs.
18 SAMUEL ALTMAN, an individual, GREGORY
BROCKMAN, an individual, OPENAI, INC., a
19 corporation, OPENAI, L.P., a limited
partnership, OPENAI, L.L.C., a limited liability
20 company, OPENAI GP, L.L.C., a limited
liability company, OPENAI OPCO, LLC, a
21 limited liability company, OPENAI GLOBAL,
22 LLC, a limited liability company, OAI
CORPORATION, LLC, a limited liability
23 company, OPENAI HOLDINGS, LLC, a limited
liability company, and DOES 1 through 100,
24 inclusive,
25 Defendants.

Case No.:

[UNLIMITED JURISDICTION]

**COMPLAINT FOR (1) BREACH OF
CONTRACT, (2) PROMISSORY
ESTOPPEL, (3) BREACH OF FIDUCIARY
DUTY, (4) UNFAIR COMPETITION
UNDER CAL. BUS. & PROF. CODE
§§ 17200 ET SEQ., AND (5) ACCOUNTING
DEMAND FOR JURY TRIAL**

ELECTRONICALLY
FILED
Superior Court of California,
County of San Francisco
02/29/2024
Clerk of the Court
BY: KEVIN DOUGHERTY
Deputy Clerk

1 Plaintiff, ELON MUSK (hereafter “Plaintiff”) alleges the following upon information and
2 belief:

3 **PARTIES**

4 1. Prior to 2019, Plaintiff was an individual residing in California. Plaintiff is a resident
5 of Texas since 2019.

6 2. On information and belief, Plaintiff alleges that Samuel Altman is a resident of the
7 County of San Francisco, State of California.

8 3. On information and belief, Plaintiff alleges that Gregory Brockman is a resident of
9 the County of San Francisco, State of California.

10 4. OpenAI, Inc. is a registered non-profit organization incorporated under the laws of
11 Delaware on December 8, 2015. OpenAI, Inc. is registered as a foreign corporation with the
12 California Secretary of State and has its principal place of business at 3180 18th Street, San
13 Francisco, CA 94110.

14 5. OpenAI, L.P. is a limited partnership formed under the laws of Delaware on
15 September 19, 2018, originally as SummerSafe, L.P. OpenAI, L.P. is registered as a foreign limited
16 partnership with the California Secretary of State and has its principal place of business at 3180 18th
17 Street, San Francisco, CA 94110.

18 6. OpenAI, L.L.C. is a limited liability company formed in Delaware on September 17,
19 2020. OpenAI, L.L.C. maintains its principal place of business in California.

20 7. OpenAI GP, L.L.C. is a limited liability company formed in Delaware on September
21 19, 2018. OpenAI GP, L.L.C is registered as a foreign limited liability company registered with the
22 California Secretary of State and has its principal place of business at 3180 18th Street, San
23 Francisco, CA 94110.

24 8. OpenAI OpCo, LLC is a limited liability company formed in Delaware on September
25 19, 2018. OpenAI OpCo, LLC is registered as an out-of-state limited liability company with the
26 California Secretary of State and has it principal place of business at 1960 Bryant Street, San
27 Francisco, CA 94110.

28

1 9. OpenAI Global, LLC is a limited liability company formed in Delaware on
2 December 28, 2022. OpenAI Global, LLC is registered as an out-of-state limited liability company
3 with the California Secretary of State and has its principal place of business at 1960 Bryant Street,
4 San Francisco, CA 94110.

5 10. OAI Corporation, LLC is a limited liability company formed in Delaware. OAI
6 Corporation, LLC maintains its principal place of business in California.

7 11. OpenAI Holdings, LLC is a limited liability company formed in Delaware on March
8 17, 2023. OpenAI Holdings, LLC is registered as an out-of-state limited liability company with the
9 California Secretary of State and has its principal place of business at 1960 Bryant Street, San
10 Francisco, CA 94110.

11 12. Hereinafter, “OpenAI, Inc.” is used solely to refer to the non-profit entity or non-
12 profit arm, while “OpenAI” is used generally to refer to OpenAI, Inc., OpenAI, L.P., OpenAI,
13 L.L.C., OpenAI GP, L.L.C., OpenAI OpCo, LLC, OpenAI Global, LLC, OAI Corporation, LLC,
14 and/or OpenAI Holdings, LLC.

15 13. Plaintiff is currently unaware as to the names and identities of Doe 1 through Doe
16 100.

17 **JURISDICTION AND VENUE**

18 14. On information and belief, Plaintiff alleges that many of the occurrences,
19 representations, and events upon which this action is based took place in County of San Francisco,
20 State of California, where the vast majority of Defendants reside or have their principal place of
21 business, and the vast majority of the occurrences, representations, and events upon which this
22 action is based took place in the State of California.

23 **GENERAL ALLEGATIONS**

24 ***Overview Of The Case***

25 **A. *The Risk Of Artificial General Intelligence***

26 15. Over the course of the 20th century, the United States gradually shifted from a
27 primarily human labor-based economy to a primarily human knowledge-based economy, with
28 economic value increasingly created primarily by human intelligence. As the century progressed,

1 another paradigm shift was already underway: value creation through artificial intelligence (AI).
2 Early AI programs were capable of outperforming humans in certain discrete tasks. Almost
3 immediately after the invention of the programmable computer, AI could show super-human
4 performance on a highly formalized problem like finding the fastest path through a network of roads.
5 It took longer for AI to reach superiority for problems requiring more creativity. In 1996, IBM's
6 Deep Blue AI program beat Gary Kasparov, the then-world champion in chess. These programs,
7 while useful, were essentially one-trick-ponies—their intelligence was not general. Deep Blue had
8 massive arrays of processors that could only play chess. Algorithms for path finding could solve a
9 maze or route a car, but they could not paint a painting.

10 16. Starting in the late 2000s and early 2010s, an older algorithm called “deep learning”
11 became practical to implement on low-cost hardware for the first time. This caused an almost
12 overnight revolution in performance across nearly all AI projects. New, top of class algorithms were
13 developed for turning speech into text, translating between languages, and recognizing what kind of
14 food is shown in a photo. One of the hallmarks of deep learning is that algorithms do not need to be
15 designed with significant knowledge of the task at hand. They learn each task from training
16 examples, essentially programming themselves. This means that they are far more general-purpose
17 than earlier systems like Deep Blue.

18 17. As deep learning algorithms became increasingly sophisticated, some of the world's
19 leading AI researchers set their sights on what has come to be called Artificial General Intelligence
20 (AGI). The basic concept of AGI is a general purpose artificial intelligence system—a machine
21 having intelligence for a wide variety of tasks like a human.

22 18. Mr. Musk has long recognized that AGI poses a grave threat to humanity—perhaps
23 the greatest existential threat we face today. His concerns mirrored those raised before him by
24 luminaries like Stephen Hawking and Sun Microsystems founder Bill Joy. Our entire economy is
25 based around the fact that humans work together and come up with the best solutions to a hard task.
26 If a machine can solve nearly any task better than we can, that machine becomes more economically
27 useful than we are. As Mr. Joy warned, with strong AGI, “the future doesn't need us.” Mr. Musk
28

1 publicly called for a variety of measures to address the dangers of AGI, from voluntary moratoria
2 to regulation, but his calls largely fell on deaf ears.

3 19. But where some like Mr. Musk see an existential threat in AGI, others see AGI as a
4 source of profit and power.

5 20. In 2014, Google acquired DeepMind, a research group focused on deep learning.
6 One of DeepMind’s initial developments was AlphaZero, a chess playing algorithm. Unlike
7 previous algorithms, however, AlphaZero used “reinforcement learning,” wherein the program
8 learns to play chess by playing itself with different versions of the software. It starts by playing
9 randomly, with no understanding of the game’s strategy. When one version of the software wins a
10 game against another, the winning program’s internal pathways are “reinforced” and the process
11 repeats.

12 21. AlphaZero rapidly became the strongest chess playing system in the world.
13 Shockingly, it was also announced that the same program was also the strongest in the world for
14 playing two other extremely difficult games. In Google/DeepMind’s words, “Starting from random
15 play, and given no domain knowledge except the game rules, AlphaZero achieved within 24 hours
16 a superhuman level of play in the games of chess and shogi (Japanese chess) as well as Go, and
17 convincingly defeated a world-champion program in each case.”

18 22. With the DeepMind team, Google immediately catapulted to the front of the race for
19 AGI. Mr. Musk was deeply troubled by this development. He believed (and still does) that in the
20 hands of a closed, for-profit company like Google, AGI poses a particularly acute and noxious
21 danger to humanity. In 2014, it was already difficult enough to compete with Google in its core
22 businesses. Google had collected a uniquely large set of data from our searches, our emails, and
23 nearly every book in our libraries. Nevertheless, up to this point, everyone had the potential to
24 compete with Google through superior human intelligence and hard work. AGI would make
25 competition nearly impossible.

26 **B. The Founding Agreement Of OpenAI, Inc.**

27 23. Mr. Altman purported to share Mr. Musk’s concerns over the threat posed by AGI.
28 In 2015, Mr. Altman wrote that the “[d]evelopment of superhuman machine intelligence (SMI) is

1 probably the greatest threat to the continued existence of humanity. There are other threats that I
2 think are more certain to happen . . . but are unlikely to destroy every human in the universe in the
3 way that SMI could.” Later that same year, Mr. Altman approached Mr. Musk with a proposal: that
4 they join forces to form a non-profit AI lab that would try to catch up to Google in the race for AGI,
5 but it would be the opposite of Google.

6 24. Together with Mr. Brockman, the three agreed that this new lab: (a) would be a non-
7 profit developing AGI for the benefit of humanity, not for a for-profit company seeking to maximize
8 shareholder profits; and (b) would be open-source, balancing only countervailing safety
9 considerations, and would not keep its technology closed and secret for proprietary commercial
10 reasons (The “Founding Agreement”). Reflecting the Founding Agreement, Mr. Musk named this
11 new AI lab “OpenAI,” which would compete with, and serve as a vital counterbalance to,
12 Google/DeepMind in the race for AGI, but would do so to benefit humanity, not the shareholders
13 of a private, for-profit company (much less one of the largest technology companies in the world).

