

EXHIBIT

9

**UNITED STATES DISTRICT COURT
WESTERN DISTRICT OF TEXAS
MIDLAND/ODESSA DIVISION**

NEURAL AI, LLC,

Plaintiff,

v.

NVIDIA CORPORATION,

Defendants.

Case No. 7:24-cv-00221-ADA-DTG

JURY TRIAL DEMANDED

**NON-PARTY TESLA, INC.’S OBJECTIONS AND RESPONSES TO PLAINTIFF
NEURAL AI, LLC’S SUBPOENA**

Non-Party Tesla, Inc. (“Tesla”) hereby serves the following objections and responses (“Responses”) to Plaintiff Neural AI, LLC’s (“Neural AI” or “Plaintiff”) (1) Subpoena to Produce Documents, Information, or Objects, and (2) Subpoena for Testimony.

PRELIMINARY STATEMENT

1. Tesla’s objections and responses to the Requests are made to the best of its current knowledge, information, belief, and understanding of the Requests. Tesla reserves the right to supplement or amend any responses should future investigation indicate that such supplementation or amendment is necessary.

2. Tesla’s responses to the Requests are made solely for the purpose of and in relation to the above-captioned action. Each response is given subject to all appropriate objections (including, but not limited to, objections concerning privilege, competency, relevancy, materiality, propriety, and admissibility). All objections are reserved and may be interposed at any time.

3. Tesla’s responses include only information that is within Tesla’s possession, custody, or control.

4. Tesla incorporates by reference each and every general objection set forth below into each and every specific response. From time to time, a specific response may repeat a general objection for emphasis or some other reason. The failure to include any general objection in any specific response shall not be interpreted as a waiver of any general objection to that response.

5. Nothing contained in these Responses and Objections or provided in response to the Requests consists of, or should be construed as, an admission relating to the accuracy, relevance, existence, or nonexistence of any alleged facts or information referenced in any Request.

GENERAL OBJECTIONS

1. Tesla objects to the Subpoena to the extent it seeks the disclosure of Tesla's highly confidential, proprietary, or trade secret technical and business information, including but not limited to internal software architectures, source code, AI/ML model designs, GPU computing infrastructure, data flow and execution flow diagrams, and engineering specifications. As a non-party to this litigation, the burden on Tesla to review and produce such highly sensitive competitive information outweighs the potential relevance to the underlying action.

2. Tesla generally objects to each Request, including the Definitions and Instructions, on the grounds and to the extent that it purports to impose obligations beyond those imposed or authorized by the Federal Rules, the Federal Rules of Evidence, the Local Rules, the Court's standing orders, any other applicable federal or state law, and any agreements between the parties. Tesla will construe and respond to the Requests in accordance with the requirements of the Federal Rules and other applicable rules or laws.

3. Tesla generally objects to Defendant's Requests on the grounds that they are overbroad, oppressive, unduly burdensome, and disproportionate to the needs of the case, particularly given that Tesla is not a party to this litigation. Tesla objects that the overbreadth

of this Subpoena subjects it to undue burden and expense in both searching for and producing the documents called for by this Subpoena.

4. Tesla generally objects to each Request, including the Definitions and Instructions, to the extent the Request seeks documents and information that are irrelevant to the claims in, or defenses to, this action, is disproportionate to the needs of the case, and/or is of such marginal relevance that its probative value is outweighed by the burden imposed on Tesla in having to provide such information, including any Request that seeks information for any time period outside that which is relevant to the claims and defenses asserted in this action particularly given that Tesla is not a party to this litigation.

5. Tesla generally objects to each Request, including the Definitions and Instructions, to the extent the Request is vague, ambiguous, unreasonably cumulative, or duplicative, including to the extent it seeks documents or communications that are otherwise responsive to other specific Requests.

6. Tesla generally objects to the Requests, including the Definitions and Instructions, on the basis that they specify an overbroad and unduly burdensome time period, or a time period when the asserted patents in the underlying action were not in force, and seek documents and things outside of the time period relevant to the claims and defenses asserted in this action.

7. Tesla generally objects to the Requests, including the Definitions and Instructions, to the extent that they are overbroad and/or unduly burdensome, including to the extent that they call for the production of “Any,” “any” or “all” documents or communications concerning the subject matter referenced therein.

8. Tesla generally objects to each Request, including the Definitions and Instructions, on the grounds and to the extent the Request purports to request the identification and disclosure of any information, communication(s), or document(s) that were prepared in anticipation of litigation or in connection with any internal investigation conducted at the direction of counsel, constitute attorney work product, reveal privileged attorney-client

communications, are covered under the common interest privilege and/or joint defense privilege, or are otherwise protected or immune from disclosure under any applicable privilege(s), law(s), or rule(s). Tesla hereby asserts all such applicable privileges and protections and excludes privileged and protected information from its responses to each Request. *See generally* Fed. R. Evid. 502.

9. Tesla generally objects to any Request to the extent it seeks production of information and/or documents that comprise or contain confidential information of a third party to whom Tesla believes it owes a duty of confidentiality or otherwise protected from disclosure by agreements between Tesla and other parties.

10. Tesla generally objects to any Request to the extent that it seeks to require Tesla to provide any information beyond what is available to Tesla at the present time after reasonable search of its own records and a reasonable inquiry of its present employees. For example, Tesla objects to any request that seeks information that is not within Tesla's possession, custody or control. Tesla also objects to any Request that seeks to impose a duty on Tesla to create materials that Tesla does not create or maintain in the ordinary course of business.

11. Tesla generally objects to each Request, including the Definitions and Instructions, to the extent that it requests information that is confidential, proprietary, or competitively sensitive.

12. Tesla generally objects to each Request to the extent that the information sought is more appropriately pursued through another discovery tool.

13. Tesla generally objects to each Request, including the Definitions and Instructions, to the extent it is argumentative, lacks foundation, or incorporates allegations and assertions that are disputed or erroneous. In furnishing the responses herein, Tesla does not concede the truth of any factual assertion or implication contained in any Request, Definition, or Instruction.

14. Tesla generally objects to any Request to the extent that it seeks information or materials that are publicly available, already in Defendant's possession, custody, or control, or are equally available to Defendant from another less burdensome source, such as parties to the litigation.

15. Tesla generally objects to any Request to the extent that it fails to describe the information requested with reasonable particularity, is indefinite as to time and scope, seeks information that is not relevant to the claims or defenses of the parties in this case, and/or is not proportional to the needs of the case.

16. Tesla generally objects to any Request to the extent that it requires Tesla to draw legal conclusions.

17. Tesla generally objects to each Request, including the Definitions and Instructions, to the extent that it purports to impose an obligation to conduct anything beyond a reasonable and diligent search of reasonably accessible files (including electronic files) where responsive documents reasonably would be expected to be found. Any Requests that seek to require Tesla to go beyond such a search are overbroad and unduly burdensome.

