

Tesla, Inc. (Petitioner)
v.
Autonomous Devices, LLC (Patent Owner)
Petitioner's Demonstratives

Case No. IPR2024-00745
U.S. Patent No. 11,663,474

Before Honorable Administrative Patent Judges
Terrence W. McMillin, Steven M. Amundson, and Russell E. Cass

Table of Contents

Overview of Challenged Claims	3
Unpatentability of Claims 56-57	9
Overview of the Prior Art	10
Obviousness.....	15
Unpatentability of the Substitute Claims (MTA).....	31
Overview of the Additional Prior Art	32
Prior Art Combinations	34
Obviousness.....	40
Lack of Written Description Support.....	69

Overview of Challenged Claims

Challenged Claims of the '474 Patent

- **Petition Grounds (IPR2024-00745)**

- **1A:** (§103) Buibas and Sinyavskiy render obvious claims 45-47, 51, 56-57, 59, 68-69
- **1B:** (§103) Buibas, Sinyavskiy, and Versace render obvious claim 60
- **2A:** (§102) Grotmol anticipates claims 45-47, 51, 59, 68-69
- **2B:** (§103) Grotmol and McClure render obvious claim 60
- **2C:** (§103) Grotmol and Buibas render obvious claims 45-47, 51, 56-57, 59, 68-69
- **2D:** (§103) Grotmol, Buibas, and McClure render obvious claim 60
- Patent Owner did not contest Petitioner's challenges to claims 45-47, 51, 59, and 68-69 in the Patent Owner Response, arguing only dependent claims **56-57**

- **Opposition Grounds (IPR2024-00745)**

- **MTA-1A:** (§103) Grotmol and Zhu render obvious substitute claims 85-91
- **MTA-1B:** (§103) Grotmol, Zhu, and Buibas render obvious substitute claims 85-91

Overview of the '474 Patent

(12) United States Patent Cosic

(10) Patent No.: US 11,663,474 B1
(45) Date of Patent: *May 30, 2023

(54) ARTIFICIALLY INTELLIGENT SYSTEMS, DEVICES, AND METHODS FOR LEARNING AND/OR USING A DEVICE'S CIRCUMSTANCES FOR AUTONOMOUS DEVICE OPERATION

(71) Applicant: Jasmin Cosic, Miami, FL (US)

(72) Inventor: Jasmin Cosic, Miami, FL (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.
This patent is subject to a terminal disclaimer.

(21) Appl. No.: 17/561,976

(22) Filed: Dec. 26, 2021

Related U.S. Application Data

(63) Continuation of application No. 16/540,972, filed on Aug. 14, 2019, now Pat. No. 11,238,344, which is a continuation of application No. 15/340,991, filed on Nov. 2, 2016, now Pat. No. 10,452,974.

(51) Int. Cl.
G06N 3/08 (2023.01)
G06N 5/022 (2023.01)

(52) U.S. Cl.
CPC G06N 3/08 (2013.01); G06N 5/022 (2013.01)

(58) Field of Classification Search
CPC G06N 5/022; G06N 3/08
See application file for complete search history.

(56) References Cited

U.S. PATENT DOCUMENTS

4,370,707 A 1/1983 Phillips et al.
4,730,315 A 3/1988 Saito et al.

4,860,203 A 8/1989 Corrigan et al.
5,560,611 A 9/1996 Uyama
5,602,982 A 2/1997 Judl et al.
5,874,834 A 2/1999 New
6,026,234 A 2/2000 Hanson et al.
6,088,731 A 7/2000 Kiraly et al.
6,106,299 A 8/2000 Ackermann et al.
6,126,330 A 10/2000 Knight
6,314,558 B1 11/2001 Angel et al.
6,408,262 B1 6/2002 Leerberg et al.
6,643,842 B2 11/2003 Angel et al.
6,728,689 B1 4/2004 Drissi et al.
6,735,632 B1 5/2004 Kiraly et al.
6,754,631 B1 6/2004 Din
6,801,912 B2 10/2004 Moskowitz et al.
6,842,877 B2 1/2005 Roberts et al.
6,915,105 B2 7/2005 Masuda
6,973,446 B2 12/2005 Mamitsuka et al.
6,996,758 B1 2/2006 Herron et al.

(Continued)

OTHER PUBLICATIONS

Tracing (software), retrieved from <URL: http://wikipedia.com> on Jan. 10, 2014, 3 pages.

(Continued)

Primary Examiner — Santiago Garcia

(74) Attorney, Agent, or Firm — Imperium Patent Works LLP

(57) ABSTRACT

Aspects of the disclosure generally relate to computing enabled devices and/or systems, and may be generally directed to devices, systems, methods, and/or applications for learning a device's operation in various circumstances, storing this knowledge in a knowledgebase (i.e. neural network, graph, sequences, etc.), and enabling autonomous operation of the device.

84 Claims, 40 Drawing Sheets





FIG. 32

Autonomous operation

TESLA-1001, FIG. 32 (cited and annotated in Pet., p. 4)

Overview of the '474 Patent




FIG. 38

TESLA-1001, FIG. 38 (cited in Reply, p. 2)

In some embodiments, DCADO Unit 100 may reside on Server 96 accessible over Network 95 as previously described. In such embodiments, any number of Loaders 98a may connect to such remote DCADO Unit 100 and the remote DCADO Unit 100 may learn their operations in circumstances including objects with various properties. In turn, any number of Loaders 98a can utilize the remote DCADO Unit 100 for autonomous operation in circumstances including objects with various properties. For example, multiple operators (i.e. Users 50, etc.) may operate their Loaders 98a that are configured to transmit their operations in circumstances including objects with various properties to a remote DCADO Unit 100. Such remote DCADO Unit 100 enables learning of the operators' collective knowledge of operating Loaders 98 in circumstances including objects with various properties. Any number of Loaders 98 can utilize such collective knowledge comprised in the remote DCADO Unit 100 for their autonomous operation. Any of the disclosed elements such as Artificial Intelligence Unit 110, Knowledgebase 530, and/or others may reside on Server 96, and any combination of local and remote elements can be implemented in alternate embodiments.