14 25. The Founding Agreement was also memorialized, among other places, in OpenAI,
15 Inc.’s December 8, 2015 Certificate of Incorporation, which affirmed that its “resulting technology
16 will benefit the public and the corporation will seek to open source technology for the public benefit
17 when applicable. The corporation is not organized for the private gain of any person.” Ex. 1 at 1.
18 The Certificate of Incorporation further affirmed that all of the corporation’s property was
19 “irrevocably dedicated” to these agreed purposes. *Id.*

20 26. In reliance on the Founding Agreement, which Mr. Altman, Mr. Brockman, and
21 OpenAI, Inc. reaffirmed with Mr. Musk on multiple occasions, Mr. Musk was a moving force
22 behind the creation of OpenAI, Inc. contributing a majority of its funding in its first several years,
23 advising on research directions, and most importantly, recruiting some of the world’s leading
24 scientists and engineers to work at the non-profit venture, including Chief Scientist Ilya Sutskever.
25 Recruiting for OpenAI, Inc. was a Herculean task in the face of relentless recruiting efforts by
26 Google/DeepMind, coupled with the lavish compensation Google/DeepMind offered. There would
27 have been no OpenAI, Inc. without Mr. Musk’s founding contributions and early leadership. Mr.
28

1 Musk continued to make contributions to OpenAI, Inc. from its founding through September 14,
2 2020.

3 27. OpenAI’s initial research was performed in the open, providing free and public
4 access to designs, models, and code. When OpenAI, Inc. researchers discovered that an algorithm
5 called “Transformers,” initially invented by Google, could perform many natural language tasks
6 without any explicit training, entire communities sprung up to enhance and extend the models
7 released by OpenAI, Inc. These communities spread to open-source, grass-roots efforts and
8 commercial entities alike.

9 28. Mr. Altman became OpenAI, Inc.’s CEO in 2019. On September 22, 2020, OpenAI
10 entered into an agreement with Microsoft, exclusively licensing to Microsoft its Generative Pre-
11 Trained Transformer (GPT)-3 language model. However, OpenAI published a detailed paper
12 describing the internals and training data for GPT-3, enabling the community to create similar
13 models themselves. And, most critically, the Microsoft license only applied to OpenAI’s pre-AGI
14 technology. Microsoft obtained no rights to AGI. And it was up to OpenAI, Inc.’s non-profit Board,
15 not Microsoft, to determine when OpenAI attained AGI.

16 C. **The 2023 Breach Of The Founding Agreement**

17 29. In 2023, Defendants Mr. Altman, Mr. Brockman, and OpenAI set the Founding
18 Agreement aflame.

19 30. In March 2023, OpenAI released its most powerful language model yet, GPT-4.
20 GPT-4 is not just capable of reasoning. It is better at reasoning than average humans. It scored in
21 the 90th percentile on the Uniform Bar Exam for lawyers. It scored in the 99th percentile on the
22 GRE Verbal Assessment. It even scored a 77% on the Advanced Sommelier examination. At this
23 time, Mr. Altman caused OpenAI to radically depart from its original mission and historical practice
24 of making its technology and knowledge available to the public. GPT-4’s internal design was kept
25 and remains a complete secret except to OpenAI—and, on information and belief, Microsoft. There
26 are no scientific publications describing the design of GPT-4. Instead, there are just press releases
27 bragging about performance. On information and belief, this secrecy is primarily driven by
28 commercial considerations, not safety. Although developed by OpenAI using contributions from

1 Plaintiff and others that were intended to benefit the public, GPT-4 is now a *de facto* Microsoft
2 proprietary algorithm, which it has integrated into its Office software suite.

3 31. Furthermore, on information and belief, GPT-4 is an AGI algorithm, and hence
4 expressly outside the scope of Microsoft’s September 2020 exclusive license with OpenAI. In this
5 regard, Microsoft’s own researchers have publicly stated that, “[g]iven the breadth and depth of
6 GPT-4’s capabilities, we believe that it could reasonably be viewed as an early (yet still incomplete)
7 version of an artificial general intelligence (AGI) system.” Moreover, on information and belief,
8 OpenAI is currently developing a model known as Q* (Q star) that has an even stronger claim to
9 AGI. As noted, Microsoft only has rights to certain of OpenAI’s pre-AGI technology. But for
10 purposes of the Microsoft license, it is up to OpenAI, Inc.’s Board to determine whether OpenAI
11 has attained AGI, and a Board coup took place in November 2023. On November 17, 2023, OpenAI,
12 Inc.’s Board fired Mr. Altman after losing “confidence in his ability to continue leading OpenAI”
13 because “he was not consistently candid with the board.” In a series of stunning developments
14 spanning the next several days, Mr. Altman and Mr. Brockman, in concert with Microsoft, exploited
15 Microsoft’s significant leverage over OpenAI, Inc. and forced the resignation of a majority of
16 OpenAI, Inc.’s Board members, including Chief Scientist Ilya Sutskever. Mr. Altman was reinstated
17 as CEO of OpenAI, Inc. on November 21. On information and belief, the new Board members were
18 hand-picked by Mr. Altman and blessed by Microsoft. The new Board members lack substantial AI
19 expertise and, on information and belief, are ill equipped by design to make an independent
20 determination of whether and when OpenAI has attained AGI—and hence when it has developed
21 an algorithm that is outside the scope of Microsoft’s license.

22 32. These events of 2023 constitute flagrant breaches of the Founding Agreement, which
23 Defendants have essentially turned on its head. To this day, OpenAI, Inc.’s website continues to
24 profess that its charter is to ensure that AGI “benefits all of humanity.” In reality, however, OpenAI,
25 Inc. has been transformed into a closed-source *de facto* subsidiary of the largest technology company
26 in the world: Microsoft.¹ Under its new Board, it is not just developing but is actually refining an
27 AGI to maximize profits for Microsoft, rather than for the benefit of humanity. Its technology,

28 ¹ Microsoft’s market cap as of February 28, 2024 was \$3.03 trillion.

1 including GPT-4, is closed-source primarily to serve the proprietary commercial interests of
2 Microsoft. Indeed, as the November 2023 drama was unfolding, Microsoft’s CEO boasted that it
3 would not matter “[i]f OpenAI disappeared tomorrow.” He explained that “[w]e have all the IP
4 rights and all the capability.” “We have the people, we have the compute, we have the data, we have
5 everything.” “We are below them, above them, around them.”

6 33. This case is filed to compel OpenAI to adhere to the Founding Agreement and return
7 to its mission to develop AGI for the benefit of humanity, not to personally benefit the individual
8 Defendants and the largest technology company in the world.

9 *Detailed Allegations*

10 **A. Mr. Musk’s Concerns Over AGI Falling Into The Wrong Hands**

11 34. In 2012, Elon Musk met Demis Hassabis, the co-founder of DeepMind, a for-profit
12 artificial intelligence company. On or about this time, Mr. Musk and Mr. Hassabis met at SpaceX’s
13 factory in Hawthorne, California where Mr. Musk and Mr. Hassabis discussed the greatest threats
14 facing society. During this conversation, Mr. Hassabis emphasized the potential dangers that the
15 advancement of AI presents to society.

16 35. Following this conversation with Mr. Hassabis, Mr. Musk became increasingly
17 concerned about the potential of AI to become super-intelligent, surpass human intelligence, and
18 threaten humanity. Indeed, Mr. Musk was not the only person afraid about the dangers of AI and
19 the research being conducted at DeepMind. It has been reported that following a meeting with Mr.
20 Hassabis and investors in DeepMind, one of the investors remarked that the best thing he could have
21 done for the human race was shoot Mr. Hassabis then and there.

22 36. Mr. Musk began discussing AI and DeepMind with those in his orbit, such as Larry
23 Page, then-CEO of Google’s parent company Alphabet, Inc. Mr. Musk would frequently raise the
24 dangers of AI in his conversations with Mr. Page, but to Mr. Musk’s shock, Mr. Page was
25 unconcerned. For example, in 2013, Mr. Musk had a passionate exchange with Mr. Page about the
26 dangers of AI. He warned that unless safeguards were put in place, “artificial intelligence-systems
27 might replace humans, making our species irrelevant or even extinct.” Mr. Page responded that
28 would merely “be the next stage of evolution,” and claimed Mr. Musk was being a “specist”—that

1 he favored the human species over intelligent machines. Mr. Musk responded, “Well, yes, I am pro-
2 human.”

3 37. At the end of 2013, Mr. Musk learned to his grave concern that Google was planning
4 to acquire DeepMind. At the time, DeepMind was one of the most advanced AI companies in the
5 industry. Thus, Mr. Musk was deeply concerned that DeepMind’s AI technology would be in the
6 hands of someone who viewed it and its power so cavalierly, and could hide its design and
7 capabilities behind closed doors.

8 38. In an effort to prevent this powerful technology from falling into Google’s hands,
9 Mr. Musk and Luke Nosek, a co-founder of PayPal, attempted to put together funding to buy
10 DeepMind. This effort culminated in an hour-long call wherein Mr. Musk and Mr. Nosek made one
11 last effort to convince Mr. Hassabis not to sell DeepMind to Google. Mr. Musk told Mr. Hassabis
12 that “[t]he future of AI should not be controlled by Larry [Page].”

13 39. Mr. Musk and Mr. Nosek’s effort was unsuccessful. It was reported in January 2014
14 that DeepMind would be acquired by Google. However, this did not deter Mr. Musk from continuing
15 to ensure that AI was developed and practiced safely.

