UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

OAKLAND DIVISION

ELON MUSK, et al., Plaintiffs, v. SAMUEL ALTMAN, et al., Defendants.	Case No. 4:24-cv-04722-YGR **JOINT LETTER BRIEF** Judge: Hon. Yvonne Gonzalez Rogers Magistrate Judge: Hon. Thomas S. Hixson Date Action Filed: August 5, 2024 Trial Date: March 30, 2026

October 6, 2025

The Honorable Thomas S. Hixson
United States District Court, Northern District of California
450 Golden Gate Avenue
San Francisco, CA 94102

Re: *Musk, et al.* v. *Altman, et al.*, Case No. 4:24-cv-04722-YGR

Dear Judge Hixson:

Pursuant to Your Honor's Discovery Standing Order, the undersigned counsel for the OpenAI Defendants and Plaintiff Elon Musk respectfully submit this joint letter addressing recent supplemental productions of text messages by Musk, Shivon Zilis, and Jared Birchall, including in response to the Court's September 22, 2025 Discovery Order. The parties attest that they met and conferred in good faith in an effort to resolve these disputes before filing this joint letter.

ATTESTATION

Pursuant to Civil Local Rule 5.1 regarding signatures, I attest that concurrence in the filing of this document has been obtained from the other signatories.

Dated: October 6, 2025

/s/ Jordan Eth
Jordan Eth

REDACTED VERSION OF DOCUMENT SOUGHT TO BE FILED UNDER SEAL

/s/ Bradley R. Wilson
Bradley R. Wilson (admitted *pro hac vice)*
Nathaniel Cullerton (admitted *pro hac vice*)
WACHTELL, LIPTON, ROSEN & KATZ
51 West 52nd Street
New York, NY 10019
Telephone: (212) 403-1000
Facsimile: (212) 403-2000

Attorneys for Defendants Samuel Altman, Gregory Brockman, OpenAI, Inc., OpenAI L.P., OpenAI, L.L.C., OpenAI GP, L.L.C., OpenAI OpCo, LLC, OpenAI Global, LLC, OAI Corporation, LLC, OpenAI Holdings, LLC, OpenAI Startup Fund Management, LLC, OpenAI Startup Fund GP I, L.L.C., OpenAI Startup Fund I, L.P., OpenAI Startup Fund SPV GP I, L.L.C., OpenAI Startup Fund SPV GP II, L.L.C., OpenAI Startup Fund SPV GP III, L.L.C., OpenAI Startup Fund SPV GP IV, L.L.C., OpenAI Startup Fund SPV I, L.P., OpenAI Startup Fund SPV II, L.P., OpenAI Startup Fund SPV III, L.P., OpenAI Startup Fund SPV IV, L.P., Aestas Management Company, LLC, and Aestas LLC

/s/ Jaymie Parkkinen
Jaymie Parkkinen, SBN 318394
Marc Toberoff, SBN 188547
TOBEROFF & ASSOCIATES, P.C.
23823 Malibu Road, Suite 50-363
Malibu, CA 90265
Telephone: (310) 246-3333
Facsimile: (310) 246-3101

Steven F. Molo (*pro hac vice*)
Robert K. Kry (*pro hac vice*)
Jennifer Schubert (*pro hac vice*)
MOLOLAMKEN LLP
430 Park Avenue
New York, NY 10022
Telephone: (212) 607-8160

Attorneys for Plaintiff Elon Musk and Non-Parties Shivon Zilis and Jared Birchall

OpenAI Defendants' Statement

Introduction. The Court in its September 22 order recognized that counsel for Elon Musk and two of his closest confidants, Shivon Zilis and Jared Birchall, had "not been transparent" in the discovery process and had allowed things to go "significantly wrong" with the production of texts from these critical custodians. Dkt. 276 at 1. The Court entered that order two weeks ago, and the problems Your Honor identified have not been fixed. The supplemental productions of Musk, Zilis, and Birchall texts are implausibly paltry, and counsel has refused to explain why or even investigate obvious problems with those productions. Making matters worse, Musk made his supplemental text productions in a spreadsheet format that strips the responsive messages from their surrounding context and dilutes their evidentiary value. This violates an agreement the parties reached months ago and departs from each side's uniform practice up to now, but Musk's counsel has refused to honor their agreement and fix the problem.