TESLA-1001, 163:16-38 (cited in MTA Reply, p. 4)

Independent Claim 45 and Dependent Claim 46

Claim 45

45. A first device comprising:
a knowledgebase that includes one or more inputs for inputting at least a portion of a circumstance representation, wherein the one or more inputs are correlated with one or more instruction sets for operating a second device;
one or more sensors;
one or more processors; and
one or more non-transitory machine readable media storing machine readable code that, when executed by the one or more processors, causes the one or more processors to perform at least:
generating a circumstance representation, wherein the generated circumstance representation represents a circumstance detected at least in part by the one or more sensors;
determining the one or more instruction sets for operating the second device at least by: inputting at least a portion of the generated circumstance representation into the one or more inputs, and using a correlation between the one or more inputs and the one or more instruction sets for operating the second device; and
at least in response to the determining, executing the one or more instruction sets for operating the second device, wherein the first device autonomously performs one or more operations defined by the one or more instruction sets for operating the second device.

Claim 46

46. The first device of claim 45, wherein the one or more inputs for inputting the at least the portion of the circumstance representation include one or more inputs for inputting at least a portion of one or more object representations, and wherein the generated circumstance representation includes one or more object representations, and wherein the first device is a first vehicle, and wherein the second device is a second vehicle.

Note: Claims 56-57 both depend from independent claim 45 and dependent claim 46.

Dependent Claims 56 and 57

Claim 56

56. The first device of claim 46, wherein the one or more inputs are further correlated with another one or more instruction sets for operating a third device, and wherein at least a portion of the one or more instruction sets for operating the second device or at least an information related to the one or more instruction sets for operating the second device is learned in a learning process that includes operating the second device at least partially by a user, and wherein at least a portion of the another one or more instruction sets for operating the third device or at least an information related to the another one or more instruction sets for operating the third device is learned in another learning process that includes operating the third device at least partially by another user.

Claim 57

57. The first device of claim 46, wherein the one or more inputs are further correlated with another one or more instruction sets for operating a third device.

Part 1: Unpatentability of Claims 56-57

Overview of the Prior Art

Petition

Overview of Buibas




FIG. 10

TESLA-1005, FIG. 10 (cited in Pet., p. 10)

Apparatus and methods for learning and training in neural network-based devices. In one implementation, the devices each comprise multiple spiking neurons, configured to process sensory input. In one approach, alternate heterosynaptic plasticity mechanisms are used to enhance learning and field diversity within the devices. The selection of alternate plasticity rules is based on recent post-synaptic activity of neighboring neurons. Apparatus and methods for simplifying training of the devices are also disclosed, including a computer-based application. A data representation of the neural network may be imaged and transferred to another computational environment, effectively copying the brain. Techniques and architectures for achieve this training, storing, and distributing these data representations are also disclosed.

TESLA-1005, Cover (cited in Pet., p. 14)

Overview of Sinyavskiy




FIG. 6A

TESLA-1006, FIG. 6A (cited in Pet., p. 12)

Generalized learning rules may be implemented. A framework may be used to enable adaptive spiking neuron signal processing system to flexibly combine different learning rules (supervised, unsupervised, reinforcement learning) with different methods (online or batch learning). The generalized learning framework may employ time-averaged performance function as the learning measure thereby enabling modular architecture where learning tasks are separated from control tasks, so that changes in one of the modules do not necessitate changes within the other. Separation of learning tasks from the control tasks implementations may allow dynamic reconfiguration of the learning block in response to a task change or learning method change in real time. The generalized spiking neuron learning apparatus may be capable of implementing several learning rules concurrently based on the desired control application and without requiring users to explicitly identify the required learning rule composition for that task.

TESLA-1006, FIG. 6A (cited in Pet., p. 11)

Overview of Grotmol




FIG. 1

TESLA-1008, FIG. 1 (cited in Pet., p. 53)

An apparatus and methods for training and/or operating a robotic device to follow a trajectory. A robotic vehicle may utilize a camera and stores the sequence of images of a visual scene seen when following a trajectory during training in an ordered buffer. Motor commands associated with a given image may be stored. During autonomous operation, an acquired image may be compared with one or more images from the training buffer in order to determine the most likely match. An evaluation may be performed in order to determine if the image may correspond to a shifted (e.g., left/right) version of a stored image as previously observed. If the new image is shifted left, right turn command may be issued. If the new image is shifted right then left turn command may be issued.

TESLA-1008, Cover (cited in Pet., p. 52)

Overview of Grotmol (cont'd)




FIG. 10B

TESLA-1008, FIG. 10B (cited in Pet., p. 53)

Obviousness of Claims 56-57

PO Raises Limited Rebuttals to the Petition's Challenges of Claims 56-57

Claim 56

56. The first device of claim 46, wherein the one or more inputs are further correlated with another one or more instruction sets for operating a third device, and wherein at least a portion of the one or more instruction sets for operating the second device or at least an information related to the one or more instruction sets for operating the second device is learned in a learning process that includes operating the second device at least partially by a user, and wherein at least a portion of the another one or more instruction sets for operating the third device or at least an information related to the another one or more instruction sets for operating the third device is learned in another learning process that includes operating the third device at least partially by another user.

Claim 57

57. The first device of claim 46, wherein the one or more inputs are further correlated with another one or more instruction sets for operating a third device.

- No dispute that Buibas and Grotmol disclose a “mixed” knowledgebase having knowledge obtained from multiple devices (e.g., a second device and a third device)
- No dispute that prior art discloses “one or more inputs” of a knowledgebase correlated with instruction sets for operating a second device
- PO instead disputes the Petition’s reliance on Buibas as rendering obvious:
 - (i) “the [same] one or more inputs” are “further correlated” with instruction sets for operating a third device; and
 - (ii) the instruction sets for operating a third device are learned in “another learning process” that includes operating the third device by another user.