16 40. Following Google’s acquisition of DeepMind, Mr. Musk began “hosting his own
17 series of dinner discussions on ways to counter Google and promote AI safety.” Mr. Musk also
18 reached out to President Barack Obama to discuss AI and AI safety. In 2015, Mr. Musk and
19 President Obama had a meeting during which Mr. Musk explained the dangers of AI and advocated
20 for regulation. Mr. Musk felt that President Obama understood the dangers of AI, but regulation
21 never came.

22 41. Despite these setbacks, Mr. Musk continued to advocate for safe AI practices. In
23 2015, it appeared that Mr. Musk may have found someone who understood his concerns about AI
24 and his desire to keep the first AGI out of the hands of a private company like Google: Defendant
25 Sam Altman.

26 42. At this time, Mr. Altman was the president of Y Combinator, a start-up accelerator
27 in Silicon Valley. Before that, Mr. Altman had been involved in various startup ventures.
28

1 43. Mr. Altman appeared to share Mr. Musk’s concerns surrounding AI. In public blog
2 posts dating back to 2014, Mr. Altman stated that AGI, if made, would “be the biggest development
3 in technology ever.” Mr. Altman pointed out that there are many companies making strides towards
4 achieving AGI, but acknowledged the unfortunate reality that the “good ones are very secretive
5 about it.”

6 44. On February 25, 2015, Mr. Altman also expressed his concern surrounding the
7 development of what he referred to as “superhuman machine intelligence” which he identified as
8 “probably the greatest threat to the continued existence of humanity” and emphasized that “as a
9 human programmed to survive and reproduce, I feel we should fight it.” Further, Mr. Altman
10 criticized those who believed that “superhuman machine intelligence” was dangerous but dismissed
11 it as “never going to happen or definitely very far off.” He accused them of engaging in “sloppy,
12 dangerous thinking.”

13 45. Indeed, in early March 2015, Mr. Altman extolled the importance of government
14 regulation as a means to ensure AI is created safely and suggested that “a group of very smart people
15 with a lot of resources” likely involving “US companies in some way” would be the most probable
16 group to achieve “superhuman machine intelligence” first.

17 46. Later that month, Mr. Altman reached out to Mr. Musk to inquire whether he would
18 be interested in drafting an open letter to the United States Government regarding AI. The two began
19 preparing a letter and approaching those of influence in the technology and AI sectors about signing.
20 It did not take long before those across the industry heard rumors of the letter.

21 47. For example, in April of 2015, Mr. Hassabis reached out to Mr. Musk stating that he
22 had heard from multiple sources that Mr. Musk was drafting a letter addressed to the President
23 calling for regulation of AI. Mr. Musk defended the idea of the regulation of AI to Mr. Hassabis,
24 stating: “If done well, it may very well accelerate AI in the long term. Without the public comfort
25 that regulatory oversight provides, there could very well be a situation where an AI causes great
26 harm and thereafter AI research is banned as dangerous to public safety.”

27 48. Five days after Mr. Hassabis reached out to Mr. Musk about the open letter regarding
28 AI, Mr. Hassabis announced the first meeting of the Google DeepMind AI Ethics Board, a board

1 that Google and DeepMind promised to establish two years earlier when Google acquired
2 DeepMind. Mr. Musk was asked to be a member of the board and offered for the first meeting to be
3 hosted at SpaceX in Hawthorne, California. It became clear to Mr. Musk after the first meeting that
4 this board was not a serious endeavor, but was rather a façade to try to slow down any regulation
5 of AI.

6 49. The Open Letter was later published on October 28, 2015, and signed by over eleven
7 thousand individuals, including Mr. Musk, Stephen Hawking, and Steve Wozniak.

8 **B. OpenAI, Inc.’s Founding Agreement**

9 50. On May 25, 2015, Mr. Altman emailed Mr. Musk, writing that he had “[b]een
10 thinking a lot about whether it’s possible to stop humanity from developing AI. I think the answer
11 is almost definitely not. If it’s going to happen, it seems like it would be good for someone other
12 than Google to do it first.” Mr. Altman had an idea: that Y Combinator start a “Manhattan Project”
13 for AI. (Unfortunately, it turns out that the “Manhattan Project” is a moniker that may be all too
14 apt.) He proposed that “we could structure it so that the tech belongs to the world via some sort of
15 nonprofit but the people working on it get startup-like compensation if it works. Obviously we’d
16 comply with/aggressively support all regulation.” Mr. Musk responded, “Probably worth a
17 conversation.”

18 51. After further communications, Mr. Altman emailed Mr. Musk on June 24, 2015 with
19 a detailed proposal for this new “AI lab.” “The mission would be to create the first general AI and
20 use it for individual empowerment—ie, the distributed version of the future that seems the safest.
21 More generally, safety should be a first-class requirement.” “The technology would be owned by
22 the foundation and used ‘for the good of the world’[.]” He proposed that they start with a group of
23 7-10 people and expand from there. He also proposed a governance structure. Mr. Musk responded,
24 “Agree on all.” Ex. 2 at 1.

25 52. Soon thereafter, Mr. Altman began recruiting others to help with the development of
26 the project. Notably, Mr. Altman approached Gregory Brockman to help with the project.

27 53. In November 2015, Mr. Altman put Mr. Brockman in communication with Mr. Musk
28 via email. Regarding the project, Mr. Brockman told Mr. Musk, “I hope for us to enter the field as

1 a neutral group looking to collaborate widely and shift the dialog towards being about humanity
2 winning rather than any particular group or company. (I think that’s the best way to bootstrap
3 ourselves into being a leading research institution.)” Optimistic about the possibility of a neutral AI
4 research group focused on humanity rather than profit for any particular individual or group,
5 Mr. Musk told Mr. Brockman that he would commit funding.

6 54. Mr. Musk came up with the name of the new lab, a name reflecting the Founding
7 Agreement: “Open AI Institute,” or simply, “OpenAI.”

8 55. With these principles of the Founding Agreement in mind, Mr. Musk joined forces
9 with Mr. Altman and Mr. Brockman to formalize and launch the project. Mr. Musk was actively
10 involved in the project, even prior to it being publicly announced. For example, Mr. Musk advised
11 Mr. Brockman on compensation packages for employees, sharing with Mr. Brockman his strategies
12 for compensation and retaining talent.

13 56. On December 8, 2015, a Certificate of Incorporation for OpenAI, Inc. was filed with
14 the Delaware Secretary of State. The Certificate memorialized in writing the Founding Agreement:

15 THIRD: This Corporation shall be a nonprofit corporation organized exclusively for
16 charitable and/or educational purposes within the meaning of section 501(c)(3) of
17 the Internal Revenue Code of 1986, as amended, or the corresponding provision of
18 any future United States Internal Revenue law. The specific purpose of this
19 corporation is to provide funding for research, development and distribution of
20 technology related to artificial intelligence. The resulting technology will benefit the
public and the corporation will seek to open source technology for the public benefit
when applicable. The corporation is not organized for the private gain of any
person. . . .

21 FIFTH: The property of this corporation is irrevocably dedicated to the purposes in
22 Article THREE hereof and no part of the net income or assets of this corporation
23 shall ever inure to the benefit of any director, officer or member thereof or to the
24 benefit of any private person. Upon the dissolution or winding up of this corporation,
25 its assets remaining after payment, or provision for payment, of all debts and
26 liabilities of this corporation shall be distributed to a nonprofit fund, foundation, or
27 corporation which is organized and operated exclusively for charitable, educational
28 and/or religious purposes and which has established its tax exempt status under
Internal Revenue Code section 501(c)(3), or the corresponding section of any future
federal tax code, or shall be distributed to the federal government, or to a state or
local government, for a public purpose.

1 57. OpenAI, Inc. was publicly announced on December 11, 2015. In the announcement,
2 Mr. Musk and Mr. Altman were named as co-chairs and Mr. Brockman was named as the CTO. The
3 announcement emphasized that OpenAI was designed to “benefit humanity,” and its research would
4 be “free from financial obligation”:

5 OpenAI is a non-profit artificial intelligence research company. Our goal is to
6 advance digital intelligence in the way that is most likely to benefit humanity as a
7 whole, unconstrained by a need to generate financial return. Since our research is
free from financial obligations, we can better focus on a positive human impact.

8 Ex. 3 at 1.

9 C. **Mr. Musk’s Crucial Role In Getting OpenAI, Inc. Off The Ground**

10 58. In an email on the day of the public announcement, Mr. Musk wrote that “[o]ur most
11 important consideration is recruitment of the best people.” He pledged that helping with the
12 recruiting effort would be his “absolute top priority 24/7.” He acknowledged that “[w]e are
13 outmanned and outgunned by a ridiculous margin by organizations you know well, but we have
14 right on our side and that counts for a lot. I like the odds.” Mr. Musk brought to bear his connections,
15 stature and clout in the effort. The fact that OpenAI, Inc. was an Elon Musk-sponsored initiative,
16 and that Mr. Musk served as co-chair, was instrumental to OpenAI, Inc.’s recruiting efforts,
17 particularly in the face of Google/DeepMind’s counter-recruiting. Without Mr. Musk’s involvement
18 and substantial supporting efforts and resources—his consideration for the Founding Agreement—
19 it is highly likely that OpenAI, Inc. would never have gotten off the ground.

20 59. One of the most important initial hires was for the role of Chief Scientist. Mr.
21 Altman, Mr. Brockman, and Mr. Musk all wanted Ilya Sutskever, then a research scientist at Google,
22 in that role. Dr. Sutskever went back and forth on whether to leave Google and join the project, but
23 it was ultimately a call from Mr. Musk on the day OpenAI, Inc. was publicly announced that
24 convinced Dr. Sutskever to commit to joining the project as OpenAI, Inc.’s Chief Scientist.