Background. On September 22, the Court entered its order obligating Musk, Zilis, and Birchall and their shared counsel to (1) serve a declaration "explaining the steps they took to comply with the subpoenas to Zilis and Birchall"; (2) "perform a linear review of all text messages between Zilis and Birchall, or between either of them and Musk"; (3) serve "detailed declarations" concerning Musk's collection and review of texts, including as to the "role that Tesla and any other Musk-controlled company has played" in that process; and (4) "perform a similar linear review of text messages between Musk and Zilis or between Musk and Birchall." Dkt. 276 at 1-2.

On September 24, Musk, Zilis, and Birchall served a declaration in purported compliance with the Court's order. *See* Ex. 1. [redacted]. OpenAI promptly asked follow-up questions (Ex. 2 at *6-9), and counsel for Musk, Zilis, and Birchall responded four days later. Counsel acknowledged that [redacted] Moreover, with respect to OpenAI's question whether anyone at other Musk-controlled companies had participated in the text review, [redacted] Ex. 2 at *1-4.

During the evening of September 25, just hours before his deposition, Musk produced an excel file with 101 individual texts from the 2015 to 2020 period that counsel claimed had been identified with search terms. These texts were reportedly extracted from an image of one of Musk's old phones preserved by Tesla that Musk's current counsel only recently discovered. On September 27, before the court-ordered linear reviews had commenced, OpenAI objected to the production format for these Musk texts and inquired about the total lack of texts between Musk and Sam Teller, another key advisor. Ex. 3. When Musk's counsel finally responded four days later, it said it was too late to fix the format or investigate why no Teller texts had been produced. Counsel refused to disclose how many Teller texts were in the review population. Ex. 4 at *2.

On October 1, Musk produced three more excel files containing 44 total texts with Zilis and Birchall (individual texts, not text chains) that were said to reflect a linear review of messages

recovered from Musk's old phone image. Later that day, Musk produced another excel file with 64 more texts with Zilis—all but three of which were fully redacted (but not logged). These appear to be the only messages Musk intends to produce based on a linear review of his current phone. Also on October 1, Zilis produced 19 text chains. On the same day, Birchall produced 9 text chains. The latter two productions apparently resulted from a linear review of Zilis's and Birchall's texts from 2015 to the present.

Deficiencies. The supplemental Musk, Zilis, and Birchall text productions are deficient in several material respects. *First*, the linear reviews were conducted in extreme haste—it is unclear by whom—and yielded an implausibly low responsiveness rate. For example, at 11:20 p.m. PDT on September 30, Musk's counsel reported that linear review of some 12,000 Musk texts had not commenced and estimated that the review would take several days. Ex. 4 at *3-4. Roughly 17 hours later, however, counsel claimed that it had completed that review and produced exactly two responsive text strings. *Id.* at *1. A responsible linear review of that size cannot be completed in such a compressed time period, much less subject to an appropriate QC review.

This hasty effort unsurprisingly generated almost nothing. The linear review of Musk's texts with Zilis and Birchall resulted in a production of only 47 texts (excluding fully redacted texts) from a universe of 15,153. That is a responsiveness rate of 0.3% for Musk's texts with two of his closest advisors over a ten-year period that includes the entire time when Musk was directly involved at OpenAI. Similarly:

- For the linear review of Zilis's texts, Plaintiffs' counsel reviewed 1,906 text chains and produced only 19—a 1% responsiveness rate.
- For the linear review of Birchall's texts, Plaintiffs' counsel reviewed 1,003 text strings and produced only 9—a 0.9% responsiveness rate.

These figures are exceedingly low in absolute terms and compared to the responsiveness rate for OpenAI's linear reviews. Counsel tries to explain the meager output of the linear reviews by claiming that Musk, Zilis, and Birchall produced a "large number" of texts based on search terms. In reality, they produced almost no texts with each other prior to the fact discovery deadline. That is why a linear review was necessary. The figures that counsel cites below are substantially inflated by Musk, Zilis, and Birchall's texts with others—principally Altman and Brockman.