18. Tesla generally objects to the Subpoena to the extent it seeks production of Tesla's proprietary source code, internal software, custom code, configuration files, build files, deployment files, runtime logs, profiler traces, or other engineering artifacts. Such materials constitute Tesla's core intellectual property and trade secrets, and their production to a non-party in a dispute between Neural AI and NVIDIA is disproportionate, unduly burdensome, and risks competitive harm.

19. Tesla's willingness to provide any document or information in response to a Request shall not be interpreted as an admission that such document or information exists, that it is relevant to a claim or defense in this action, or that it is admissible for any purpose. Tesla does not waive its right to object to the admissibility of any document or information produced by any party on any ground.

OBJECTIONS TO THE DEFINITIONS AND INSTRUCTIONS

1. Tesla objects to Plaintiff's definition of "NVIDIA GPUs" as overbroad, unduly burdensome, and disproportionate to the needs of the case. The definition encompasses virtually every NVIDIA GPU product ever manufactured across eight architecture generations, including hundreds of individual product SKUs spanning consumer, enterprise, data center, and embedded platforms. This sweeping definition, combined with the document requests, would require Tesla to search for, collect, and review documents relating to any and all NVIDIA hardware it has ever used, regardless of relevance to the patents-in-suit. Further, such information can be obtained through other, less burdensome and more appropriate means, including from parties to the litigation.

2. Tesla objects to Plaintiff's definitions of "and" and "or" as overbroad, unduly burdensome, impermissibly vague, and not proportional to the needs of the case, to the extent they purport to change the customary and usual meaning of these terms and alter the meaning of a phrase to impose requirements in excess of those under the Federal Rules of Civil Procedure and the Local Rules of this Court.

3. Tesla objects to Plaintiff's definitions of "any" and "each" as overbroad, unduly burdensome, impermissibly vague, and not proportional to the needs of the case, to the extent it purports to seek information that is unrelated to the present case, to the extent would capture documents of no evidentiary value and impose an undue burden on a non-party, and to the extent it exceeds the obligations imposed by the Federal Rules, the local rules of this Court, and any orders this Court entered in this case.

4. Tesla objects to Plaintiff's definitions of "concerning," "related to," "relating to," and "regarding" as overbroad, vague, and disproportionate to the needs of the case. The definitions encompass over twenty verbs including "alluding to," "contradicting," "mentioning," and "memorializing," which would capture documents of no evidentiary value and impose an undue burden on a non-party.

5. Tesla objects to Plaintiff's definition of "document(s)" as overly broad and unduly burdensome, to the extent it purports to seek information that is unrelated to the present case, and to the extent it exceeds the obligations imposed by the Federal Rules, the local rules of this Court, and any orders this Court entered in this case. Tesla recognizes that Plaintiff's definition of "document" does not include "Source Code" and Plaintiff has provided a separate definition for "Source Code."

6. Tesla objects to Plaintiff's definition of "persons" as overbroad, unduly burdensome, vague, ambiguous, and not proportional to the needs of the case to the extent it purports to include "formal or informal entities and organizations" and extends to "public and private corporations, partnerships, professional corporations, limited liability companies, business trusts, banking institutions, associations, firms, joint ventures, commissions, bureaus, departments, and any other legal entity, including any divisions, subsidiaries, departments, and other units thereof" regardless of relevance to any claim or defense in this case. Tesla further objects to the inclusion of "informal entities and organizations" and "any other legal entity" as unbounded in scope and undefined, as these terms could be interpreted to encompass virtually any grouping of individuals or organizational unit without limitation. The breadth of this definition, when applied across the Requests, would impose an undue burden on Tesla—a non-party—by requiring it to search for and identify documents involving an unlimited universe of persons and entities with no meaningful nexus to the claims or defenses in this action. Tesla will interpret the term "persons" according to its customary usage in the context of a particular Request and limit any response accordingly.

7. Tesla objects to Plaintiff's definition of "Source Code" as overbroad and unduly burdensome and as calling for information not relevant to the case. The definition encompasses "all versions and revisions," "all associated files," "scripts, header files, makefiles, configuration files, and documentation"—effectively requiring Tesla to potentially produce entire software development repositories relating to any system that touches an NVIDIA GPU. Tesla objects to this definition as it would require the disclosure of sensitive Tesla trade secrets

and intellectual property in a dispute to which Tesla is not a party and without making any showing of relevance or that the request is proportional to the needs of the case.

8. Tesla objects to Plaintiff's definition of "You," and "Your" and its inclusion of "but not limited to its predecessors, successors, parents, subsidiaries, divisions, affiliates, and all past or present directors, officers, partners, managers, employees, contractors, agents, representatives, accountants, consultants, in-house and outside counsel" as overbroad and unduly burdensome to the extent the terms are meant to include any individual(s), entit(ies), or any other person(s) over which Tesla exercises no control and to the extent Defendant purports to use the terms to impose obligations on Tesla that go beyond the requirements of the Federal Rules and the Local Rules. Tesla further objects to these definitions to the extent that Defendant purports to use these defined terms to seek documents that are not relevant to the claims and defenses in this action, including seeking documents from Tesla's subsidiaries and with respect to products not at issue in this litigation. Tesla will construe the terms "You" and "Your" so as to include only the following: Tesla, Inc. and its employees. Tesla further objects to the definition and its inclusion of "in-house and outside counsel" to the extent that it seeks information protected by attorney-client or work product privilege.

9. Tesla objects to Plaintiff's Instructions to the extent they impose obligations that go beyond the requirements of the Federal Rules and the Local Rules.

10. Tesla objects to Plaintiff's instruction regarding production format, metadata, native format, and load file specifications as overbroad, unduly burdensome, and disproportionate to the needs of this case.

11. Tesla objects to Plaintiff's instruction that Tesla identify materials not in its possession, custody, or control as overbroad, vague, and disproportionate to the needs of this case to the extent it purports to impose interrogatory-style narrative obligations on Tesla to explain the absence of documents or to speculate regarding the existence, destruction, or location of documents.

12. Tesla objects to Plaintiff's instruction regarding continuing obligations to supplement to the extent it purports to impose obligations beyond those required by the Federal Rules for non-parties.

13. Tesla objects to Plaintiff's Requests to the extent they fail to impose a meaningful geographic scope limitation, and further objects to the extent it purports to encompass "non-U.S. activity that directly supports or enables U.S. operations or usage." This language is undefined, unbounded, and would require Tesla, a non-party, to make subjective legal and factual determinations regarding the geographic nexus of its global computing operations. As written, the Requests could be construed to require Tesla to search for, collect, and produce documents from any facility, system, or operation worldwide that has any arguable connection to the United States, imposing a burden and expense on a non-party that is grossly disproportionate to the needs of this case. Tesla will construe the Requests as limited to activity occurring within the United States.

14. Tesla objects to Plaintiff's "Instructions" to the extent they impose obligations that go beyond the requirements of the Federal Rules and the Local Rules.

15. Tesla objects to Plaintiff's instruction regarding production format and load file specifications as overbroad, unduly burdensome, and disproportionate to the needs of this case. Tesla further objects to the extent that this instruction purports to impose obligations that go beyond the requirements of the Federal Rules and the Local Rules, as applicable. Tesla will construe this instruction in accordance with the Federal Rules and the Local Rules.