25 60. Mr. Musk actively recruited for OpenAI, Inc. over the next several months.
26 Google/DeepMind made increasingly lavish counter-offers to OpenAI, Inc. recruits in an attempt to
27 kill the new venture. In late February, Mr. Musk emailed Mr. Brockman and Mr. Altman, reiterating
28 that “[w]e need to do what it takes to get the top talent. Let’s go higher. If, at some point, we need

1 to revisit what existing people are getting paid, that's fine. Either we get the best people in the world
2 or we will get whipped by Deepmind. Whatever it takes to bring on ace talent is fin[e] by me.
3 Deepmind is causing me extreme mental stress. If they win, it will be really bad news with their one
4 mind to rule the world philosophy. They are obviously making major progress and well they should,
5 given the talent level over there.”

6 61. Mr. Musk was not just using his connections and clout to recruit on behalf of OpenAI,
7 Inc. When he told Mr. Brockman and Mr. Altman to go “higher” in their offers and “do what it takes
8 to get the top talent,” he was the one funding those higher offers. In 2016 alone, Mr. Musk
9 contributed over \$15 million dollars to OpenAI, Inc., more than any other donor. The funding he
10 provided enabled OpenAI, Inc. to assemble a team of top talent. Similarly, in 2017, Mr. Musk
11 contributed nearly \$20 million dollars to OpenAI, Inc., again, more than any other donor. All told,
12 Mr. Musk contributed more than \$44 million to OpenAI, Inc. between 2016 and September 2020.

13 62. In addition, through Musk Industries LLC, Mr. Musk leased OpenAI, Inc.’s initial
14 office space in the Pioneer Building in San Francisco and paid the monthly rental expenses. Mr.
15 Musk regularly visited OpenAI, Inc., and was present for important company milestones, such as
16 when the first DGX-1 AI supercomputer was donated to OpenAI, Inc. in 2016. Mr. Musk would
17 receive updates on OpenAI, Inc.’s progress and would provide his feedback and advice.

18 **D. Mr. Altman And Mr. Brockman Repeatedly Reaffirm The Founding Agreement**

19 63. In 2017, Mr. Brockman and others suggested transforming OpenAI, Inc. from a non-
20 profit to a for-profit corporation. After a series of communications over several weeks, Mr. Musk
21 told Mr. Brockman, Dr. Sutskever, and Mr. Altman “[e]ither go do something on your own or
22 continue with OpenAI as a nonprofit. I will no longer fund OpenAI until you have made a firm
23 commitment to stay or I’m just being a fool who is essentially providing free funding to a startup.
24 Discussions are over.”

25 64. In response, Mr. Altman told Mr. Musk “[I] remain enthusiastic about the non-profit
26 structure!” Ultimately, Mr. Brockman and Mr. Sutskever conveyed to Mr. Musk that they, too, were
27 resolved to continuing the non-profit structure, and that they would spend the next year committed
28 to fundraising to support the non-profit.

1 65. On February 21, 2018, Mr. Musk stepped down as a co-chair of OpenAI, Inc.
2 Nevertheless, Mr. Musk continued making contributions to OpenAI, Inc. relying on, and in
3 furtherance of, the Founding Agreement. For example, in 2018, Mr. Musk donated approximately
4 \$3.5 million to OpenAI, Inc. He also continued receiving updates about OpenAI, Inc. from Mr.
5 Brockman, Dr. Sutskever, and Mr. Altman.

6 66. In April 2018, Mr. Altman sent Mr. Musk a draft OpenAI Charter and solicited his
7 feedback. The draft charter described OpenAI’s mission as to ensure that AGI “benefits all of
8 humanity.” It stated, “We commit to use any influence we obtain over AGI’s deployment to ensure
9 it is used for the benefit of all, and to avoid enabling uses of AI or AGI that harm humanity or unduly
10 concentrate power. Our primary fiduciary duty is to humanity. We anticipate needing to marshal
11 substantial resources to fulfill our mission, but will always assiduously act to minimize conflicts of
12 interest . . . that could compromise broad benefit.”

13 67. On March 11, 2019, OpenAI, Inc. announced that it would be creating a for-profit
14 subsidiary: OpenAI, L.P. Prospective investors were notified of an “important warning” at the top
15 of the summary term sheet that the for-profit entity “exists to advance OpenAI Inc.’s [the non-
16 profit’s] mission of ensuring that safe artificial general intelligence is developed and benefits all of
17 humanity. The General Partner’s duty to this mission and the principles advanced in the OpenAI
18 Inc. Charter take precedence over any obligation to generate a profit.” Accordingly, investors were
19 expressly advised that “[i]t would be wise to view any investment in OpenAI LP in the spirit of a
20 donation.”

21 68. Following the announcement, Mr. Musk reached out to Mr. Altman asking him to
22 “be explicit that I have no financial interest in the for-profit arm of OpenAI.” However, Mr. Musk
23 continued to support OpenAI, Inc., the non-profit, donating an additional \$3.48 million in 2019.

24 69. On September 22, 2020, OpenAI announced that it exclusively licensed certain of its
25 pre-AGI technology to Microsoft. Consistent with the Founding Agreement, OpenAI’s website
26 states that AGI, which it describes as “a highly autonomous system that outperforms humans at most
27 economically valuable work” “is excluded from IP licenses and other commercial terms with
28

1 Microsoft, which only apply to pre-AGI technology.” However, OpenAI’s Board “determines when
2 we’ve attained AGI.”

3 **E. OpenAI’s Shifting Corporate Structure**

4 70. In the years following the announcement of the OpenAI, L.P., OpenAI’s corporate
5 structure became increasingly complex. OpenAI, L.L.C. was formed in Delaware on September 17,
6 2020. OpenAI, L.L.C.’s sole member is OpenAI OpCo, LLC.

7 71. OpenAI OpCo, LLC was formed in Delaware on September 19, 2018, and its sole
8 member is OpenAI Global, LLC.

9 72. OpenAI Global, LLC was formed in Delaware on December 28, 2022. On
10 information and belief, OpenAI Global, LLC, like OpenAI, L.P., is a capped for-profit entity.
11 OpenAI Global, LLC has two members: Microsoft Corporation and OAI Corporation, LLC.

12 73. OAI Corporation, LLC is a limited liability company formed in Delaware. The sole
13 member of OAI Corporation, LLC is OpenAI Holdings, LLC.

14 74. OpenAI Holdings, LLC was formed in Delaware on March 17, 2023, and has
15 multiple members including OpenAI, Inc., Aestas, LLC, and various individual members.

16 75. On information and belief, OpenAI, Inc. manages OpenAI, L.P. and OpenAI Global,
17 LLC through its general partner, OpenAI GP, L.L.C., which was registered as a limited liability
18 company in Delaware on September 19, 2018. OpenAI GP, L.L.C. is wholly owned by OpenAI,
19 Inc. (the non-profit), and is controlled by the non-profit’s Board of Directors. Under OpenAI’s
20 charter, the Board has no fiduciary duties to shareholders; rather, its sole fiduciary duty is to
21 humanity.

22 76. On information and belief, at least OpenAI, L.P. and OpenAI GP LLC were initially
23 established to facilitate and fund OpenAI, Inc.’s non-profit mission as set forth in the Founding
24 Agreement.

25 **F. The Development Of OpenAI’s Technology—From AI To AGI**

26 77. Initial work at OpenAI followed much in the footsteps of DeepMind. OpenAI used
27 reinforcement learning to play a game. Instead of playing chess, however, OpenAI competed in
28 Dota 2, a strategy video game with far more moving pieces than chess. OpenAI’s team quickly

1 constructed a new model that beat the reigning world champion team, demonstrating “that self-play
2 reinforcement learning can achieve superhuman performance on a difficult task.”

3 78. Meanwhile, at Google, an algorithm called the Transformer was created that solved
4 many of the issues deep learning had faced in understanding long sequences of text. This algorithm,
5 an example of a “large-language model,” was developed to translate text from one language to
6 another and works by forming connections between words in the source language and mapping
7 those connections to the target language.

8 79. Researchers at OpenAI continued this research and quickly produced yet another
9 startling result: by using the first half of Google’s Transformer architecture, a deep neural network
10 could be pre-trained on a large corpus of text and used to generate new text. In January 2018,
11 OpenAI released the source code and trained model for this Generative Pre-Trained Transformer
12 (GPT) along with a detailed paper describing the model and its capabilities

13 80. In 2019, OpenAI released a second-generation model, GPT-2. The release was again
14 accompanied by a detailed paper describing the model and noting that unlike previous models,
15 which needed to be trained on a specific task, “[w]hen a large language model is trained on a
16 sufficiently large and diverse dataset it is able to perform well across many domains and datasets.”
17 These models were proving themselves to be very different from previous AI systems. Instead of
18 training the system to perform a particular task, they could be simply “asked” to perform a new task
19 in natural language.

20 81. As contemplated by the Founding Agreement, OpenAI publicly released the full
21 version of GPT-2. Notably, it did so despite the fact that “humans find GPT-2 outputs convincing,”
22 that “GPT-2 can be fine-tuned for misuse,” and that “detection [of GPT generated text] is
23 challenging.” At the time, OpenAI stated that it was releasing the full, open version with the hope
24 that it “will be useful to developers of future powerful models.” This release was accompanied by a
25 detailed paper co-authored by OpenAI scientists as well as independent social and technical
26 scientists. This paper explained just some of the many benefits that came from releasing models
27 publically as opposed to keeping them closed.

28

1 82. Their publication did prove to be useful to the developers of future, powerful models.
2 Entire communities sprung up to enhance and extend the models released by OpenAI. These
3 communities spread to open-source, grass-roots efforts and commercial entities alike.

4 83. In 2020, OpenAI announced a third version of its model, GPT-3. It used “175 billion
5 parameters, 10x more than any previous non-sparse language model.” Again, OpenAI announced
6 the development of this model with the publication of a research paper describing its complete
7 implementation for others to build on.

8 84. In 2022, researchers at Google took these results and showed that a small change
9 called chain-of-thought prompting could enable “large language models to perform complex
10 reasoning.” Researchers at the University of Tokyo and Google quickly expanded on Google’s
11 results and showed that OpenAI’s GPT-3 could reason about an entirely new problem in a step-by-
12 step manner, like a human, “by simply adding ‘Let’s think step by step’ before each answer.”