Reply, 1-3

The '474 Patent Discloses Nothing Significant About Correlating Instruction Sets for Operating Multiple Devices with the Same Inputs

Dr. Saber:

69. Second, in my opinion, Petitioner's contention of combining capabilities by a network merge (i.e. merging weight vectors as Buibas specifically defines a "network merge" as discussed above) is not, and indeed does not even mention, further correlating the same one or more inputs with another instruction set, let alone an instruction set for operating a third device as required by claim 56. Merging two weight vectors as Buibas specifically discloses in its "Network Merges" involves combining weights of weight vectors and does not involve any instruction sets, much less *further correlating the same*

EX-2001, ¶69

The '474 Patent:

Referring to FIG. 17, the disclosed artificially intelligent devices, systems, and methods for learning and/or using a device's circumstances for autonomous device operation may include various artificial intelligence models and/or techniques. The disclosed systems, devices, and methods are independent of the artificial intelligence model and/or technique used and any model and/or technique can be utilized to facilitate the functionalities described herein. Examples of these models and/or techniques include deep learning, supervised learning, unsupervised learning, neural networks (i.e. convolutional neural network, recurrent neural network, deep neural network, etc.), search-based, logic and/or fuzzy logic-based, optimization-based, tree/graph/other data structure-based, hierarchical, symbolic and/or sub-symbolic, evolutionary, genetic, multi-agent, deterministic, probabilistic, statistical, and/or other models and/or techniques.

TESLA-1001, 103:44-55 (cited in Reply, p. 4)


Dr. Janet:

Based on my review of the specification and file history, I likewise identified no such criticality or significance. Apart from the benefits of a mixed knowledgebase generally (which Buibas discloses), the '474 Patent identifies nothing unpredictable or uniquely advantageous about a *single* set of inputs of the knowledgebase being correlated with *multiple* instruction sets. On the contrary, the '474 Patent concedes that the particular model or structure of the knowledgebase employed is *independent* of the disclosed functions, and hence, any suitable model can be used to carry out the functions of the '474 Patent. TESLA-1001, 103: 44-55 ("The disclosed systems, devices, and methods are *independent* of the artificial intelligence model and/or technique used and any model and/or technique can be utilized to facilitate the functionalities described herein.").

TESLA-1022, ¶12 (cited in Reply, p. 4)

Buibas's Disclosure of Knowledgebase Inputs Correlated with Instruction Sets for Operating Multiple Devices is at Least as Extensive as the '474 Patent's


Buibas Discloses Neural Networks Having Input Neurons Correlated with Instruction Sets Produced by an Output Layer



Buibas, FIG. 10

Pet., 22

PO and Dr. Saber Identify the Basic Feedforward Neural Network Depicted in FIG. 17 of the '474 Patent as Providing Support for a Knowledgebase Having Input Neurons Correlated With Instruction Sets for Operating Multiple Devices




'474 Patent, FIG. 17 (Model A)

Reply, 4-5; MTA, 9-11; EX2001, 79-80, 113, 127, 139, 141, 154, 164, 192, 203

Buibas's Merge Operations Combine Fully-Defined Brain Images and Produce an Offspring Whose Inputs are Correlated with Instruction Sets for Operating Multiple Devices

Buibas:



TESLA-1005, FIG. 4 (cited in Pet., p. 43)

In some implementations, the network merge may comprise non-interleaved combination as illustrated by the merges 430 in FIG. 4, where the parents A provides the portion 432 and the parent B provides the portion 434 to the off-spring 436. In some implementations (not shown), the portion 432 may comprises the same number of weights as the portion 434. **The non-interleaved network merge, such as illustrated in FIG. 4, may be used, for example, to combine capabilities of two (or more) parents or to add a capability to one of the parent network.** TESLA-1005, ¶[0096].

TESLA-1005, [0096] (cited in TESLA-1003, ¶99)

Dr. Janet (First Declaration):

99. From my analysis of the references, the Buibas-Sinyavskiy combination renders obvious claim [56]. For example, as described above with respect to [45.1], Buibas-Sinyavskiy disclose that the input interface/connections (*one or more inputs*) of a partition of a brain image (*knowledgebase*) for a particular task or behavior are correlated with the instruction sets produced by that partition. See above, element [45.1]. Buibas further discloses techniques for “merging” brain images (*the one or more inputs are further correlated with another one or more instruction sets for operating a third device*) to combine behavior modules learned on two or more robots (a *second device* and a *third device*). *Id.* For example, “non-interleaved” network merges can be used “to combine capabilities of two (or more) parents or to add a capability to one of the parent network [*sic*].” TESLA-1005, ¶[0096], FIG. 4; *generally id.*, ¶¶[0090]-[0098], [0124]-[0127], [0189], [0225], FIG. 9. Accordingly, it would have been understood and obvious to a POSITA that the parents to a merged network that forms a partition of a brain image on the first device can be trained on separate devices, e.g., a *second device* and a *third device*,

TESLA-1003, ¶99 (cited in Pet, pp. 41-43)

Buibas' Merge Operations Involve Merging More Than Weights of Brain Images

Buibas:

- “[A] user may extract the state of a neural network (or other useful training-related information) directly from the device. Thus, the artificial mind and its traits may be copied, stored, and later retrieved. This state information may be shared with other users.” TESLA-1005, [0070].
- “[P]artial [brain] images may be used to augment or upgrade (downgrade) a pre-loaded brain image or a stored brain image. ... [T]wo or more partial brain images may be combined to form full brain images.” TESLA-1005, [0188].
- “The user may also augment the old brain image with new functionality from a partial new brain. Alternatively, the user may select multiple ones from different new functionalities of partial brain images to create a custom hybrid brain.” TESLA-1005, [0225].

Dr. Janet (Second Declaration):

15. Patent Owner and Dr. Saber focus narrowly on Buibas' description of FIG. 4, which explains that “*In some implementations*, the network lifecycle *may* comprise network *weight updates* comprising two or more sources, also referred herein to as ‘merges.’” TESLA-1005, [0094], FIG. 4; *generally id.*, [0090]-[0115], FIGS. 3-5. While it is true that the discussion of FIG. 4 centers on merging network weights, Buibas never suggests that a network merge operates *only* to combine the *weights* of the parent networks as opposed to the networks themselves. A POSITA would not have read the disclosure of FIG. 4 in isolation without considering Buibas' disclosure in its entirety, which clarifies that not only weights, but “brain images” that fully define the networks for each task can be the subject of a network merge, thereby forming an offspring that combines the capabilities of two or more parent brain images as a whole. See, e.g., TESLA-1005, [0225] (“The user may also augment the old brain image with new functionality from a partial new brain. ... *Users may also merge brain full images The merging of brain images involves the combination of one brain image with another brain image.*”), [0188] (“[T]wo or more *partial brain images may be combined to form full brain images.*”), [0189] “combining brain images”; “form full brain images from partial brain images”), TESLA-1022, ¶15 (cited in Reply, p. 3)