16. Tesla objects to Plaintiff's instruction regarding the production of documents in native format as overbroad, vague, and disproportionate to the needs of this case to the extent it purports to unilaterally define the criteria for native productions that differ from standard industry practice.

17. Tesla objects to Plaintiff's instruction regarding metadata load file requirements as overbroad, unduly burdensome, and disproportionate to the needs of this case to the extent it purports to unilaterally dictate specific metadata that may not be reasonably available,

automatically generated, or maintained in the ordinary course of business without undue burden or expense. Tesla further objects to the extent it purports to unilaterally impose rigid metadata reporting requirements, which as previously stated, may or may not be technically feasible.

18. All Tesla's responses herein and/or related documents will be provided subject to any protective order entered in this case by the Court, or, if no protective order is entered, as HIGHLY CONFIDENTIAL ATTORNEYS' EYES ONLY.

19. The foregoing general reservations and objections are incorporated into each of the responses and objections to the specific Request set forth below.

OBJECTIONS AND RESPONSES TO DOCUMENT REQUESTS

DOCUMENT REQUEST NO. 1:

Documents sufficient to identify all software, frameworks, libraries, APIs, scripts, Source Code, configuration files, and custom code You use to perform computations on NVIDIA GPUs.

RESPONSE TO DOCUMENT REQUEST NO. 1:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions and Instructions, as though fully set forth in this Response.

Tesla objects to this Request as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including Tesla or third-party software, frameworks, libraries, APIs, scripts, Source Code, configuration files, and custom code.

Tesla objects to this Request on the grounds that it is overbroad, unduly burdensome, unnecessary, oppressive, and not proportional to the needs of this case, particularly given that Tesla is not a party to this litigation. The Request seeks identification of "all software, frameworks, libraries, APIs, scripts, Source Code, configuration files, and custom code" used to perform computations on NVIDIA GPUs, a scope that could potentially encompass virtually any software system Tesla operates. Tesla further objects to the terms "all" and "Source Code" as overbroad and unduly burdensome as they would require an exhaustive identification effort that is disproportionate to the discovery need in a dispute to which Tesla is not a party.

Tesla objects to this Request to the extent that this catch-all request would potentially require Tesla to produce an unknowable quantum of engineering documentation representing Tesla's valuable intellectual property. The Request fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation.

Tesla further objects to this Request to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance

or that the request is proportional to the needs of the case. Tesla objects to this Request to the extent it seeks the disclosure of Tesla confidential trade secrets and proprietary technical information, internal software architectures, and custom-built AI/ML frameworks that constitute core intellectual property and are not relevant to this case. Production of such materials, particularly on a third-party, would risk competitive harm to Tesla, a non-party.

Tesla further objects to this Request as vague and ambiguous as to what "computations" are relevant to the underlying litigation, and as to the meaning and scope of the terms "software, frameworks, libraries, APIs, scripts, Source Code, configuration files, and custom code."

Tesla further objects to this Request to the extent it seeks information beyond the use of NVIDIA products.

Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla further objects to the extent this Request seeks to impose a duty on Tesla to create materials or compile information that Tesla does not create or maintain in the ordinary course of business. Tesla further objects on the grounds that information regarding NVIDIA's software, frameworks, and libraries is more readily available from NVIDIA, the Defendant in this action.

Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to this Request to the extent that it calls for documents not within Tesla's possession, custody or control or kept in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Request and the burden it imposes on Tesla.

DOCUMENT REQUEST NO. 2:

Documents sufficient to show whether You use NVIDIA's Aerial, Clara Parabricks, cuBLAS, cuDNN, cuFFT, cuQuantum, cuSOLVER, cuSPARSE, Drive, DriveWorks, Holoscan, Isaac, Isaac Lab, Maxine, Memory Map, Merlin, Metropolis, Modulus, Monai, Morpheus, NeMo, PyTorch, RAPIDS, Riva, Runtime Driver, TensorFlow, TensorRT, Triton, VSS (Deepstream), or any other NVIDIA software as part of computations You perform using NVIDIA GPUs.

RESPONSE TO DOCUMENT REQUEST NO. 2:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Request as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including Tesla's use of software specifically listed in this Request or "any other NVIDIA software."

Tesla objects to this Request on the grounds that it is overbroad, unduly burdensome, unnecessary, and not proportional to the needs of this case. The Request identifies over thirty specific NVIDIA software products and then adds the catch-all phrase "or any other NVIDIA software," rendering the Request virtually unlimited in scope.

Tesla objects to this Request to the extent that this catch-all request would potentially require Tesla to produce an unknowable quantum of engineering documentation representing Tesla's valuable intellectual property. The Request fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation.

Tesla further objects to this Request to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case. Tesla objects to this Request to the extent it seeks the disclosure of Tesla's highly confidential and proprietary information

regarding its internal software stack, computing infrastructure, and technology choices, which constitute competitively sensitive business information. Disclosure of which specific NVIDIA tools Tesla does or does not use would reveal Tesla's internal technology strategy and competitive posture.

Tesla further objects on the grounds that information regarding NVIDIA's software products and their usage by customers is more appropriately sought from NVIDIA, the Defendant in this action, which possesses licensing records, telemetry data, and customer usage information.

Tesla objects that the catch-all phrase "any other NVIDIA software" renders the Request vague, ambiguous, and boundless in scope, as it effectively extends the demand to every NVIDIA product, tool, or library, without temporal or functional limitation.

Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla further objects on the grounds that the information sought is more readily obtainable from the parties to this litigation or from other less burdensome sources, thus disproportionate to the needs of the case.

Tesla further objects to the extent that this Request is duplicative of Request No. 1. Tesla objects to this Request on the grounds that it is unduly burdensome and oppressive to the extent that it seeks information and documents that are equally available to the parties in this litigation.

Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to this Request to the extent that it calls for documents not within Tesla's possession, custody or control or kept in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Request and the burden it imposes on Tesla.

DOCUMENT REQUEST NO. 3:

Documents sufficient to show whether You use sample Source Code provided by NVIDIA as part of computations You perform using NVIDIA GPUs.

RESPONSE TO DOCUMENT REQUEST NO. 3:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Request on the grounds that it is overbroad, unduly burdensome, unnecessary, and not proportional to the needs of this case. The term "sample Source Code" is vague and ambiguous, as it is unclear what specific sample code is being referenced, what universe of NVIDIA sample code is at issue, or how Tesla would be expected to identify whether any portion of its codebase derives from, incorporates, or was inspired by NVIDIA sample code. Tesla objects to this Request to the extent that compliance would require Tesla to conduct an exhaustive audit of its entire codebase against an undefined body of NVIDIA sample code, a task that is extraordinarily burdensome and disproportionate for a non-party.

Tesla objects to this Request to the extent it seeks the disclosure of Tesla's proprietary source code and internal software, which constitute trade secrets and core intellectual property. Tesla further objects on the grounds that NVIDIA sample source code is publicly available and/or in NVIDIA's possession, custody, or control, and information regarding its distribution to customers is more appropriately sought from the Defendant. Tesla further objects to this Request to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case.