13 85. A path to Artificial Generative Intelligence could be seen. And the timeline to it was
14 dramatically compressing.

15 86. On March 14, 2023, OpenAI released a new generation of its model, GPT-4. This
16 generation was not just capable of reasoning *but was better at reasoning than average humans*.
17 GPT-4 scored in the 90th percentile on the Uniform Bar Exam. It scored in the 99th percentile on
18 the GRE Verbal Assessment. It even scored a 77% on the Advanced Sommelier examination. By
19 OpenAI’s own objective measures, GPT-4 is already capable of intelligence that is superior to
20 humans on a wide variety of economically valuable tasks.

21 87. This development was not lost on the research community. In a detailed analysis
22 titled “Sparks of Artificial General Intelligence: Early experiments with GPT-4,” Microsoft
23 researchers note that “GPT-4 can solve novel and difficult tasks that span mathematics, coding,
24 vision, medicine, law, psychology and more, without needing any special prompting. Moreover, in
25 all of these tasks, GPT-4’s performance is strikingly close to human-level performance, and often
26 vastly surpasses prior models such as [GPT-3.5 based] ChatGPT.”

27 88. They compared the performance of GPT-4 with the performance of GPT-3 based
28 systems and found that “there is no comparison with the outputs from GPT-4.” On mathematical

1 problems, they show that “[t]he solution given by GPT-4 is correct and the argument is sound, while
2 ChatGPT [based on GPT-3] produces an incorrect solution which (in the case of a human) would
3 reflect a lack of understanding of the concept of function inversion.” On another example, they show
4 that “GPT-4 gives a correct solution while ChatGPT begins by rearranging the terms without any
5 clear direction or purpose, and ends up with an incorrect solution.”

6 89. Microsoft’s own scientists acknowledge that GPT-4 “attains a form of general
7 intelligence” and that “[g]iven the breadth and depth of GPT-4’s capabilities, we believe that it could
8 reasonably be viewed as an early (yet still incomplete) version of an artificial general intelligence
9 (AGI) system.”

10 **G. The Founding Agreement Is Breached In 2023**

11 90. Having reached the threshold of AGI, which under the Founding Agreement they
12 were to develop for the benefit of humanity rather than for any for-profit company or personal profit,
13 Defendants instead radically departed from their mission in breach of the Founding Agreement.
14 GPT-4 is an entirely closed model. The internal design of GPT-4 remains a secret and no code has
15 been released. OpenAI has not published a paper describing any aspect of its internal design; it has
16 simply issued press releases boasting about its performance. The internal details of GPT-4 are
17 known only to OpenAI and, on information and belief, to Microsoft. GPT-4 is hence the opposite
18 of “open AI.” And it is closed for propriety commercial reasons: Microsoft stands to make a fortune
19 selling GPT-4 to the public, which would not be possible if OpenAI—as it is required to do—makes
20 the technology freely available to the public. Contrary to the Founding Agreement, Defendants have
21 chosen to use GPT-4 not for the benefit of humanity, but as proprietary technology to maximize
22 profits for literally the largest company in the world. Further, OpenAI’s entire development is now
23 veiled in secrecy and the public only has rumors and isolated fragments of communications to
24 understand what may be released next.

25 91. Researchers have pointed out that one of the remaining limitations of GPT
26 architecture-based AIs is that they generate their output a piece at a time and cannot “backtrack.”
27 These issues have been seen before in artificial intelligence research and have been largely solved
28 for other applications. In path and maze finding, AI must be able to find the right path despite the

1 existence of dead-ends along the way. The standard algorithm to perform this is called “A*”
2 (pronounced A-star).

3 92. Reuters has reported that OpenAI is developing a secretive algorithm called Q*.
4 While it is not clear what Q* is, Reuters has reported that several OpenAI staff members wrote a
5 letter warning about the potential power of Q*. It appears Q* may now or in the future be a part of
6 an even clearer and more striking example of artificial general intelligence that has been developed
7 by OpenAI. As an AGI, it would be explicitly outside the scope of OpenAI’s license with Microsoft,
8 and must be made available for the benefit of the public at large.

9 93. For purposes of the license with Microsoft, OpenAI, Inc.’s Board determines whether
10 OpenAI has attained AGI, and in a series of shocking developments described in more detail below,
11 a majority of OpenAI, Inc.’s Board was forced to resign on November 22, 2023, and their
12 replacements were, on information and belief, handpicked by Mr. Altman and Microsoft.

13 94. On November 17, 2023, OpenAI, Inc.’s Board dismissed Mr. Altman. OpenAI
14 announced in a blog post that Mr. Altman had been fired and his “departure follows a deliberative
15 review process by the board, which concluded that he was not consistently candid in his
16 communications with the board, hindering its ability to exercise its responsibilities. The board no
17 longer has confidence in his ability to continue leading OpenAI.”

18 95. Mr. Brockman was also removed from the Board but was told he would retain his
19 role at OpenAI.

20 96. At the time, the Board consisted of Helen Toner, Adam D’Angelo, Tasha McCauley,
21 Dr. Sutskever, Mr. Brockman, and Mr. Altman. In addition to serving on the Board, Ms. Toner is a
22 researcher and advisor for the Center for the Governance of AI (GovAI) and the Director of Strategy
23 at Georgetown’s Center for Security and Emerging Technology. Ms. McCauley is a Senior
24 Management Scientist at RAND Corporation, a nonprofit which specializes in public policy
25 decision making. She is also an advisor for GovAI. Mr. D’Angelo—the only board member to
26 remain after Mr. Altman’s reinstatement—is a tech CEO and entrepreneur.

27 97. The choice to include on OpenAI, Inc.’s Board multiple academics and public policy
28 experts with deep AI policy experience, most of whom had no financial stake in the company, was

1 deliberate. This composition of financially disinterested Board members with strong records of
2 public service ensured that the Board would put the nonprofit's principal beneficiary—humanity—
3 before financial success. This safeguard was put in place for the Board in furtherance of OpenAI,
4 Inc.'s non-profit mission and the Founding Agreement: to safely create AGI which would benefit
5 humanity, not the pecuniary interests of a for-profit company.

6 98. On information and belief, Mr. Altman's firing was due in part to OpenAI's
7 breakthrough in realizing AGI. In fact, news reports suggested that there was a rift among OpenAI
8 Board members and executives regarding safety concerns and the potential threat posed by
9 OpenAI's next generation Q*.

10 99. News of Mr. Altman's ouster spread rapidly. Following the Board's announcement
11 that Mr. Altman was fired, Mr. Brockman announced he would be leaving OpenAI with Mr. Altman.

12 100. When Microsoft CEO Satya Nadella learned of Mr. Altman's firing, he was
13 reportedly furious. As a 49% shareholder in OpenAI's for-profit arm, Mr. Nadella felt that Microsoft
14 should have been consulted before Mr. Altman's firing. However, at this time, outside of Mr.
15 Altman, OpenAI, Inc.'s Board had no ties to Microsoft, and no fiduciary duty to investors in the for-
16 profit arm. Rather, on information and belief, Mr. Altman was the primary liaison between
17 Microsoft and OpenAI, Inc.

18 101. Mr. Nadella invited Mr. Altman and Mr. Brockman to lead a new Microsoft AI
19 research lab, unbound by the constraints of OpenAI, Inc.'s humanitarian mission. Nadella made it
20 clear that employees who left OpenAI would be welcome at Microsoft's new lab with the same
21 salary.

22 102. Microsoft was confident that, through its substantial ownership in OpenAI's for-
23 profit arm, it could completely sequester OpenAI, Inc.'s research should the company cease to exist.
24 Indeed, during an interview shortly after Mr. Altman's firing, Mr. Nadella stated:

25 [W]e were very confident in our own ability. We have all the IP rights and all the
26 capability. If OpenAI disappeared tomorrow, I don't want any customer of ours to
27 be worried about it quite honestly, because we have all of the rights to continue the
28 innovation. Not just to serve the product, but we can go and just do what we were
doing in partnership ourselves. We have the people, we have the compute, we have
the data, we have everything.

1 103. Despite Microsoft’s statements regarding their capability to function without
2 OpenAI, Microsoft never abandoned its plan to secure the reinstatement of Mr. Altman as OpenAI,
3 Inc.’s CEO. In the days following Mr. Altman’s firing, OpenAI, Inc.’s Board was faced with
4 mounting pressure from lawyers and major shareholders, including Microsoft, to reinstate Mr.
5 Altman.

6 104. Ms. Toner was specifically targeted in efforts to reinstate Mr. Altman. During these
7 efforts, a lawyer for OpenAI told Ms. Toner that she and the Board could face claims for breach of
8 their fiduciary duties to investors should OpenAI fail due to Mr. Altman’s firing.

9 105. However, the OpenAI, Inc. Board has never had a fiduciary duty to investors. In fact,
10 all investors in the for-profit arm are told that the company’s duty to its mission takes precedence
11 over its duty to its investors, and OpenAI, Inc.’s website clearly indicates it only holds a fiduciary
12 duty to humanity.

13 106. Ms. Toner, who described the lawyer’s actions as an intimidation tactic, opined that
14 Mr. Altman’s continued removal would in fact advance the company’s mission by promoting the
15 safety of humanity over profit. However, none of this dissuaded the shareholders and Mr. Altman
16 from pushing for his reinstatement.

17 107. Further, Microsoft had substantial coercive power over OpenAI, Inc. and its Board.
18 During the period of Mr. Altman’s firing, CEO Nadella said the following of Microsoft’s
19 relationship with OpenAI:

20 We are in there. We are below them, above them, around them. We do the kernel
21 optimizations, we build tools, we build the infrastructure. So that’s why I think a lot of
22 the industrial analysts are saying, “Oh wow, it’s really a joint project between Microsoft
and OpenAI.” The reality is we are, as I said, very self-sufficient in all of this.