Buibas' Merged Brain Images Correlate the Same One or More Inputs with Instruction Sets for Operating a Second and Third Device

Dr. Saber (First Declaration):

69. Second, in my opinion, Petitioner's contention of combining capabilities by a network merge (i.e. merging weight vectors as Buibas specifically defines a "network merge" as discussed above) is not, and indeed does not even mention, further correlating the same one or more inputs with another instruction set, let alone an instruction set for operating a third device as required by claim 56. Merging two weight vectors as Buibas specifically discloses in its "Network Merges" involves combining weights of weight vectors and does not involve any instruction sets, much less *further correlating the same* EX-2001, ¶¶69

Dr. Janet (Second Declaration):

21. Patent Owner's and Dr. Saber's further assumption that the inputs to a *partitioned* brain image must be considered "separate[ly]" from each other is also illogical and inconsistent with Buibas' express disclosure. Buibas depicts a partitioned neuromorphic apparatus in FIG. 11D, for example, that includes just a single input interface 1160 and a single output interface 1170. The input interface 1160 is correlated by way of downstream connections to the respective outputs of each partition, which are all further communicated to the output interface 1170. The outputs of each partition would include instruction sets for operating a first parent (e.g., a second device) and additional instruction sets for operating a second parent (e.g., a third device) in the scenario of a non-interleaved merge, and in this case, the input interface 1160 is naturally correlated with the instruction sets for operating both devices consistent with the features recited in claims 56 and 57:

TESLA-1022, ¶¶21 (cited in Reply, p. 12)

Buibas' Merged Brain Images Correlate the Same One or More Inputs with Instruction Sets for Operating a Second and Third Device (cont'd)

Buibas:




FIG. 11D

TESLA-1005, FIG. 11D (cited in Reply, p. 13)

Buibas' Merged Brain Images Correlate the Same One or More Inputs with Instruction Sets for Operating a Second and Third Device (cont'd)

Dr. Saber's testimony is inconsistent with Patent Owner's incorrect assumption that the inputs to a partitioned brain image must be considered separately such that they could not collectively provide the claimed "one or more inputs"

Dr. Saber (First Deposition):

Q. So when the system is determining the instruction sets for operating a third device, it uses the same one or more input neurons that it used for determining the instruction sets for operating the second device; right?

A. Well, it's referring to the one or more input neurons, but obviously, you would do this intelligently. You wouldn't -- you know, it depends on what the input neurons are. You know, you would do this intelligently.

But, yes, it is using a correlation between the one or more input neurons to the neural network. So there may be a hundred inputs, but you might be using 80 of those for something and 60 for another.

TESLA-1028, 103:13-104:5 (cited in Reply, p. 12)

Buibas' Merge Operations Would Result in Operable Networks that Produce Instruction Sets for Operating a Second Device and Third Device

Dr. Saber (First Declaration):

77. Further, additional instruction sets cannot simply be added to existing instruction sets that are already correlated with existing same inputs in a knowledgebase (i.e. neural network) without appropriately retraining the network. As of the critical date (see my discussion above at Sections III and V), additional instruction sets could only be added to existing instruction sets that are already correlated with existing inputs by appropriately retraining the knowledgebase with training sets that include the additional instructions sets.

This is so because, during retraining, proper weights for correlations/connections are calculated/adjusted through backpropagation process (see Ex. 1001, 18:64-67) for specific instruction sets associated with specific same inputs in each training set. Backpropagation, generally speaking, typically involves an iterative process of training a neural network where a training input is provided to a given network and the corresponding calculated output is computed via feed forward operations through the network. The error between the calculated and expected

EX-2001, ¶77

Dr. Janet (Second Declaration):

25. To start, I note that this argument does not apply to non-interleaved merges in which the system would produce a partitioned brain image comprised of individual parent networks integrated into separate portions of the offspring image. TESLA-1005, [0092]-[0098], [0225], FIG. 4; above, §II.A.3. For a partitioned offspring, no changes to the weights, neurons, or topology are necessary to carry each parent network into the offspring, as each partition simply corresponds to a parent network, and each partition provides a substantially identical structure and capability to its parent. Further, the offspring would provide one or more inputs

26. Even for interleaved merges that combine the parents into a single (non-partitioned) network, Buibas explains that the offspring is still operable to provide a combined set of capabilities of the parents without re-training. See TESLA-1005, [0103] (“utilize interleaved merge to produce an off-spring exhibiting combined set of capabilities”); [0099]-[0115], [0225], FIG. 5. Patent Owner’s suggestion that re-training would be necessary add instruction sets is belied by Buibas’ detailed explanation of the interleaved network merge that makes it possible to combine the capabilities of the parents even without further training. *Id.* A POSITA would have no reason to doubt Buibas’ representations in this regard. Patent Owner’s TESLA-1022, ¶¶25-26 (cited in Reply, pp. 15-16)

Buibas' Merge Operations Would Result in Operable Networks that Produce Instruction Sets for Operating a Second Device and Third Device (cont'd)

Buibas:

[0103] The interleaved network merge may be used in some implementations to generate new weight vector for a network configured to be use by off-spring configured to implement (a subset) of parent functionality. By way of illustration, two or more trash collector robotic apparatus, employed in a home or an office building, and comprising similar hardware (e.g., sensors and actuators) but trained under different conditions and, hence, comprising different sets of network weights, may utilize interleaved merge to produce an off-spring exhibiting combined set of capabilities.

TESLA-1005, [0103] (cited in Reply, pp. 15-16)

brain image. Thus, a user may revert to a previous brain for any reason. After backup, the old brain image may be entirely replaced with the new full brain image. The user may also augment the old brain image with new functionality from a partial new brain. Alternatively, the user may select multiple ones from different new functionalities of partial brain images to create a custom hybrid brain. Users may also merge brain full images from the options available in the Socialize menu. The merging of brain images involves the combination of one brain image with another brain image. The merge process is may be guided by any of the methods discussed above. Further, two or more successive merges of the same or different types may be used to create new brain images (e.g. the same “father” may be used in two successive merges once with the “mother” and then again with the offspring). In some cases, a full brain image that is incompatible with a given robotic device is made compatible through one or more brain merges and/or augmentations.