Tesla further objects on the grounds that the information sought is more readily obtainable from the parties to this litigation or from other less burdensome sources, thus disproportionate to the needs of the case. Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla

further objects to the extent this Request seeks to impose a duty on Tesla to create materials or compile analyses that Tesla does not create or maintain in the ordinary course of business. Tesla further objects to the extent that this Request is duplicative of Requests Nos. 1 and 2.

Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to this Request to the extent that it calls for documents not within Tesla's possession, custody or control or kept in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Request and the burden it imposes on Tesla.

DOCUMENT REQUEST NO. 4:

Documents sufficient to show whether and how any software You use to perform computations on NVIDIA GPUs calls, invokes, interfaces with, wraps, depends on, sits on top of, modifies, extends, or implements functionality provided by CUDA, cuDNN, TensorRT, CUDA libraries, CUDA drivers, CUDA runtime, CUDA applications or frameworks or any other NVIDIA software.

RESPONSE TO DOCUMENT REQUEST NO. 4:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Request as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including whether and how Tesla software "perform[s] computations on NVIDIA GPUs calls, invokes, interfaces with, wraps, depends on, sits on top of, modifies, extends, or implements functionality provided by CUDA, cuDNN, TensorRT, CUDA libraries, CUDA drivers, CUDA runtime, CUDA applications or frameworks or any other NVIDIA software."

Tesla objects to this Request on the grounds that it is overbroad, unduly burdensome, unnecessary, oppressive, and not proportional to the needs of this case. Tesla objects to this Request as being vague and ambiguous to the extent it seeks documents showing "whether and how" Tesla's software "calls, invokes, interfaces with, wraps, depends on, sits on top of, modifies, extends, or implements" NVIDIA functionality.

Tesla objects to this Request to the extent that this catch-all request would potentially require Tesla to produce an unknowable quantum of engineering documentation representing Tesla's valuable intellectual property. The Request fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation.

Tesla objects to this Request to the extent it implicates a scope that would require Tesla to map and document every interaction between its proprietary software systems and NVIDIA's computing stack across its entire business. The enumeration of nine distinct verbs describing software interaction, combined with the catch-all phrase "or any other NVIDIA software," renders the Request virtually unlimited in scope.

Tesla further objects to this Request to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case. Tesla objects to this Request to the extent it seeks the disclosure of Tesla's confidential trade secrets and proprietary technical information, including internal software architectures, custom integrations, dependency structures, and engineering designs that constitute core intellectual property.

Tesla objects to this Request as compound and vague. The Request combines at least nine distinct NVIDIA technologies with at least nine distinct functional relationships, creating a matrix of discrete inquiries posed as a single request.

Tesla further objects to the terms "sits on top of," "wraps," and "interfaces with" as vague, ambiguous, and susceptible to multiple technical interpretations that render meaningful compliance impossible without speculation. The catch-all phrase "any other NVIDIA

software" renders the Request boundless in scope, extending the demand to every NVIDIA product, tool, or library, without temporal or functional limitation.

Tesla further objects to the extent this Request seeks to impose a duty on Tesla to create materials, such as architectural mappings, dependency analyses, or software interaction diagrams, that Tesla does not create or maintain in the ordinary course of business. Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla further objects to the extent that this Request is duplicative of Requests Nos. 1, 2, and 3.

Tesla further objects on the grounds that information regarding NVIDIA's CUDA platform, software libraries, and their interfaces is extensively documented in NVIDIA's public developer resources and is more appropriately sought from the Defendant. Tesla further objections to this Request to the extent it seeks information beyond the use of NVIDIA products.

Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to this Request to the extent that it calls for documents not within Tesla's possession, custody or control or kept in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Request and the burden it imposes on Tesla.

DOCUMENT REQUEST NO. 5:

Documents sufficient to show the architecture, design, data flow, control flow, and execution flow of any system in which You use NVIDIA GPUs to perform computations, including diagrams, technical specifications, design documents, Powerpoints, slide decks, internal and external presentations, Source Code, configuration files, build files, deployment files, runtime logs, and profiler traces.

RESPONSE TO DOCUMENT REQUEST NO. 5:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Request as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including “the architecture, design, data flow, control flow, and execution flow of any system in which [Tesla] use[s] NVIDIA GPUs to perform computations, including diagrams, technical specifications, design documents, Powerpoints, slide decks, internal and external presentations, Source Code, configuration files, build files, deployment files, runtime logs, and profiler traces.”

Tesla objects to this Request on the grounds that it is overbroad, unduly burdensome, unnecessary, oppressive, and grossly disproportionate to the needs of this case. The Request seeks the "architecture, design, data flow, control flow, and execution flow" of every system in which Tesla uses NVIDIA GPUs, along with virtually every category of engineering artifact, "diagrams, technical specifications, design documents, Powerpoints, slide decks, internal and external presentations, Source Code, configuration files, build files, deployment files, runtime logs, and profiler traces." Tesla objects to this Request to the extent that this catch-all request would potentially require Tesla to produce an unknowable quantum of engineering documentation and source code representing Tesla's valuable intellectual property. The Request fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation.

Tesla further objects to this Request to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case. Tesla objects to this Request to the extent it seeks the disclosure of Tesla's highly confidential trade secrets and proprietary

technical information, including proprietary source code, internal software architectures, and custom-built AI/ML frameworks that constitute core intellectual property. Production of such materials, particularly on a third-party, would risk competitive harm to Tesla, a non-party.

Tesla further objects to this Request to the extent it seeks information beyond the use of NVIDIA products.

Tesla further objects to the extent that the scope of this Request is overbroad and bears no reasonable relationship to the claims or defenses in the underlying action. Tesla further objects to the extent this Request seeks to impose a duty on Tesla to create materials that Tesla does not create or maintain in the ordinary course of business. Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla further objects to the extent that this Request is duplicative of prior Requests, including Requests Nos. 1, 2, 3, and 4.

Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to this Request to the extent that it calls for documents not within Tesla's possession, custody or control or kept in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla will not produce any documents in response to this Request.

DOCUMENT REQUEST NO. 6:

Documents sufficient to show whether computations You performed using NVIDIA GPUs involved artificial neural networks, neural-network computational layers or computations with outputs as inputs for other neurons or layers.

RESPONSE TO DOCUMENT REQUEST NO. 6:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Request as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including whether computations Tesla performed using NVIDIA GPUs involved artificial neural networks, neural-network computational layers or computations with outputs as inputs for other neurons or layers.

Tesla objects to this Request on the grounds that it is overbroad, unduly burdensome, unnecessary, and not proportional to the needs of this case. Tesla further objects to this Request to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case.

Tesla further objects to the terms "artificial neural networks, neural-network computational layers or computations with outputs as inputs for other neurons or layers" as vague and ambiguous to the extent they could encompass virtually any computation Tesla performs on NVIDIA hardware.

Tesla further objects on the grounds that the information sought is more readily obtainable from the parties to this litigation or from other less burdensome sources, thus disproportionate to the needs of the case.

Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla further objects to the extent this Request seeks to impose a duty on Tesla to create materials that Tesla does not create or maintain in the ordinary course of business. Tesla further objects on the grounds that information regarding NVIDIA GPUs' neural network capabilities is publicly available and more appropriately sought from NVIDIA. Tesla further objects to the extent that this Request is duplicative of prior Requests, including Requests Nos. 1, 4, and 5.

Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to this Request to the extent that it

calls for documents not within Tesla's possession, custody or control or kept in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Request and the burden it imposes on Tesla.

DOCUMENT REQUEST NO. 7:

Documents sufficient to show whether You use a pointer to data stored in memory (e.g. memory bank or partition), using as an input to a subsequent computational layer the pointer to output data from a GPU computation, using pointers in neural network computations, swapping an input pointer with the pointer to data output from a GPU computation, pointer swapping, pointer rotation, buffer swapping, ping-pong buffers, double or triple buffering, alternating input/output buffers, or any other technique in which output data from one computation, layer, iteration, time step, or cycle becomes input data for a later computation, layer, iteration, time step, or cycle.

RESPONSE TO DOCUMENT REQUEST NO. 7:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Request as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, whether Tesla uses a pointer to data stored in memory.

Tesla objects to this Request to the extent that it is overbroad, unduly burdensome, unnecessary, oppressive, and not proportional to the needs of this case. Tesla objects to this Request to the extent that it identifies a sweeping range of fundamental computing techniques, pointer usage, pointer swapping, pointer rotation, buffer swapping, ping-pong buffers, double or triple buffering, alternating input/output buffers, and then adds the catch-all phrase "or any

other technique in which output data from one computation, layer, iteration, time step, or cycle becomes input data for a later computation, layer, iteration, time step, or cycle."

Tesla objects to this Request to the extent it seeks the disclosure of Tesla's confidential trade secrets and proprietary technical information, including internal memory management strategies, GPU optimization techniques, custom buffer management implementations, and low-level engineering designs that constitute core intellectual property. Tesla further objects to this Request to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case.

Tesla further objects to the extent this Request seeks to impose a duty on Tesla to create materials, such as memory architecture diagrams, pointer flow analyses, or buffer management documentation, that Tesla does not create or maintain in the ordinary course of business.

Tesla further objects on the grounds that the information sought is more readily obtainable from the parties to this litigation or from other less burdensome sources, thus disproportionate to the needs of the case.

Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine.

Tesla further objects on the grounds to the extent that the techniques described, pointer usage, buffer swapping, double buffering, ping-pong buffers, are well-known, standard computing techniques whose operation is extensively and publicly documented in computer science literature, NVIDIA's developer guides, and CUDA programming documentation, and are more appropriately explored through NVIDIA's own materials.

Tesla objects to this Request to the extent that this catch-all request would potentially require Tesla to produce an unknowable quantum of engineering documentation representing Tesla's valuable intellectual property. The Request fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying

litigation, rendering meaningful compliance impossible without speculation. Tesla objects to this Request to the extent it uses generic technical terminology to describe fundamental and ubiquitous computing operations that are not unique to any particular proprietary technology, patented method, or party to this litigation. As drafted, this Request is overbroad, unduly burdensome, and disproportionate to the needs of this case, particularly as directed to a non-party.

Tesla further objects to the extent that this Request is duplicative of prior Requests, including Requests Nos. 1, 4, 5, and 6.

Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to this Request to the extent that it calls for documents not within Tesla's possession, custody or control or kept in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Request and the burden it imposes on Tesla.

DOCUMENT REQUEST NO. 8:

Documents sufficient to show whether You store input data, output data, intermediate results, tensors, activations, weights, parameters, internal variables, GPU programs, kernels, textures, shaders, or other GPU-computation-related data in separate, partitioned, logical, physical, first/second, input/output, texture, shader, shared, global, device, host, pinned, GPU RAM, GPU cache(s), or unified memory regions (shared by CPU and GPU) when performing computations using NVIDIA GPUs.

RESPONSE TO DOCUMENT REQUEST NO. 8:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Request as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including whether Tesla stores certain GPU-computation-related data in certain memory regions.

Tesla further objects to this Request to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case.

Tesla objects to this Request to the extent that it is overbroad, unduly burdensome, unnecessary, oppressive, and grossly disproportionate to the needs of this case to the extent that it would potentially require Tesla to produce an unknowable quantum of engineering documentation representing Tesla's valuable intellectual property. The Request fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation.

Tesla objects to this Request to the extent it seeks the disclosure of Tesla's confidential trade secrets and proprietary intellectual property, including GPU memory allocation strategies, custom memory optimization techniques, tensor management implementations, and internal computing architectures. Tesla further objects to the catch-all phrase "or other GPU-computation-related data" as vague, ambiguous, and unbounded.

Tesla objects to this Request to the extent it uses generic technical terminology to describe fundamental and ubiquitous computing operations that are not unique to any particular proprietary technology, patented method, or party to this litigation. The Request fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation. As drafted, this Request is therefore facially overbroad, unduly burdensome, and disproportionate to the needs of this case, particularly as directed to a non-party.

Tesla further objects to the extent this Request seeks to impose a duty on Tesla to create materials, such as memory allocation maps, data storage analyses, or memory partition documentation, that Tesla does not create or maintain in the ordinary course of business. Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine.

Tesla further objects on the grounds that NVIDIA's own documentation, developer guides, CUDA programming manuals, and hardware specifications describe in detail the memory architecture, memory types, and memory management capabilities of NVIDIA GPUs and are publicly available and more appropriately sought from the Defendant. Tesla further objects to the extent that this Request is duplicative of prior Requests, including Requests Nos. 1, 4, 5, and 7.

Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to this Request to the extent that it calls for documents not within Tesla's possession, custody or control or kept in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Request and the burden it imposes on Tesla.

DOCUMENT REQUEST NO. 9:

Documents sufficient to show how input data is received, acquired, stored, transferred, copied, streamed, prefetched, staged, queued, or loaded from CPU memory, host memory, system memory, storage, sensors, cameras, or other input sources to NVIDIA GPU memory including GPU RAM (e.g. GPU HBM, GDDR) and/or GPU cache(s) before, during, or in parallel with computations You perform using NVIDIA GPUs.

RESPONSE TO DOCUMENT REQUEST NO. 9:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Request as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including how Tesla handles input data related to GPU computations.

Tesla objects to this Request to the extent that it is overbroad, unduly burdensome, unnecessary, oppressive, and not proportional to the needs of this case. Tesla objects to the Request to the extent that it seeks documents showing "how" input data is "received, acquired, stored, transferred, copied, streamed, prefetched, staged, queued, or loaded" from an extensive and open-ended list of sources, "CPU memory, host memory, system memory, storage, sensors, cameras, or other input sources," to NVIDIA GPU memory, "before, during, or in parallel with" computations. This scope encompasses virtually every data pipeline and data ingestion pathway across Tesla's entire computing infrastructure that touches any NVIDIA hardware.