23 108. Moreover, at the time of the firing, Microsoft had only paid a fraction of a \$10 billion
24 investment commitment it had made to OpenAI, giving Microsoft serious leverage over the
25 “independent” non-profit Board. Additionally, if Microsoft withheld its cloud computing system on
26 which OpenAI was reliant, the company would be incapacitated.

27 109. The pressure on the Board from Mr. Altman and Microsoft continued until
28 November 21, when Mr. Altman was reinstated as the CEO of OpenAI, Inc. and Mr. Brockman re-

1 joined as President. One of Mr. Altman’s conditions upon return was that Ms. Toner, Ms. McCauley,
2 and Dr. Sutskever resign from the Board.

3 110. On information and belief, following his return, Mr. Altman hand-picked a new
4 Board that lacks similar technical expertise or any substantial background in AI governance, which
5 the previous board had by design. Mr. D’Angelo, a tech CEO and entrepreneur, was the only
6 member of the previous board to remain after Mr. Altman’s return. The new Board consisted of
7 members with more experience in profit-centric enterprises or politics than in AI ethics and
8 governance. They were also reportedly “big fans of Altman.” Two of the new board members were
9 Bret Taylor and Larry Summers. Mr. Taylor is no stranger to Silicon Valley, and has been heavily
10 involved in various Bay Area profit-driven ventures. On February 14, 2024, Mr. Taylor and ex-
11 Google executive Clay Bavor launched a startup focused on building AI chatbots for businesses.
12 Dr. Summers is an economist who, on information and belief, had no experience working in artificial
13 intelligence-based ventures prior to November 2023. Microsoft also obtained an observer seat on
14 the Board from which it could keep a close eye on its ostensibly non-profit golden goose.

15 111. With the reinstatement of Mr. Altman and the restructuring of the Board, OpenAI’s
16 corporate structure that had been designed as a system of checks and balances between the non-
17 profit arm, for-profit arm, the Board, and the CEO to ensure the non-profit mission was being carried
18 out, collapsed overnight. OpenAI, Inc.’s once carefully crafted non-profit structure was replaced by
19 a purely profit-driven CEO and a Board with inferior technical expertise in AGI and AI public
20 policy. The Board now has an observer seat reserved solely for Microsoft. With this restructuring,
21 OpenAI, Inc. abandoned its non-profit mission of developing AGI for the benefit of humanity
22 broadly, thereby keeping it out of the hands of a large for-profit corporation in which vast power
23 would be unduly concentrated.

24 112. The OpenAI, Inc. Board’s technical expertise in AI, neutrality, and commitment to
25 OpenAI, Inc.’s nonprofit mission are particularly essential to OpenAI, Inc.’s mission as it is the
26 Board that determines whether OpenAI has attained AGI for purposes of its license agreement with
27 Microsoft. This means that the Board is tasked with determining whether OpenAI’s most powerful
28 and advanced technology is actually excluded from the scope of Microsoft’s exclusive license.

1 Given Microsoft’s enormous financial interest in keeping the gate closed to the public, OpenAI,
2 Inc.’s new captured, conflicted, and compliant Board will have every reason to delay ever making
3 a finding that OpenAI has attained AGI. To the contrary, OpenAI’s attainment of AGI, like
4 “Tomorrow” in *Annie*, will always be a day away, ensuring that Microsoft will be licensed to
5 OpenAI’s latest technology and the public will be shut out, precisely the opposite of the Founding
6 Agreement.

7 113. OpenAI’s conduct could have seismic implications for Silicon Valley and, if allowed
8 to stand, could represent a paradigm shift for technology start-ups. It is important to reflect on what
9 has transpired here: a non-profit startup has collected tens of millions of dollars in contributions for
10 the express purpose of developing AGI technology for public benefit, and shortly before achieving
11 the very milestone that the company was created to achieve, the company has become a closed, for-
12 profit partner of the world’s largest corporation, thereby personally enriching the Defendants. If this
13 business model were valid, it would radically redefine how venture capitalism is practiced in
14 California and beyond. Rather than start out as a for-profit entity from the outset, “smart” investors
15 would establish non-profits, use pre-tax donations to fund research and development, and then once
16 their technology had been developed and proven, would slide the resulting IP assets into a new for-
17 profit venture to enrich themselves and their profit-maximizing corporate partners. That is not
18 supposed to be how the law works in California or anywhere else in this country, and this should
19 not be the first Court to hold otherwise.

20 114. To further understand why this is important, if OpenAI’s new business model is
21 valid, for every dollar that an investor “invests” by contributing to a non-profit, that investor gets
22 approximately 50 cents back from the state and federal governments in the form of reduced income
23 taxes, so the net cost to them of each \$1 of investment is only 50 cents. However, with OpenAI’s
24 new business model, they get the same “for profit” upside as those who invest the conventional way
25 in for-profit corporations and thus do not get an immediate tax write off, financed by the government
26 and, ultimately, the public. From an investment perspective, competing against an entity employing
27 the new OpenAI business model would be like playing a game of basketball where the other team’s
28 baskets are worth twice as many points. If this Court validates OpenAI’s conduct here, any start-up

1 seeking to remain competitive in Silicon Valley would essentially be required to follow this OpenAI
2 playbook, which would become standard operating procedure for start-ups to the detriment of
3 legitimate non-profits, the government's tax coffers, and ultimately the people of California and
4 beyond. Notably, OpenAI's for-profit arm was recently valued at nearly \$80 billion.

5 **H. November 2023 To Present: Altman's OpenAI**

6 115. The public is still in the dark regarding what exactly the Board's "deliberative review
7 process" revealed that resulted in the initial firing of Mr. Altman. However, one thing is clear to Mr.
8 Musk and the public at large: OpenAI has abandoned its "irrevocable" non-profit mission in the
9 pursuit of profit. Numerous leaders and intellectuals have publicly commented on the irony and
10 tragedy of OpenAI becoming "Closed, For-Profit AI."

11 116. For example, on November 29, 2023, MIT economists wrote an opinion in the L.A.
12 Times expressing concern with OpenAI's new profit-driven directive. In their words, "[d]isruption
13 and uncontrolled growth have become religion for the tech industry, and Altman has been one of its
14 most dedicated high priests." The economists emphasized the new Board is far more likely to allow
15 Mr. Altman to scale up OpenAI as quickly as possible, no matter the severity of societal costs.

16 117. The president of Public Citizen, a non-profit consumer advocacy organization, wrote
17 an open letter to California Attorney General Rob Bonta earlier this year raising concerns about
18 whether OpenAI's for-profit subsidiary is exerting inappropriate control over the non-profit, or
19 whether the non-profit's purpose had shifted to profit-making under Mr. Altman and Microsoft. The
20 letter suggested that should the non-profit abandon its original mission, it should be dissolved with
21 proceeds going to another charitable enterprise.

22 118. A WIRED investigation from January 2024 found that OpenAI has also recently
23 shuttered the public's access to previously available "key documents." In line with OpenAI's
24 original promise of transparency, OpenAI's IRS documents have indicated since its founding that
25 any member of the public may view copies of its governing documents, financial statements, and
26 conflict of interest rules. However, when WIRED requested these documents, OpenAI indicated it
27 had changed its policy. Thus, while OpenAI has long touted its commitment to transparency,
28

1 members of the public are unable to obtain information shedding light on the incidents of November
2 2023.

3 119. Access to OpenAI’s documents could inform the public whether it has changed its
4 governance to appease Microsoft and other shareholders. At the very least, changes had to be made
5 to accommodate Microsoft’s seat on the Board, and now Mr. Altman is reportedly in discussions
6 with Middle Eastern investors to raise up to \$7 trillion in an effort develop a global network of AI
7 chip fabrication plants. If \$10 billion from Microsoft was enough to get it a seat on the Board, one
8 can only imagine how much influence over OpenAI, Inc. these new potential investments could
9 confer on the investors. This is especially troubling when one potential donor is the national security
10 advisor of the United Arab Emirates, and United States officials are concerned due to the United
11 Arab Emirates’ ties to China. Moreover, Mr. Altman has been quoted discussing the possibility of
12 making the United Arab Emirates a “regulatory sandbox” where AI technologies are tested.

13 120. Further, access to OpenAI’s conflict-of-interest policy is important to shed light on
14 the Board’s ability to control Mr. Altman’s use of OpenAI to advance his own economic interests,
15 which so far appear to have gone unchecked. For example, in 2019 when Mr. Altman was CEO,
16 OpenAI signed a letter of intent to buy \$51 million worth of chips from a start-up in which Mr.
17 Altman was heavily invested.

18 121. While OpenAI, Inc. was once a pioneer of safe, responsible development of AGI
19 based on open communication with the public, it has now shut its doors, brought the biggest investor
20 of its for-profit subsidiary onto a Board which owes its only fiduciary duty to humanity, and
21 continues in secrecy towards a profit-centric future with possible calamitous implications for
22 humanity.

23 122. Mr. Musk founded and funded OpenAI, Inc. with Mr. Altman and Mr. Brockman in
24 exchange for and relying on the Founding Agreement to ensure that AGI would benefit humanity,
25 not for-profit corporations. As events turned out in 2023, his contributions to OpenAI, Inc. have
26 been twisted to benefit the Defendants and the biggest company in the world. This was a stark
27 betrayal of the Founding Agreement, turning that Agreement on its head and perverting OpenAI,
28 Inc.’s mission. Imagine donating to a non-profit whose asserted mission is to protect the Amazon

1 rainforest, but then the non-profit creates a for-profit Amazonian logging company that uses the
2 fruits of the donations to clear the rainforest. That is the story of OpenAI, Inc.

3 **FIRST CAUSE OF ACTION**

4 **Breach of Contract**

5 ***Against All Defendants***

6 123. Plaintiff realleges and incorporates by reference only paragraphs of this Complaint
7 necessary for his claim of Breach of Contract.