TESLA-1005, [0225] (cited in Reply, pp. 15-16)

Merging Weight Vectors Would Still Produce a Brain Image Having One or More Inputs Correlated with Instruction Sets for Operating Second/Third Devices

Dr. Saber (First Declaration):

68. First, a network merge is a term Buibas uses for a merge of weight vectors of the neural networks. The entire section of Buibas titled “*Network Merges*” (Ex. 1005, ¶[0094-0098]) specifically describes merging of weight vectors. Further, Buibas specifically states that “merges” are “network weight updates comprising two or more sources.” Ex. 1005, ¶[0094]. Even Further, Petitioner’s own citations of Buibas clearly show that a network merge is a merge of weight vectors of the neural networks (see Ex. 1005, ¶[0028] stating “FIG. 4 is a functional block diagram illustrating a *network merge* in accordance with one or more implementations” and *showing merging of weight vectors*; see also Ex. 1005, FIG. 4 (showing an example of merging of weight vectors and using the resulting merged weight vector to update weights of a neural network)).

EX-2001, ¶68

Dr. Janet (Second Declaration):

28. Even if one accepts for the sake of argument that Buibas’ merge operation is limited to merging only network weights as Patent Owner argues, Buibas would still render claims 56 and 57 obvious. Patent Owner’s arguments to the contrary are premised on a false dichotomy it creates between “weights” of the neural network and “instruction sets.” See, e.g., POR, 14 (“Thus, merging of weight vectors results in another vector of weights, not instruction sets.”). A POSITA would not have conflated weights for instruction sets, but the framework of Patent Owner’s argument does not reflect how a POSITA would have analyzed the issue or Buibas’ disclosure. Although weights and instruction sets are not the same, they are nonetheless interrelated, particularly in the context of Buibas’ neural networks that are configured to output instruction sets (e.g., motor commands) for operating a device. TESLA-1005, [0153], [0159]. Weights define the strength of connections between successive neurons, while instruction sets are produced at the output. Weights and instruction sets are thus interrelated insofar as the output (instruction sets) produced for a given input depends on the weights of connections between the input and the output. TESLA-1005, [0062], [0073]-[0078], FIGS. 1, 10. A POSITA

TESLA-1022, ¶28 (cited in Reply, pp. 17-18)

Merging Weight Vectors Would Still Produce a Brain Image Having One or More Inputs Correlated with Instruction Sets ... (cont'd)

Dr. Janet (Second Declaration):

29. Moreover, merely because a “merge” at its lowest/most granular level is carried out by combining weight vectors does not mean that the merge does not also provide a combined set of instruction sets at the output of the merged network. All computer functions involve low-level operations that can be sterilely characterized in terms of manipulations to bits, bytes, weights, or the like, but the higher-level effect is typically much more meaningful, as it is here. In particular, as Buibas explains, the effect of a merge operation is to provide a “combined set of capabilities” in the offspring. TESLA-1005, [0092], [0103], [0225]. Accordingly, a POSITA would have understood and found obvious from Buibas that merging weight vectors serves to combine the capabilities of the parents by weighting connections between the input and output neurons such that the output neurons would be configured to produce outputs corresponding to the capabilities of the parents, e.g., instructions sets for operating a second device and additional instruction sets for operating a third device. Additionally, the Patent Owner appears

TESLA-1022, ¶29 (cited in Reply, pp. 18-19)

opinion. EX2003, 77:4-82. Although I indicated that weight vectors would be shared in a network merge, I also clearly indicated that network merges can involve far more than simply sharing weights (consistent with Buibas’ explicit disclosure – e.g., “neuron and synaptic weights, biases, activation functions, topology”). *Id.* (Q: “So you could extract a portion of a neural network and fold that into another neural network ...” A: “... you can pull a weight vector, but that weight vector is not going to tell you much about that neural network. It’s just going to be a block of values.”; Q: “Okay. And so when you pull those weights out, you describe them as a means of correlating input values to output values or instruction sets. So the weights are used in that process; is that correct?” A: “If I pull weights only, it’s a mystery what that would do. You can’t just pull the weights.”).

TESLA-1022, ¶29 (cited in Reply, pp. 18-19)

Merging Weight Vectors Would Still Produce a Brain Image Having One or More Inputs Correlated with Instruction Sets ... (cont'd)

Buibas:

[0103] The interleaved network merge may be used in some implementations to generate new weight vector for a network configured to be use by off-spring configured to implement (a subset) of parent functionality. By way of illustration, two or more trash collector robotic apparatus, employed in a home or an office building, and comprising similar hardware (e.g., sensors and actuators) but trained under different conditions and, hence, comprising different sets of network weights, may utilize interleaved merge to produce an off-spring exhibiting combined set of capabilities.

TESLA-1005, [0103] (cited in Reply, pp. 18-19)

[0153] The neuromorphic apparatus 1130 is also configured to provide output (e.g., an indication of recognized object or a feature, or a motor command, e.g., to zoom/pan the image array) via the interface 1144.

TESLA-1005, [0153] (cited in Reply, pp. 17-18)

[0159] The neuromorphic apparatus 1150 may provide output (e.g., an indication of recognized object or a feature, or a motor command, e.g., to zoom/pan the image array) via the interface 1170. In some implementations, the apparatus 1150 may perform all of the I/O functionality using single I/O block (not shown).

TESLA-1005, [0159] (cited in Reply, pp. 17-18)

Claiming an Additional Correlation to Instruction Sets for Operating a Third Device Amounts to a Patentably Insignificant Duplication of Parts

In re Harza

Insofar as claim 1 is concerned, we agree with the Board of Appeals. The basic structure of this claim is not patentably distinguishable from that taught by Gardner. The only distinction to be found is in the recitation in claim 1 of a plurality of ribs on each side of the web whereas Gardner shows only a single rib on each side of the web. [1] It is well settled that the mere duplication of parts has no patentable significance unless a new and unexpected result is produced, and we are of the [***8] opinion that such is not the case here. The other limitations defined by claim 1 are manifest in Gardner. The webs as recited in the claim and as shown in the reference are similar, and the rib of Gardner, besides being "coextensive in length" with the web, has a neck and "an enlargement at the end * * * [of the neck] remote from * * * [the] web." Further, it is obvious that the ribs of Gardner extend outwardly from the web a distance "less than the width of the web," and that the ribs and web are of the claimed rigidity and flexibility.