Tesla objects to this Request to the extent it seeks the disclosure of Tesla's confidential trade secrets and proprietary technical information, including internal data pipeline architectures, sensor fusion systems, data preprocessing workflows, custom data loading and prefetching implementations, and GPU optimization strategies that constitute core intellectual property. Tesla specifically objects to the reference to "sensors, cameras, or other input sources" as a transparent attempt to compel disclosure of Tesla's most valuable and competitively sensitive proprietary intellectual property. Tesla further objects to this Request to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case.

Tesla objects to this Request to the extent it uses generic technical terminology to describe fundamental and ubiquitous computing operations that are not unique to any particular proprietary technology, patented method, or party to this litigation. The Request fails to identify

with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation. As drafted, this Request is therefore overbroad, unduly burdensome, and disproportionate to the needs of this case, particularly as directed to a non-party.

Tesla further objects to the catch-all phrase "or other input sources" as vague, ambiguous, and unbounded. Tesla further objects to the extent this Request seeks to impose a duty on Tesla to create materials, such as data flow diagrams, pipeline architecture documents, or data transfer analyses, that Tesla does not create or maintain in the ordinary course of business. Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine.

Tesla further objects on the grounds that the data transfer and memory management mechanisms described—CPU-to-GPU transfers, memory staging, prefetching, streaming—are standard computing operations whose architecture is publicly documented by NVIDIA in its CUDA programming guides and developer documentation and are more appropriately sought from the Defendant. Tesla further objects to the extent that this Request is duplicative of prior Requests, including Requests Nos. 4, 5, 7, and 8.

Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to this Request to the extent that it calls for documents not within Tesla's possession, custody or control or kept in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Request and the burden it imposes on Tesla.

DOCUMENT REQUEST NO. 10:

Documents sufficient to show how output data from a GPU computation(s), intermediate results of GPU computations, tensors, buffers, activations, variables, or other computation results are stored, transferred, copied, streamed, written back, returned, accumulated, reused, or made available including asynchronously from NVIDIA GPU memory to CPU memory, host memory, system memory, storage, display, network, or another memory location before, during, or in parallel with computations You perform using NVIDIA GPUs and also including in the opposite direction, copying data from CPU or host or other memory to a queue for GPU computation while other GPU computations are occurring.

RESPONSE TO DOCUMENT REQUEST NO. 10:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions and Instructions, as though fully set forth in this Response.

Tesla objects to this Request as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including how Tesla handles output data related to GPU computations.

Tesla objects to this Request on the grounds that it is overbroad, unduly burdensome, unnecessary, oppressive, and grossly disproportionate to the needs of this case. The Request seeks documents showing how output data, intermediate results, tensors, buffers, activations, variables, and "other computation results" are "stored, transferred, copied, streamed, written back, returned, accumulated, reused, or made available (including asynchronously)" of unbounded scope, including different memory domains such as "GPU memory to CPU memory, host memory, system memory, storage, display, network, or another memory location" and then extends the Request to encompass data movement "in the opposite direction" as well. Tesla objects to the catch-all phrases "or other computation results" and "or another memory location" as vague, ambiguous, and unbounded, rendering the Request limitless in scope. Tesla objects to this Request to the extent that this catch-all request would potentially require Tesla to produce an unknowable quantum of engineering documentation representing Tesla's valuable

intellectual property. The Request fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation.

Tesla further objects to this Request to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case. Tesla objects to this Request to the extent it seeks the disclosure of Tesla's confidential trade secrets and proprietary intellectual property, including GPU-CPU data transfer strategies, asynchronous computing pipelines, custom memory management implementations, and real-time inference data flow architectures. Production of such materials in a dispute between Neural AI and NVIDIA would risk competitive harm to Tesla, a non-party.

Tesla objects to this Request to the extent it uses generic technical terminology to describe fundamental and ubiquitous computing operations that are not unique to any particular proprietary technology, patented method, or party to this litigation. The Request fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation. As drafted, this Request is therefore facially overbroad, unduly burdensome, and disproportionate to the needs of this case, particularly as directed to a non-party.

Tesla further objects to the extent this Request seeks to impose a duty on Tesla to create materials, such as bidirectional data flow diagrams, memory transfer analyses, or asynchronous pipeline documentation, that Tesla does not create or maintain in the ordinary course of business. Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine.

Tesla further objects on the grounds that the data transfer mechanisms described, GPU-to-CPU transfers, asynchronous memory operations, streaming, write-back, are standard computing operations extensively documented in NVIDIA's public CUDA programming

guides and developer resources and are more appropriately sought from the Defendant. Tesla further objects to the extent that this Request is substantially duplicative of prior Requests, including Requests Nos. 4, 5, 7, 8, and 9.

Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to this Request to the extent that it calls for documents not within Tesla's possession, custody or control or kept in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Request and the burden it imposes on Tesla.

DOCUMENT REQUEST NO. 11:

Documents sufficient to show how computations You perform using NVIDIA GPUs are scheduled, ordered, controlled, queued, synchronized, parallelized, launched, interrupted, resumed, or executed, including through kernels, CUDA streams, CUDA graphs, events, threads, controllers, schedulers, compilers, runtimes, inference engines, run lists, run engines, or custom software.

RESPONSE TO DOCUMENT REQUEST NO. 11:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions and Instructions, as though fully set forth in this Response.

Tesla objects to this Request as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including "how computations [Tesla] perform[s] using NVIDIA GPUs are scheduled, ordered, controlled, queued, synchronized, parallelized, launched, interrupted, resumed, or executed, including through kernels, CUDA streams, CUDA graphs, events, threads, controllers, schedulers, compilers, runtimes, inference engines, run lists, run engines, or custom software."

Tesla objects to this Request on the grounds that it is facially overbroad, unduly burdensome, unnecessary, oppressive, and not proportional to the needs of this case. The Request seeks documents showing how Tesla's GPU computations are "scheduled, ordered, controlled, queued, synchronized, parallelized, launched, interrupted, resumed, or executed" through an exhaustive and open-ended list of mechanisms, "kernels, CUDA streams, CUDA graphs, events, threads, controllers, schedulers, compilers, runtimes, inference engines, run lists, run engines, or custom software." Tesla objects to the inclusion of "custom software" as a catch-all that would require disclosure of Tesla's proprietary computing systems in their entirety. Tesla objects to this Request to the extent that this catch-all request would potentially require Tesla to produce an unknowable quantum of engineering documentation representing Tesla's valuable intellectual property. The Request fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation.

Tesla further objects to this Request to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case. Tesla objects to this Request to the extent it seeks the disclosure of Tesla's confidential trade secrets and proprietary intellectual property, including custom GPU scheduling systems, proprietary inference engine designs, compute orchestration strategies, compiler optimizations, and proprietary runtime environments that constitute core competitive advantages. Production of such materials in a dispute between Neural AI and NVIDIA would risk competitive harm to Tesla, a non-party.

Tesla further objects to the extent this Request seeks to impose a duty on Tesla to create materials, such as scheduling architecture analyses, execution flow documentation, or orchestration diagrams, that Tesla does not create or maintain in the ordinary course of business.