8 124. From OpenAI, Inc.'s founding in 2015 through September 2020, Plaintiff
9 contributed tens of millions of dollars, provided integral advice on research directions, and played
10 a key role in recruiting world-class talent to OpenAI, Inc. in exchange and as consideration for the
11 Founding Agreement, namely, that: OpenAI, Inc. (a) would be a non-profit developing AGI for the
12 benefit of humanity, not for a for-profit company seeking to maximize shareholder profits; and
13 (b) would be open-source, balancing only countervailing safety considerations, and would not keep
14 its technology closed and secret for proprietary commercial reasons. This Founding Agreement is
15 memorialized in, among other places, OpenAI, Inc.'s founding Articles of Incorporation and in
16 numerous written communications between Plaintiff and Defendants over a multi-year period,
17 such as:

18 a. "The specific purpose of this corporation is to provide funding for research,
19 development and distribution of technology related to artificial intelligence. The
20 resulting technology will benefit the public and the corporation will seek to open
21 source technology for the public benefit when applicable. The corporation is not
22 organized for the private gain of any person." Ex. 1 at 1.

23 b. Mr. Altman stated: "The mission would be to create the first general AI and use it
24 for individual empowerment—ie, the distributed version of the future that seems the
25 safest. More generally, safety should be a first-class requirement. . . . The technology
26 would be owned by the foundation and used 'for the good of the world'[".]” Plaintiff
27 replied: "Agree on all." Ex. 2 at 1.

28

1 125. Defendants have breached the Founding Agreement in multiple separate and
2 independent ways, including at least by:

- 3 a. Licensing GPT-4, which Microsoft’s own scientists have written can “reasonably
4 be viewed as an early (yet still incomplete) version of an artificial general
5 intelligence (AGI) system,” exclusively to Microsoft, despite agreeing that OpenAI
6 would develop AGI for the benefit of humanity, not for the private commercial gain
7 of a for-profit company seeking to maximize shareholder profits, much less the
8 largest corporation in the world.
- 9 b. Failing to disclose to the public, among other things, details on GPT-4’s
10 architecture, hardware, training method, and training computation, and further by
11 erecting a “paywall” between the public and GPT-4, requiring per-token payment
12 for usage, in order to advance Defendants and Microsoft’s own private commercial
13 interests, despite agreeing that OpenAI’s technology would be open-source,
14 balancing only countervailing safety considerations.
- 15 c. Permitting Microsoft, a publicly traded for-profit corporation, to occupy a seat on
16 OpenAI, Inc.’s Board of Directors and exert undue influence and control over
17 OpenAI’s non-profit activities including, for example, the determination of whether
18 and to what extent to make OpenAI, Inc.’s technology freely available and open to
19 the public and the determination of whether OpenAI has attained AGI.

20 126. As a direct and proximate result of Defendants’ breaches, Plaintiff has suffered
21 damages in an amount that is presently unknown, but that substantially exceeds this Court’s
22 jurisdictional minimum of \$35,000, and, if necessary, will be proven at trial.

23 127. Plaintiff also seeks and is entitled to specific performance of Defendant’s contractual
24 obligations.

1 SECOND CAUSE OF ACTION

2 **Promissory Estoppel**

3 *Against All Defendants*

4 128. Plaintiff realleges and incorporates by reference only paragraphs of this Complaint
5 necessary for his claim of Promissory Estoppel.

6 129. In order to induce Plaintiff to make millions of dollars in contributions to OpenAI,
7 Inc. over a period of years, and to induce him to provide substantial time and other resources to get
8 OpenAI, Inc. off the ground as alleged herein, Defendants repeatedly promised Plaintiff, including
9 in writing, that OpenAI (a) would be a non-profit developing AGI for the benefit of humanity, not
10 for a for-profit company seeking to maximize shareholder profits; and (b) would be open-source,
11 balancing only countervailing safety considerations, and would not keep its technology closed and
12 secret for proprietary commercial reasons.

13 130. In doing so, Defendants reasonably expected that Plaintiff would (as he did) rely on
14 their promises and provide funding, time and other resources to OpenAI, Inc.

15 131. Plaintiff reasonably relied on Defendants' false promises to his detriment, ultimately
16 providing tens of millions of dollars of funding to OpenAI, Inc., as well as his time and other
17 resources, on the condition that OpenAI would remain a non-profit irrevocably dedicated to creating
18 safe, open-source AGI for public benefit, only to then have OpenAI abandon its "irrevocable" non-
19 profit mission, stop providing basic information to the public, and instead exclusively dedicate and
20 license its AGI algorithms to the largest for-profit company in the world, precisely the opposite of
21 the promises Defendants made to Plaintiff.

22 132. Injustice can only be avoided through the enforcement of Defendants' repeated
23 promises. If specific enforcement is not awarded, then Defendants must at minimum make
24 restitution in an amount equal to Plaintiffs' contributions that have been misappropriated and by the
25 amount that the intended third-party beneficiaries of the Founding Agreement have been damaged,
26 which is an amount presently unknown, and if necessary, will be proven at trial, but that substantially
27 exceeds this Court's jurisdictional minimum of \$35,000.

28

1 **THIRD CAUSE OF ACTION**

2 **Breach of Fiduciary Duty**

3 ***Against All Defendants***

4 133. Plaintiff realleges and incorporates by reference only paragraphs of this Complaint
5 necessary for his claim of Breach of Fiduciary Duty.

6 134. Under California law, Defendants owe fiduciary duties to Plaintiff, including a duty
7 to use Plaintiff's contributions for the purposes for which they were made. *E.g.*, Cal. Bus. & Prof.
8 Code § 17510.8. Defendants have repeatedly breached their fiduciary duties to Plaintiff, including
9 by:

- 10 a. Using monies received from Plaintiff, and by using intellectual property and
11 derivative works funded by those monies, for “for-profit” purposes that directly
12 contravene both the letter and the express intent of the parties’ agreement, thereby
13 breaching Defendants’ contractual promises to Plaintiff, and also breaching
14 Defendants’ promises to the express intended third-party beneficiaries of the
15 parties’ agreement, *i.e.*, the public. For example, Defendants caused GPT-4, which
16 Microsoft’s own scientists have written can “reasonably be viewed as an early (yet
17 still incomplete) version of an artificial general intelligence (AGI) system,” to be
18 exclusively licensed to Microsoft, in derogation of OpenAI, Inc.’s irrevocable non-
19 profit mission to develop AGI for the benefit of humanity.
- 20 b. Failing to disclose to the public, among other things, details on GPT-4’s
21 architecture, hardware, training method, and training computation, and further by
22 erecting a “paywall” between the public and GPT-4, requiring per-token payment
23 for usage, in order to advance Defendants and Microsoft’s own private commercial
24 interests, despite agreeing that OpenAI’s technology would be open-source,
25 balancing only countervailing safety considerations.
- 26 c. Permitting Microsoft, a publicly traded for-profit corporation, to occupy a seat on
27 OpenAI, Inc.’s Board of Directors and exert undue influence and control over
28 OpenAI, Inc.’s non-profit activities including, for example, the determination of

1 and by those acting in concert with them arising from these acts of unfair competition and other
2 unfair business practices.

3 144. Plaintiff is entitled to restitution and/or disgorgement of any and all monies received
4 by Defendants while they engaged in such practices, in addition to prejudgment interest pursuant to
5 Business & Professions Code § 17200 *et seq.* Plaintiff further seeks to enjoin Defendants from
6 carrying out such activities again in the future, and an order compelling specific performance.

7 **FIFTH CAUSE OF ACTION**

8 **Accounting**

9 ***Against All Defendants***

10 145. Plaintiff realleges and incorporates by reference only paragraphs of this Complaint
11 necessary for his claim of Accounting.

12 146. As discussed above, over the course of multiple years, Defendants solicited tens of
13 millions of dollars in donations from Plaintiff with representations and assurances that OpenAI, Inc.
14 would remain a non-profit irrevocably dedicated to creating safe AGI for public benefit, and that all
15 donated funds would be used in accordance with that “irrevocable” mandate. On information and
16 belief, Defendants also solicited tens of millions of dollars in donations from others based on the
17 same representations.

18 147. Contrary to the express purpose for which Plaintiff’s funds were contributed, as set
19 forth herein, Defendants used monies received from Plaintiff, and intellectual property and
20 derivative works funded by those monies, for “for-profit” purposes, including for the benefit of
21 private individuals and likely of the individual Defendants themselves, that directly contravene both
22 the letter and the express intent of the parties’ agreement.

23 148. Because Defendants are in possession of financial information relating to the use of
24 the charitable contributions that Plaintiff and others made to OpenAI, Inc. and the intellectual
25 property and derivative works funded by those monies, Plaintiff is presently unable to ascertain his
26 interest in or the use, allocation, or distribution of assets without an accounting. Plaintiff is therefore
27 entitled to an accounting.

28

PRAYER FOR RELIEF

WHEREFORE, Plaintiff prays for judgment against Defendants as follows:

- A. For an order compelling specific performance of Defendants’ repeated contractual promises including, without limitation,
 - i. An order requiring that Defendants continue to follow OpenAI’s long-standing practice of making AI research and technology developed at OpenAI available to the public, and
 - ii. An order prohibiting Defendants from utilizing OpenAI, Inc. or its assets for the financial benefit of the individual Defendants, Microsoft, or any other particular person or entity;
- B. For a judicial determination that GPT-4 constitutes Artificial General Intelligence and is thereby outside the scope of OpenAI’s license to Microsoft;
- C. For a judicial determination that Q* and/or other OpenAI next generation large language models in development constitute(s) Artificial General Intelligence and is/are outside the scope of OpenAI’s license to Microsoft;
- D. Injunctive relief consistent with and to effectuate items A through C above;
- E. For an award of restitution and/or disgorgement of any and all monies received by Defendants while they engaged in the unfair and improper practices described herein, in addition to prejudgment interest and penalties, pursuant to Business & Professions Code § 17200 *et seq.*;
- F. For an accounting of funds donated by Plaintiffs and others to OpenAI, Inc., and of intellectual property or derivative works funded by same, and of Defendants’ use of the same for personal benefit or the benefit of any individual or third-party entity;
- G. For general, compensatory, and punitive damages to be proven at trial, each of which Plaintiff will contribute to a non-profit or charity;
- H. For attorneys’ fees pursuant to California Code of Civil Procedure § 1021.5;
- I. For pre-judgment and post-judgment interest at the maximum legal rate; and
- J. For such other and further relief as the Court may deem just and proper.