In re Harza, 274 F.2d 669 (CCPA 1960) (cited in Pet, pp. 43-44)

Dr. Janet:

100. From my analysis of the references, the Buibas-Sinyavskiy combination renders obvious claim [57] for each of the reasons described above with respect to claim [56]. *See above*, claim [56]. Moreover, the mere duplication of parts has no patentable significance unless a new and unexpected result is produced (i.e., the instruction sets for operating the third device of this claim when compared to the instruction sets for operating the second device of claim 1 produce no unexpected result). Here, the '474 Patent describes nothing at all unexpected about sharing instruction sets from second and third devices as opposed to just a second device, and the additional feature claimed in claim [57] is obvious.

TESLA-1003, ¶100 (cited in Pet, pp. 43-44)

Buibas Discloses and Renders Obvious “Another Learning Process That Includes Operating the Third Device at Least Partially by Another User” (Claim 56)

Dr. Janet (Second Declaration):

32. Based on the teachings of Buibas, a POSITA would have recognized and found obvious that when the parents to a merged brain image are networks that have been trained by different users on different devices (i.e., a second device and a third device), the instruction sets produced by the merged brain image would include instruction sets for operating a second device learned in a learning process that includes operating the second device by a first user and additional instruction sets for operating a third device learned in another learning process that includes operating the third device by another user. Considering that the instruction sets for operating the second device and the instruction sets for operating the third device can pertain to different tasks/behaviors, and that different users can independently train the different tasks/behaviors on different devices, a POSITA would have appreciated that the training routines conducted on the two different devices constitute distinct “learning processes.”

33. Despite its critiques of Buibas, Patent Owner does not appear to identify any specific factor or criterion that would need to be met to demonstrate disclosure of separate learning processes. See POR, 27-32. Based on my review, the term “learning process” is never defined in the ‘474 Patent and is instead used in just one paragraph of the ‘474 Patent’s lengthy disclosure. TESLA-1001, 103:10-19. In a related IPR involving similar subject matter and claims that likewise referred to multiple learning processes, Patent Owner’s declarant, Dr. Saber, acknowledged that different learning processes may involve different users, different devices, and operating device(s) at different times. TESLA-1032, 38:18-39:2, 49:19-51:12, 52:11-55:21. Here, all of those factors are present in Buibas, which allows users to independently carry out different training routines on different devices at different times. A POSITA would have understood from these considerations that the instruction sets learned on a parent third device in Buibas are learned in another learning process separate from the learning process on the second device by different users.

TESLA-1022, ¶¶32-33 (cited in Reply, pp. 21-22)

Part II: Unpatentability of Substitute Claims Presented with PO's Motion to Amend

Overview of the Additional Prior Art

Opposition to Motion to Amend

Overview of Zhu




FIG. 9

Zhu's bounding boxes identify objects detected in a vicinity of an autonomous vehicle

Referring back to FIG. 1, the autonomous vehicle 102 may also include an electronic representation of a driving environment for maneuvering in the driving environment, and for determining whether there are one or more objects proximate to the autonomous vehicle 102 in the driving environment. For example, autonomous driving computer system 144 may include detailed map information 114 that defines one or more driving environments. The detailed map information 114 may include various maps that identify the shape and elevation of roadways, lane lines, intersections, crosswalks, speed limits, traffic signals, buildings, signs, real time traffic information, or other such objects and information. The

TESLA-1023, 10:29-40 (cited in TESLA-1022, ¶46)

TESLA-1023, FIG. 9 (cited in TESLA-1022, ¶41)

Combinations of Grotmol-Zhu and Grotmol-Zhu-Buibas

Zhu's Additional Features Would Have Been Combined into Grotmol

Dr. Janet:

44. A POSITA reviewing the teachings of Grotmol and Zhu would have found it obvious to apply Zhu's suggestion for extracting additional features of detected objects—such as a classification/type of the object, and the object size and location (e.g., using a bounding box). In the combination, Grotmol's autonomous vehicle would then predictably use the additional features of the extracted objects suggested by Zhu to learn/characterize images and circumstances encountered by the vehicle during training, and to compare current circumstances obtained during autonomous operation to previously learned circumstances to determine instruction sets for operating the device consistent with Grotmol's existing methodologies for training and autonomously operating a device based on a wide range of features of detected objects and sensory inputs. TESLA-1008, 21:64-22:18, 25:30-42, 26:5-20, TESLA-1022, ¶44 (cited in Opp, pp. 2-3)

A POSITA Would Have Been Motivated to Combine Grotmol and Zhu for Multiple Reasons

Dr. Janet:

45. A POSITA would have been motivated to combine Grotmol and Zhu for multiple reasons. To start, a POSITA would have sought to classify and label the type of objects detected in the vicinity of Grotmol's autonomous vehicle based on Zhu's teachings to provide additional semantic information to Grotmol's controller that would allow it to more accurately and efficiently characterize circumstances encountered by the vehicle and to compare current circumstances to prior circumstances learned in the past (e.g., as represented in a trained neural network). TESLA-1023, 11:16-55. For example, semantic labels indicating the type of detected objects would allow the robot to readily characterize a current scene at a higher level of abstraction than the raw pixels of an image, and the semantic labels could be processed more efficiently by a neural network than the raw pixels of an image to determine the similarity of the current scene to prior learned circumstances (e.g., determine whether a tree, a ball, and an object are present in both the current scene and a prior circumstance). A POSITA likewise would have sought to apply TESLA-1022, ¶¶45 (cited in Opp, pp. 3-4)

Dr. Janet:

46. A POSITA also would have been motivated to combine Grotmol and Zhu to gain the benefit of Zhu's cross-validation techniques that allow the accuracy of different sensors on the vehicle to be evaluated. TESLA-1008, 38:43-39:12, FIGS. 20A-B; TESLA-1023, 10:29-57, 11:56-62, 14:43-15:65, 22:64-23:10.