Second, the production of Musk's 2015 to 2020 texts resulting from search terms contains a major gap that has not been explained: Musk produced zero one-on-one text messages with Teller, Musk's longtime Chief of Staff. The OpenAI Defendants asked about this gap within 48 hours, but Musk's counsel refused to investigate, saying it is too "late in the process" to do so. Counsel now claims there is "no evidence that Musk ever texted Teller about OpenAI," but the evidence is abundant. Teller was Musk's right-hand for years, including during OpenAI's 2015 formation; Musk's 2017 negotiations with the other founders about a for-profit pivot; and Musk's 2018 separation from OpenAI. [redacted] And Teller is also on multiple group texts with Musk about OpenAI that have been produced by others.

One possible explanation for the missing Musk/Teller texts is that Musk's search terms did not work. Both Zilis and Birchall produced key texts via linear reviews that should have been produced earlier (because they contain search terms such as "OpenAI" or "Altman") but were not. Another possibility is that responsive Musk/Teller texts were filtered out in the Tesla screening process. We do not know because Musk's counsel has refused to look into the problem.

Third, Musk's recent text productions violate the parties' longstanding agreement that texts would be "unitized into 24-hour periods of communications . . . between message participants" and produced in that format. The parties agreed to use that format in May, and prior to September 25, every text that OpenAI, Musk, Zilis, or Birchall produced complied with the agreement. Counsel does not dispute any of this, nor does it present any evidence that complying with the parties' agreement was impossible. Nevertheless, Musk's recent productions adopted a new format (apparently dictated by Tesla) with individual texts presented as line entries in an excel workbook, with no surrounding context, and with no ability to view attached files. The problem with this alternative approach is illustrated by this example: Musk recently produced an isolated text message from March 2019 in which Altman asked Musk for a "few mins to talk" about a "new openai structure" and solicited Musk's feedback on that structure. But Musk's excel file omits subsequent texts, produced by OpenAI, in which Musk and Altman set a time for the call. *Compare* Ex. 5, *with* Ex. 6. This crucial context shows that Musk discussed with Altman in advance one of the main transactions he now challenges, but it is omitted from Musk's production because of his improper production format. That is highly prejudicial to OpenAI.

Requested Relief. The Court should order Musk, Zilis, and Birchall to (1) submit for inspection *in camera* all texts deemed non-responsive in the linear reviews from the key months of August 2017, September 2017, and February 2018; (2) direct trial counsel of record to redo all linear reviews previously ordered by the Court; (3) direct trial counsel of record to conduct a linear review of Musk's texts with Teller; and (4) reproduce all of Musk's texts in a unitized 24-hour format. In addition, given the repeated failure by counsel to ensure that Musk, Zilis, and Birchall fulfill their obligations, OpenAI respectfully submits that the Court should consider appointing a special master or special counsel to review the Musk, Zilis, and Birchall text messages to ensure that all of their relevant documents are produced.

Plaintiff's Statement

Introduction. Plaintiff's counsel diligently undertook the linear reviews this Court ordered, and through extensive efforts, completed those reviews by the fact discovery deadline. Mindful of this Court's concerns about transparency, Plaintiff was very transparent throughout that process, providing daily updates with statistical information on Plaintiff's progress. Those linear reviews confirm the adequacy of Plaintiff's earlier search term reviews. For Zilis, the linear reviews resulted in only a few additional responsive text strings (19) beyond the large number Zilis had already produced using search terms (1,168). For Musk and Birchall, the linear reviews similarly resulted in only small numbers of additional text strings (10 and 9 respectively) beyond the moderate number already produced (209 and 31). That the search term reviews and linear reviews both yielded only modest numbers of responsive texts for Musk and Birchall is not evidence of a broken discovery process. Rather, it simply confirms that those two custodians did not use texts to communicate about OpenAI very often.

A chart showing the complete production history is attached as Appendix A. As that chart shows, Musk, Zilis, and Birchall have now produced about 1,400 text strings, almost exactly the same number as the 1,400 text strings that OpenAI has produced in this case. Dkt. 240 at 3.