Tesla objects to this Request to the extent it uses generic technical terminology to describe fundamental and ubiquitous computing operations that are not unique to any particular proprietary technology, patented method, or party to this litigation. The Request fails to

identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation. As drafted, this Request is therefore facially overbroad, unduly burdensome, and disproportionate to the needs of this case, particularly as directed to a non-party.

Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla further objects on the grounds that information regarding CUDA streams, CUDA graphs, kernel launching, GPU scheduling, and compute execution is extensively documented in NVIDIA's public developer documentation, CUDA programming guides, and technical specifications and is more appropriately sought from the Defendant. Tesla further objects to the extent that this Request is substantially duplicative of prior Requests, including Requests Nos. 1, 4, 5, 7, and 8.

Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to this Request to the extent that it calls for documents not within Tesla's possession, custody or control or kept in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Request and the burden it imposes on Tesla.

DOCUMENT REQUEST NO. 12:

Documents sufficient to show whether and how user inputs, user commands, configuration changes, parameter changes, model changes, computational-element changes, input changes, interruptions, or display/output changes affect computations You perform using NVIDIA GPUs and/or queue them for GPU computation.

RESPONSE TO DOCUMENT REQUEST NO. 12:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions and Instructions, as though fully set forth in this Response.

Tesla objects to this Request as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including “whether and how user inputs, user commands, configuration changes, parameter changes, model changes, computational-element changes, input changes, interruptions, or display/output changes affect computations [Tesla] perform[s] using NVIDIA GPUs and/or queue them for GPU computation.”

Tesla objects to this Request on the grounds that it is overbroad, unduly burdensome, unnecessary, oppressive, and not proportional to the needs of this case. The Request seeks documents showing "whether and how" an open-ended and effectively limitless list of changes, "user inputs, user commands, configuration changes, parameter changes, model changes, computational-element changes, input changes, interruptions, or display/output changes," affect Tesla's GPU computations or queue them for processing. This scope is virtually unlimited. Tesla objects to this Request to the extent that this catch-all request would potentially require Tesla to produce an unknowable quantum of engineering documentation representing Tesla's valuable intellectual property. The Request fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation.

Tesla objects to the terms "computational-element changes," "input changes," and "display/output changes" as vague, ambiguous, and undefined. It is unclear what constitutes a "computational-element change" or how broadly "input changes" is intended to sweep.

Tesla further objects to this Request to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case. Tesla objects to this Request to the extent it seeks the disclosure of Tesla's confidential trade secrets and proprietary technical

information, including user interface architectures, real-time computing systems, dynamic GPU scheduling implementations, model update pipelines, interactive inference systems, and human-machine interface designs. Production of such materials in a dispute between Neural AI and NVIDIA would risk competitive harm to Tesla, a non-party.

Tesla further objects to the extent this Request seeks to impose a duty on Tesla to create materials, such as user interaction flow analyses, change-impact documentation, or input-to-computation mapping diagrams, that Tesla does not create or maintain in the ordinary course of business. Tesla objects to this Request to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine.

Tesla objects to this Request to the extent it uses generic technical terminology to describe fundamental and ubiquitous computing operations that are not unique to any particular proprietary technology, patented method, or party to this litigation. The Request fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation. As drafted, this Request is therefore facially overbroad, unduly burdensome, and disproportionate to the needs of this case, particularly as directed to a non-party.

Tesla further objects on the grounds that the computing concepts described in this Request, user input handling, parameter configuration, compute queuing, are standard GPU computing operations documented in NVIDIA's public developer resources and are more appropriately sought from the Defendant. Tesla further objects to the extent that this Request is substantially duplicative of prior Requests, including Requests Nos. 1, 4, 5, 7, and 11.

Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to this Request to the extent that it calls for documents not within Tesla's possession, custody or control or kept in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Request and the burden it imposes on Tesla.

OBJECTIONS AND RESPONSES TO DEPOSITION TOPICS

DEPOSITION TOPIC NO. 1:

The NVIDIA software and libraries You use to perform computations, including but not limited to NVIDIA's Aerial, Clara Parabricks, cuBLAS, cuDNN, cuFFT, cuQuantum, cuSOLVER, cuSPARSE, Drive, DriveWorks, Holoscan, Isaac, Isaac Lab, Maxine, Memory Map, Merlin, Metropolis, Modulus, Monai, Morpheus, NeMo, PyTorch, RAPIDS, Riva, Runtime Driver, TensorFlow, TensorRT, Triton, VSS (Deepstream).

RESPONSE TO DEPOSITION TOPIC NO. 1:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Topic as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including Tesla's use of a long list of software and libraries. Tesla further objects to this Topic as vague and ambiguous as to what "computations" are relevant to the underlying litigation.

Tesla objects to this Topic on the grounds that it is overbroad, unduly burdensome, unnecessary, and not proportional to the needs of this case. The Topic identifies almost thirty NVIDIA software and libraries, placing undue burden on Tesla, a non-party.

Tesla further objects to this Topic to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case. Tesla objects to this Topic to the extent it seeks the disclosure of Tesla's highly confidential and proprietary information regarding its internal software stack, computing infrastructure, and technology choices, which constitute competitively sensitive business information. Disclosure of which specific NVIDIA tools Tesla does or does not use would reveal Tesla's internal technology strategy and competitive posture.

Tesla further objects to this Topic on the grounds that information regarding NVIDIA's software products and their usage by customers is more appropriately sought from NVIDIA, the Defendant in this action, which possesses licensing records, telemetry data, and customer usage information.

Tesla further objects to this Topic on the grounds that the information sought is more readily obtainable from the parties to this litigation or from other less burdensome sources, thus disproportionate to the needs of the case.

Tesla objects to this Topic to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to the extent this Topic seeks to impose a duty on Tesla to create materials or compile analyses that Tesla does not create or maintain in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Topic and the burden it imposes on Tesla.

DEPOSITION TOPIC NO. 2:

The NVIDIA sample Source Code You use, in whole or in part, to conduct computations.

RESPONSE TO DEPOSITION TOPIC NO. 2:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Topic as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action. Tesla further objects to this Topic as vague and ambiguous as to what "computations" are relevant to the underlying litigation.

Tesla objects to this Topic on the grounds that it is overbroad, unduly burdensome, unnecessary, and not proportional to the needs of this case. The term "sample Source Code" is vague and ambiguous, as it is unclear what specific sample code is being referenced, what universe of NVIDIA sample code is at issue, or how Tesla would be expected to identify whether any portion of its codebase derives from, incorporates, or was inspired by NVIDIA sample code. Tesla objects to this Topic to the extent that compliance would require Tesla to conduct an exhaustive audit of its entire codebase against an undefined body of NVIDIA sample code, a task that is extraordinarily burdensome and disproportionate for a non-party.