47. Zhu also provides implementation details for sensors present on an autonomous vehicle and situations where this sensory input can be used to control an autonomous vehicle. TESLA-1024, 8:25-46, 10:20-28. Specifically, Zhu illustrates that a vehicle's sensors can be used to continuously capture data, such that evolving changes in the vehicle's circumstance can be detected expediently. *Id.* TESLA-1022, ¶¶46-47 (cited in Opp, pp. 4-5)

A POSITA Would Have Expected Success in Combining Grotmol and Zhu

Dr. Janet:

49. A POSITA would have reasonably expected success implementing the combination considering that Grotmol and Zhu both relate to processing sensory information on autonomous vehicles, and Grotmol likewise describes techniques for extracting features of circumstances and objects detected in the vehicle's environment. TESLA-1008, 18:29-37, 33:49-55; TESLA-1023, 14:43-17:28, FIGS. 9-10. Further, Zhu's autonomous vehicles include sensors explicitly disclosed by Grotmol for providing input to a neural network (e.g., cameras, SONAR, LiDAR). TESLA-1008, 8:57-9:40, 10:16-22, 38:33-51; TESLA-1023, 1:18-27, 8:25-46. Additionally, Grotmol also provides a suitable mechanism (e.g., an adaptive controller) for implementing Zhu's robotic behaviors (e.g., labeling objects). TESLA-1008, 21:64-22:18, 25:30-42, 26:5-20, 27:24-40, FIGS. 10A-11B, 15. Another prior art patent (Gabardos, TESLA-1030) to the same assignee (Brain Corp.) as Grotmol likewise confirms that the use of a bounding box to indicate size and location of objects in a system substantially similar to that described in Grotmol was known before earliest effective filing date of the '474 Patent. TESLA-1030, 30:18-33, 38:56-65, FIGS. 10A-11B. TESLA-1022, ¶49 (cited in Opp, pp. 5-6)

Other Brain Corp. prior art references confirm that bounding boxes were known before the Critical Date. Brain Corp. is the assignee of Grotmol and Buibas.

Gabardos:

The apparatus 1000 may comprise an adaptive controller 1012 configured to determine and provide control output to channels 1020, 1022, 1024. In one or more implementations, individual channels 1020, 1022, 1024 may comprise one or more of motor actuators, electrical components (e.g., LED), transducers (e.g., sound transducer), and/or other electrical, mechanical, and/or electromechanical components. Output 1014 of the feature extractor 1010 may be provided to adaptive controller 1012. In some implementations of robotic navigation applications, the output 1014 may comprise, e.g., position of an object of interest represented by, e.g., coordinates (x, y, width, height) of a bounding box around the object, Extracted contours of obstacles from the image, information related to motion of the robotic platform (e.g., speed, direction of motion of the rover 160 in FIG. 1) and/or other information related to environment of the robot and/or state of the robotic platform.

TESLA-1030, 30:18-33 (cited in TESLA-1022, ¶49)

A POSITA Would Have Also Combined Grotmol, Zhu, and Buibas

Dr. Janet (First Declaration):

151. For example, in a predictable combination based on Grotmol in view of Buibas, Buibas' brain images would be implemented as part of the artificial intelligence systems operating on Grotmol's mobile robotic apparatus (e.g., a neural network). TESLA-1008, 19:10-18, 26:5-20, 34:3-12. Grotmol's adaptive controller for instance, could include a spiking neural network configured according to a brain image to operate in lieu of or in conjunction with the RKNN classifiers. Adaptive controller configurations including brain images and/or RKNN classifiers trained on different robots and by different users could then be shared and even combined on different robots, as suggested by Buibas. TESLA-1005, ¶¶[0070]-[0071], [0184], [0188], [0196], [0201]. Multiple reasons would have motivated a POSITA to pursue this combination.

TESLA-1003, ¶151 (cited in Pet., p. 78)

Dr. Janet (Second Declaration):

2. Ground MTA-1B: Grotmol-Zhu-Buibas

50. It would have been obvious to further combine Grotmol-Zhu with Buibas for the reasons that I explained that it would be obvious to combine Grotmol with Buibas in my Prior Declaration, e.g. such that the adaptive controllers on Grotmol's device would be implemented using brain images like those described in Buibas, where the brain images could be independently trained, shared, and merged by different users on different devices to flexibly provide different sets of capabilities and learned behaviors to Grotmol's device. TESLA-1003, ¶¶150-157. A POSITA would have been motivated to combine Grotmol-Zhu with Buibas, and would have reasonably expected success, for the reasons discussed in my Prior Declaration.

TESLA-1022, ¶50 (cited in Opp, p. 6)

A POSITA Would Have Also Combined Grotmol, Zhu, and Buibas (continued)

Dr. Janet (First Declaration):

152. First, a POSITA would have been motivated to implement spiking neural networks and brain images in the combination to realize the known benefits of such techniques like those expressly recognized in Buibas, including their ability to be readily trained by users, to evolve, and to be merged and partitioned to implement multiple tasks/behaviors on a robot. See, e.g., TESLA-1005, [0073]-[0099], FIGS. 3-4.

153. Second, a POSITA would have been motivated to implement Grotmol such that its adaptive controllers (e.g., containing spiking neural networks and/or RKNN classifiers) would include inputs and instructions sets learned by training multiple devices to expand the capabilities of the device and to more efficiently obtain training data from a fleet of robots without requiring all training to be performed locally on a single device. Indeed, this arrangement would provide increased scalability for operations at a larger scale (e.g., fleet operations).

TESLA-1003, ¶¶152-153 (cited in Pet., p. 78-79)

154. Third, a POSITA would have been motivated to pursue the combination to generate a revenue stream through the sharing of learned knowledge. For example, Buibas discloses that “subscription[s]” can be used to access trained brain images created by other users, thus providing a source of recurring revenue and increased profitability from these subscription models. TESLA-10059, ¶[0071]. For example, Grotmol’s techniques include a “network interface” and “remote server” that would have been capable of sharing brain images. TESLA-1005, 8:16-26, 46:27-35, FIG. 7.

155. Fourth, a POSITA would have been motivated to pursue this combination because doing so is merely the application of known techniques as suggested by Buibas to a known system (e.g., Grotmol’s robots) to yield predictable results.