Background. By September 9, following traditional search term reviews, Musk had produced 108 text strings, Zilis 1,168 text strings, and Birchall 31 text strings. App. A. On September 12, the Court ordered Musk to search for old phone data from counsel in other cases. Dkt. 263. On September 19, Plaintiff's counsel learned that Tesla had located a backup of old phone data. Plaintiff's counsel reviewed that backup data using search terms and produced 101 responsive text messages on September 25. Of those 101 messages, 17 were texts with Altman that OpenAI admittedly did not produce from Altman's own text messages. Ex. 7 at 1-2.

On September 22, the Court ordered Plaintiff's counsel to conduct linear reviews of all texts between Musk and Zilis, Musk and Birchall, and Birchall and Zilis, without the use of search terms. Dkt. 276. The Court also ordered Plaintiff's counsel to serve a declaration describing prior discovery efforts. *Id.* Plaintiff's counsel served that declaration on September 24. Ex. 1. The next day, OpenAI sent a list of 36 follow-up questions (including subparts). On September 29, Plaintiff's counsel responded to every one of those follow-up questions. Ex. 2.

Meanwhile, Plaintiff's counsel obtained the necessary data exports for the linear reviews and had that data processed for review. Plaintiff's counsel then conducted the linear reviews from September 28 to October 1, devoting extensive resources to complete the reviews by the October 1 deadline. For Zilis's texts, counsel reviewed 1,452 text strings with Musk and 454 text strings with Birchall, of which 19 were responsive and not previously produced. Ex. 4. For Birchall's texts, counsel reviewed 580 text strings with Musk and 423 text strings with Zilis, of which 9 were responsive and not previously produced. *Id.* For Musk, counsel reviewed 13,154 individual text messages with Zilis and 1,999 individual text messages with Birchall, of which 10 text strings were responsive. *Id.* Plaintiff's counsel produced the additional texts on September 30 and October 1. App. A. Throughout, Plaintiff's counsel kept OpenAI apprised with detailed statistical information, sending updates on September 29, September 30, and October 1. Ex. 4 at 1, 3-4, 7.

Argument. Plaintiff's counsel took this Court's concerns to heart and conducted linear reviews that were diligent, thorough, and transparent. The results of those reviews confirm the adequacy of Plaintiff's prior discovery efforts. OpenAI urges that the linear review results are "implausibly paltry." To the contrary, the fact that the linear reviews identified only a small number of additional responsive texts confirms that the search term reviews Plaintiffs already conducted were adequate. OpenAI offers no evidence or even explanation as to why the responsiveness rates are "implausible." For several reasons, they are not.

First, OpenAI ignores that these were ***linear reviews***, conducted without the benefit of any search terms whatsoever. Even for a search term review, a responsiveness rate in the 5% to 15% range would not be uncommon. Where reviews are conducted without any search term filtering at all, responsiveness rates less than 1% are not at all surprising.

Second, the results are qualitatively consistent with the three custodians' involvement with OpenAI. Zilis was a board member of OpenAI from around 2019 to 2023 and thus understandably had a large number of responsive texts with other OpenAI personnel (but not Musk or Birchall). By contrast, Musk resigned from OpenAI in February 2018 and had little connection to OpenAI thereafter. Birchall never had any significant connection to OpenAI at all; he merely helped coordinate Musk's financial contributions in early years and was tangentially involved in the 2017 restructuring discussions. Musk, Birchall, and Zilis had plenty of other reasons to text each other, both work-related and personal. There is simply no reason to assume that they must have been texting each other frequently about the OpenAI Phase I discovery topics in this case.

Third, the fact that the search term reviews and linear reviews both independently yielded consistent results confirms that the reason for the low responsiveness rates is that Musk, Zilis, and Birchall did not text each other about OpenAI very often. The whole point of the linear reviews was to confirm or refute the adequacy of the search term reviews. That has now been done, and the linear reviews confirmed that the prior searches were sufficient.