Tesla further objects to this Topic to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case. Tesla objects to this Topic to the extent it seeks the disclosure of Tesla's proprietary source code and internal software, which constitute trade secrets and core intellectual property. Tesla further objects on the grounds that NVIDIA sample source code is publicly available and/or in NVIDIA's possession, custody, or control, and information regarding its distribution to customers is more appropriately sought from the Defendant.

Tesla further objects to this Topic on the grounds that the information sought is more readily obtainable from the parties to this litigation or from other less burdensome sources, thus disproportionate to the needs of the case.

Tesla objects to this Topic to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to the extent this Topic seeks to impose a duty on Tesla to create materials or compile analyses that Tesla does not create or maintain in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Topic and the burden it imposes on Tesla.

DEPOSITION TOPIC NO. 3:

Your customizations and/or data inputs to NVIDIA software that alter the way in which NVIDIA software performs computations and/or a description of the data input to NVIDIA software on which computations are run.

RESPONSE TO DEPOSITON TOPCI NO. 3:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Topic as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including how Tesla handles input data related to GPU computations. Tesla further objects to this Topic as vague and ambiguous as to what "computations" are relevant to the underlying litigation. Tesla further objects to this Topic to the extent it seeks information beyond the use of NVIDIA products.

Tesla objects to this Topic to the extent that it is overbroad, unduly burdensome, unnecessary, oppressive, and not proportional to the needs of this case. Tesla objects to the Topic to the extent that it seeks information showing Tesla's "customizations" or "alter the way in which NVIDIA software performs computations" without providing specificity as to what scope of customizations or alterations are relevant to the underlying action. The Topic fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation. As drafted, this Topic would extend to the entirety of Tesla's computing operations and is therefore overbroad, unduly burdensome, and disproportionate to the needs of this case, particularly as directed to a non-party.

Tesla further objects to this Topic to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case.

Tesla objects to this Topic to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d). Tesla further objects to the extent this Topic seeks to impose a duty on Tesla to create materials or compile analyses that Tesla does not create or maintain in the ordinary course of business.

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Topic and the burden it imposes on Tesla.

DEPOSITION TOPIC NO. 4:

Identification of Your software that uses NVIDIA GPUs to perform computations.

RESPONSE TO DEPOSITION TOPIC NO. 4:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Topic as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including Tesla or third-party software that uses NVIDIA GPUs. Tesla further objects to this Topic as vague and ambiguous as to what "computations" are relevant to the underlying litigation. Tesla further objects to this Topic to the extent it seeks information beyond the use of NVIDIA products.

Tesla objects to this Topic on the grounds that it is overbroad, unduly burdensome, unnecessary, oppressive, and not proportional to the needs of this case, particularly given that Tesla is not a party to this litigation. The Topic seeks identification of software used to perform

computations on NVIDIA GPUs without any specificity, a scope that could potentially encompass virtually any software system Tesla operates. The Topic fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation. As drafted, this Topic would extend to the entirety of Tesla's computing operations and is therefore overbroad, unduly burdensome, and disproportionate to the needs of this case, particularly as directed to a non-party.

Tesla objects to this Topic to the extent it seeks Tesla confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case. Tesla objects to this Topic to the extent it seeks the disclosure of Tesla confidential trade secrets and proprietary technical information, including proprietary source code, internal software architectures, and custom-built AI/ML frameworks that constitute core intellectual property and are not relevant to this case. Disclosure of such information, particularly on a third-party, would risk competitive harm to Tesla, a non-party.

Tesla further objects to the extent this Topic seeks to impose a duty on Tesla to create materials or compile analyses that Tesla does not create or maintain in the ordinary course of business. Tesla objects to this Topic to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d).

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Topic and the burden it imposes on Tesla.

DEPOSITION TOPIC NO. 5:

Using Your software, the ways in which output data from a GPU computation(s), including intermediate results of GPU computations are stored, referenced by a pointer, transferred, copied, streamed, written back, returned, accumulated, reused, or made available

including asynchronously from NVIDIA GPU memory to CPU memory, host memory, system memory, storage, display, network, or another memory location before, during, or in parallel with computations performed using NVIDIA GPUs

RESPONSE TO DEPOSITON TOPIC NO. 5:

Tesla restates and incorporates its Preliminary Statement, General Objections, Objections to Definitions, and Objections to Instructions, as though fully set forth in this Response.

Tesla objects to this Topic as irrelevant to the extent it seeks information that lacks any connection to any specific claims or defenses at issue in the underlying action, including how Tesla handles output data related to GPU computations. Tesla further objects to this Topic as vague and ambiguous as to what "computations" are relevant to the underlying litigation. Tesla further objects to this Topic to the extent it seeks information beyond the use of NVIDIA products.

Tesla objects to this Topic on the grounds that it is overbroad, unduly burdensome, unnecessary, oppressive, and grossly disproportionate to the needs of this case. Tesla objects to this Topic to the extent it uses generic technical terminology to describe fundamental and ubiquitous computing operations that are not unique to any particular proprietary technology, patented method, or party to this litigation. The Topic fails to identify with sufficient particularity what specific implementation or application of these techniques is at issue in the underlying litigation, rendering meaningful compliance impossible without speculation. As drafted, this Topic would extend to the entirety of Tesla's computing operations and is therefore overbroad, unduly burdensome, and disproportionate to the needs of this case, particularly as directed to a non-party.

Tesla further objects to this Topic to the extent it seeks documents containing confidential, proprietary or trade secret information without making any showing of relevance or that the request is proportional to the needs of the case. Tesla objects to this Topic to the extent it seeks the disclosure of Tesla's confidential trade secrets and proprietary intellectual

property, including GPU-CPU data transfer strategies. Disclosure of such confidential materials in a dispute between Neural AI and NVIDIA would risk competitive harm to Tesla, a non-party.

Tesla further objects to the extent this Topic seeks to impose a duty on Tesla to create materials or compile analyses that Tesla does not create or maintain in the ordinary course of business. Tesla objects to this Topic to the extent it seeks information protected by the attorney-client privilege or the work-product doctrine. Tesla further objects to this Request as Neural AI has failed to comply with the requirements of Fed. R. Civ. P. 45(d).

Subject to and without waiving the foregoing general and specific objections, Tesla is willing to meet and confer regarding the scope of this Topic and the burden it imposes on Tesla.

Dated: July 21, 2026

/s/ Jun Zheng_____

Jun Zheng
TX Bar No. 24102681
zhengjun@tesla.com
Tesla, Inc.
1 Tesla Rd
Austin, TX 78725
(512) 417-3528

Gina H. Cremona
CA Bar No. 305392
gcremona@tesla.com
Tesla, Inc.
1501 Page Mill Rd.
Palo Alto, CA 94304
(650) 647-0015

Ashraf Fawzy
DC Bar No. 989132
afawzy@tesla.com
Tesla, Inc.
800 Connecticut Ave. NW
Washington, DC 20006
(202) 905-9221

Attorneys for TESLA, INC.

CERTIFICATE OF SERVICE

I hereby certify that a true copy of the above document was served upon Plaintiff Neural AI's counsel of record via electronic mail on July 21, 2026.

/s/ Jun Zheng
Jun Zheng