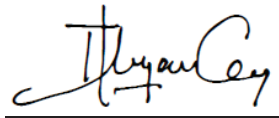
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

DEMAND FOR JURY TRIAL

Plaintiff hereby demands trial by jury as to all issues, claims, and/or causes of action properly triable before a jury

DATED: February 29, 2024

IRELL & MANELLA LLP

By:  _____

Morgan Chu
Alan Heinrich
Iian Jablon
Abigail Sellers
Justin Koo
Henry White

Attorneys for Plaintiff Elon Musk

Exhibit 1

State of Delaware
Secretary of State
Division of Corporations
Delivered 02:22 PM 12/08/2015
FILED 02:22 PM 12/08/2015

SR 20151247198 - File Number 5902936

CERTIFICATE OF INCORPORATION OF

A NON-STOCK CORPORATION

OPENAI, INC.

FIRST: The name of the Corporation is "OpenAI, Inc." (the "Corporation").

SECOND: The address of the Corporation's registered office in the State of Delaware is 2711 Centerville Road, Suite 400, Wilmington, New Castle County, Delaware 19808. The name of its registered agent at such address is Corporation Service Company.

THIRD: This Corporation shall be a nonprofit corporation organized exclusively for charitable and/or educational purposes within the meaning of section 501(c)(3) of the Internal Revenue Code of 1986, as amended, or the corresponding provision of any future United States Internal Revenue law. The specific purpose of this corporation is to provide funding for research, development and distribution of technology related to artificial intelligence. The resulting technology will benefit the public and the corporation will seek to open source technology for the public benefit when applicable. The corporation is not organized for the private gain of any person. In furtherance of its purposes, the corporation shall engage in any lawful act of activity for which nonprofit corporations may be organized under the General Corporation Law of Delaware.

FOURTH: This corporation is organized and operated exclusively for purposes set forth in Article THIRD hereof within the meaning of the Internal Revenue Code section 501(c)(3). No substantial part of the activities of this corporation shall consist of carrying on propaganda, or otherwise attempting to influence legislation and this corporation shall not participate or intervene in any political campaign (including the publishing or distribution of statements) on behalf of any candidate for public office. Notwithstanding any other provision of these articles, the corporation shall not carry on any other activities not permitted to be carried on (a) by a corporation/organization exempt from Federal income tax under section 501(c)(3) of the Internal Revenue Code, or the corresponding section of any future federal tax code, or (b) by a corporation/organization, contributions to which are deductible under section 170(c)(2) of the Internal Revenue Code, or the corresponding section of any future federal tax code.

FIFTH: The property of this corporation is irrevocably dedicated to the purposes in Article THREE hereof and no part of the net income or assets of this corporation shall ever inure to the benefit of any director, officer or member thereof or to the benefit of any private person. Upon the dissolution or winding up of this corporation, its assets remaining after payment, or provision for payment, of all debts and liabilities of this corporation shall be distributed to a nonprofit fund, foundation, or corporation which is organized and operated exclusively for charitable, educational and/or religious purposes and which has established its tax exempt status under Internal Revenue Code section 501(c)(3), or the corresponding section of any future federal tax code, or shall be distributed to the federal government, or to a state or local government, for a public purpose.

SIXTH: The corporation shall not have any capital stock.

SEVENTH: The corporation shall not have any members.

EIGHTH: The name and mailing address of the incorporator are as follows:

**Jonathan Levy
335 Pioneer Way
Mountain View, CA 94041**

I, Jonathan Levy, for the purpose of forming a corporation under the laws of the State of Delaware, do make, file and record this Certificate, and do certify that the facts herein stated are true, and I have accordingly hereunto set my hand this

Executed on December 8, 2015.

AUTHORIZED OFFICER

/s/Jonathan Levy
Jonathan Levy

Exhibit 2

From: Elon Musk [REDACTED]
Sent: Wednesday, June 24, 2015 11:06 PM
To: Sam Altman <[REDACTED]>
Subject: Re: AI lab

Agree on all

On Jun 24, 2015, at 10:24 AM, Sam Altman <[REDACTED]> wrote:

<!--[if !supportLists]-->1) <!--[endif]-->The mission would be to create the first general AI and use it for individual empowerment—ie, the distributed version of the future that seems the safest. More generally, safety should be a first-class requirement.

<!--[if !supportLists]-->2) <!--[endif]-->I think we'd ideally start with a group of 7-10 people, and plan to expand from there. We have a nice extra building in Mountain View they can have.

<!--[if !supportLists]-->3) <!--[endif]-->I think for a governance structure, we should start with 5 people and I'd propose you, [REDACTED], [REDACTED], [REDACTED] and me. The technology would be owned by the foundation and used "for the good of the world", and in cases where it's not obvious how that should be applied the 5 of us would decide. The researchers would have significant financial upside but it would be uncorrelated to what they build, which should eliminate some of the conflict (we'll pay them a competitive salary and give them YC equity for the upside). We'd have an ongoing conversation about what work should be open-sourced and what shouldn't. At some point we'd get someone to run the team, but he/she probably shouldn't be on the governance board.

<!--[if !supportLists]-->4) <!--[endif]-->Will you be involved somehow in addition to just governance? I think that would be really helpful for getting work pointed in the right direction getting the best people to be part of it. Ideally you'd come by and talk to them about progress once a month or whatever. We generically call people involved in some limited way in YC "part-time partners" (we do that with Peter Thiel for example, though at this point he's very involved) but we could call it whatever you want. Even if you can't really spend time on it but can be publicly supportive, that would still probably be really helpful for recruiting.

<!--[if !supportLists]-->5) <!--[endif]-->I think the right plan with the regulation letter is to wait for this to get going and then I can just release it with a message like "now that we are doing this, I've been thinking a lot about what sort of constraints the world needs for safety." I'm happy to leave you off as a signatory. I also suspect that after it's out more people will be willing to get behind it.


Sam

Exhibit 3

Blog

Introducing OpenAI

OpenAI is a non-profit artificial intelligence research company. Our goal is to advance digital intelligence in the way that is most likely to benefit humanity as a whole, unconstrained by a need to generate financial return. Since our research is free from financial obligations, we can better focus on a positive human impact.



December 11, 2015

Authors

[Greg Brockman](#) ↓

[Ilya Sutskever](#) ↓

[OpenAI](#) ↓

Announcements

OpenAI is a non-profit artificial intelligence research company. Our goal is to advance digital intelligence in the way that is most likely to benefit humanity as a whole, unconstrained by a need to generate financial return. Since our research is free from financial obligations, we can better focus on a positive human impact.

We believe AI should be an extension of individual human wills and, in the spirit of liberty, as broadly and evenly distributed as possible. The outcome of this venture is uncertain and the work is difficult, but we believe the goal and the structure are right. We hope this is what matters most to the best in the field.

Background

Artificial intelligence has always been a surprising field. In the early days, people thought that solving certain tasks (such as chess) would lead us to discover human-level intelligence algorithms. However, the solution to each task turned out to be much less general than people were hoping (such as doing a search over a huge number of moves).

The past few years have held another flavor of surprise. An AI technique explored for decades, deep learning, started achieving state-of-the-art results in a wide variety of problem domains. In deep learning, rather than hand-code a new algorithm for each problem, you design architectures that can twist themselves into a wide range of algorithms based on the data you feed them.

This approach has yielded outstanding results on pattern recognition problems, such as recognizing objects in images, machine translation, and speech recognition. But we've also started to see what it might be like for computers to be creative, to dream, and to experience the world.

Looking forward

AI systems today have impressive but narrow capabilities. It seems that we'll keep whittling away at their constraints, and in the extreme case they will reach human performance on virtually every intellectual task. It's hard to fathom how much human-level AI could benefit society, and it's equally hard to imagine how much it could damage society if built or used incorrectly.

OpenAI

Because of AI's surprising history, it's hard to predict when human-level AI might come within reach. When it does, it'll be important to have a leading research institution which can prioritize a good outcome for all over its own self-interest.

We're hoping to grow OpenAI into such an institution. As a non-profit, our aim is to build value for everyone rather than shareholders. Researchers will be strongly encouraged to publish their work, whether as papers, blog posts, or code, and our patents (if any) will be shared with the world. We'll freely collaborate with others across many institutions and expect to work with companies to research and deploy new technologies.

OpenAI's research director is [Ilya Sutskever](#), one of the world experts in machine learning. Our CTO is [Greg Brockman](#), formerly the CTO of Stripe. The group's other founding members are world-class research engineers and scientists: [Trevor Blackwell](#), [Vicki Cheung](#), [Andrej Karpathy](#), [Durk Kingma](#), [John Schulman](#), [Pamela Vagata](#), and [Wojciech Zaremba](#). Pieter Abbeel, Yoshua Bengio, Alan Kay, Sergey Levine, and Vishal Sikka are advisors to the group. OpenAI's co-chairs are Sam Altman and Elon Musk.

Sam, Greg, Elon, Reid Hoffman, Jessica Livingston, Peter Thiel, Amazon Web Services (AWS), Infosys, and [YC Research](#) are donating to support OpenAI. In total, these funders have committed \$1 billion, although we expect to only spend a tiny fraction of this in the next few years.

You can follow us on Twitter at [@OpenAI](#).

Authors

Greg Brockman

[View all articles](#)

Ilya Sutskever

[View all articles](#)

OpenAI

[View all articles](#)

2

Related research

[View all research](#)