Fourth, the overall numbers refute OpenAI's accusations. Musk, Zilis, and Birchall produced about 1,400 text strings overall, almost exactly the same number as OpenAI. App. A; Dkt. 240 at 3. Zilis alone produced over 1,100 text strings. App. A. OpenAI cannot fairly assert that Plaintiff's text productions are paltry when they are in line with OpenAI's own numbers.

OpenAI's remaining complaints are meritless. ***First***, OpenAI's complaint that Plaintiff's counsel reviewed the Musk-Zilis texts too quickly ignores that the "12,000" figure refers to individual consecutive text messages, not text strings. Those messages amounted to only 927 one-day text strings. Plaintiff's counsel's review team worked diligently to review those text strings in less than one day. That review was facilitated by the fact that nearly all of the 2020-2025 Musk-Zilis texts were very obviously nonresponsive.

Second, Plaintiff's counsel reasonably declined OpenAI's last-minute September 27 request to investigate Musk's texts with Sam Teller. Ex. 4 at 2. There is no evidence that Musk ever texted Teller about OpenAI. An investigation would have required yet another linear review of Musk's texts with another party. That would have required a new data export and new data processing, which likely would have taken over a week even apart from the review. Counsel were

already very busy complying with this Court's existing linear review order and preparing for and/or attending three depositions during the last few days of fact discovery. A new investigation into the Musk-Teller texts was not warranted at that very late date.

Third, OpenAI's scattered references to Tesla's data pre-screening processes are irrelevant because none of the linear review sets was pre-screened. As per the Court's order, the linear review sets were all turned over to Plaintiff's counsel for review directly.

Finally, OpenAI's formatting objections are misplaced. The parties never ultimately agreed on an ESI protocol for text messages. Even so, Plaintiff's counsel produced the vast majority of texts (1,335 out of 1,446) in OpenAI's preferred short message format, which includes surrounding one-day text strings. For the linear reviews, the Zilis and Birchall text strings were produced in that format. App. A. The Musk linear review texts were produced as an Excel table that contained equivalent information (although certain nonresponsive surrounding texts were redacted on personal privacy grounds, consistent with both parties' practice throughout the case). *Id.* The only production for which counsel did not produce surrounding nonresponsive texts was the Musk "old phone" search-term set on September 25. *Id.* Those texts were extracted from legacy backup data for an old phone, and the data export that Plaintiff's counsel received did not permit the production of surrounding nonresponsive texts. Ex. 4 at 2-3.

Plaintiff's counsel worked diligently and in good faith to conduct the linear reviews this Court ordered to cross-check the adequacy of their prior discovery efforts. Those additional reviews have now served their purpose and confirmed that the prior work done was sufficient. OpenAI has identified no deficiency in that process and shown no other basis for further relief.

Appendix A – Musk, Zilis, and Birchall Text Productions

Custodian	Date	Production	Count	Search Type	Format
Musk (new phone)	7/13/2025	2024MUSK-004	85	search terms	full strings
Zilis	8/22/2025	2024MUSK-011_Zilis	66	search terms	full strings
Musk (new phone)	8/25/2025	2024MUSK-012	23	search terms	full strings
Zilis	9/1/2025	2024MUSK-016_Zilis	813	search terms	full strings
Birchall	9/5/2025	2024MUSK-017	26	search terms	full strings
Zilis	9/8/2025	2024MUSK-018_Zilis	41	search terms	full strings
Zilis	9/9/2025	2024MUSK-019_Zilis	248	search terms	full strings
Birchall	9/9/2025	2024MUSK-020	5	search terms	full strings
Musk (old phone)	9/25/2025	2024MUSK-026	101	search terms	single texts (Excel)
Zilis	9/30/2025	2024MUSK-036_Zilis	12	linear review	full strings
Zilis	10/1/2025	2024MUSK-038_Zilis	7	linear review	full strings
Birchall	10/1/2025	2024MUSK-039	9	linear review	full strings
Musk (old phone)	10/1/2025	2024MUSK-040	8	linear review	full strings (Excel)
Musk (new phone)	10/1/2025	2024MUSK-041	2	linear review	full strings (Excel)
Total			**1,446**		