156. Fifth, a POSITA would have been motivated to pursue this combination because it merely instantiates use of known techniques as taught in Buibas to improve a similar system as taught in Grotmol.

TESLA-1003, ¶¶154-156 (cited in Pet., p. 79-80)

Obviousness based on Grotmol-Zhu and Grotmol-Zhu-Buibas

Substitute Claim 85

Original Limitations from Claim 45:

85. A first device comprising: a knowledgebase that includes one or more inputs for inputting at least a portion of a circumstance representation, wherein the one or more inputs are correlated with one or more instruction sets for operating a second device; one or more sensors; one or more processors; and one or more non-transitory machine readable media storing machine readable code that, when executed by the one or more processors, causes the one or more processors to perform at least: generating a circumstance representation, wherein the generated circumstance representation represents a circumstance detected at least in part by the one or more sensors; determining the one or more instruction sets for operating the second device at least by: inputting at least a portion of the generated circumstance representation into the one or more inputs, and using a correlation between the one or more inputs and the one or more instruction sets for operating the second device; and at least in response to the determining, executing the one or more instruction sets for operating the second device, wherein the first device autonomously performs one or more operations defined by the one or more instruction sets for operating the second device; ...

Substitute Claim 85 (cont'd)

New Limitations Presented with PO's Motion to Amend:

85. ... wherein the knowledgebase includes a neural network, wherein the one or more inputs are one or more input neurons of the neural network, wherein the machine readable code, when executed by the one or more processors, further causes the one or more processors to perform at least: generating a stream of circumstance representations comprising: the generated circumstance representation that (i) includes multiple object representation data structures that each comprises: a label for object type, coordinates for object location, and a computer model for object size, and (ii) represents multiple objects detected at a first time, and another circumstance representation that (i) includes another multiple object representation data structures that each includes: a label for object type, coordinates for object location, and a computer model for object size, and (ii) represents the multiple objects detected at a second time, wherein the determining the one or more instruction sets for operating the second device includes: inputting the multiple object representation data structures of the stream into the one or more input neurons, and using a correlation between the one or more input neurons and the one or more instruction sets for operating the second device; and determining another one or more instruction sets for operating a third device at least by: inputting the another multiple object representation data structures of the stream into the one or more input neurons, and using a correlation between the one or more input neurons and the another one or more instruction sets for operating the third device.

The Board Preliminarily Determined that the Substitute Claims do not Distinguish the Opposition's Prior Art

Preliminary Guidance:

For the reasons explained in the Opposition, at this stage, and based on the current record, we determine that Petitioner has shown a reasonable likelihood that proposed substitute independent claim 85, as well as proposed substitute dependent claims 86–91, would have been obvious based on Grotmol and Zhu.

For the reasons explained in the Opposition, at this stage, and based on the current record, we determine that Petitioner has shown a reasonable likelihood that proposed substitute independent claim 85, as well as proposed substitute dependent claims 86–91, would have been obvious based on Grotmol, Zhu, and Buibas.

Preliminary Guidance, pp. 13-14

Substitute Claim Limitation 85.8

85.8: wherein the knowledgebase includes a neural network, wherein the one or more inputs are **one or more input neurons** of the neural network

A POSITA would have Understood that Neural Networks Include Input Neurons

Fisher:




FIG. 10A

TESLA-1029, FIG. 10A (cited in TESLA-1022, ¶52)

Dr. Janet:

classifier)”), 34:3-17, 36:19-31. A POSITA would have known that neural networks conventionally included input neurons (as evidenced by Buibas), and it would have been obvious for a neural network in Grotmol’s adaptive controller to include input neurons that would allow it to receive information characterizing a sensory input from which a control output could be generated. TESLA-1005, ¶¶[0068], [0073]-[0076], [0136]-[0138], [0144]-[0154], FIGS. 1, 10, 11. For example, Fisher, incorporated by reference in Grotmol (TESLA-1008, 9:53-60) discloses a spiking neural network 1040 having input neurons 1042. TESLA-1029, [0227]-[0228], FIG. 10A. Any of a variety of neural network architectures would be obvious and suitable to apply in the combination (e.g., as discussed in Bichler and Gocsek below). TESLA-1035, ¶¶[0025]-[0040], [0119]-[0122], FIG. 7; TESLA-1036, ¶¶[0044]-[0047], FIG. 5.

TESLA-1022, ¶52 (cited in Opp, pp. 7-8)

Buibas' Neural Network Includes Input Neurons

Buibas:




FIG. 10

TESLA-1005, FIG. 10 (cited Opp, p. 8)


Dr. Janet:

55. In addition to Grotmol's disclosures of neural networks, Buibas further describes a spiking neural network that would have been obvious to provide in Grotmol's adaptive controller in some implementations of the combination, which network includes input neurons for receiving information characterizing sensory input(s). TESLA-1005, ¶¶[0068], [0073]-[0076], [0136]-[0138], [0144]-[0154], FIGS. 1, 10, 11.

TESLA-1022, ¶55 (cited in Reply, p. 8)

Patent Owner's Support for "Input Neurons" is Indistinguishable from the Prior Art

The '474 Patent:



TESLA-1001, FIG. 17 (cropped to show "Model A")

Buibas:




FIG. 10

TESLA-1001, FIG. 10 (cited Opp, p. 8)

Motion to Amend:

The following limitations are additions to Substitute Claim 85, and Dr. Saber's testimony confirms that the inventor was in possession of these limitations (and the claim as a whole) at the time of filing the '991 Application. *See* Ex. 2001, ¶135. The '976 Application discloses element [85.8]. *See* [92:8-9], [Fig. 17 – Model A showing input nodes/neurons], [94:1-2 – “For example, a neural network may

Motion to Amend, p. 15

Dr. Janet:

69. The Patent Owner and its declarant, Dr. Saber, have explained that even a basic neural network like the feedforward network Model A depicted in FIG. 17 of the '474 Patent has input neurons correlated with instruction sets for operating a device merely because the input neurons have “downstream connections with output nodes/neurons that include/produce instruction sets for operating a device.” MTA, 15-16 (identifying written description support in Model A); EX2001, pp. 79-80, 113, TESLA-1022, ¶69 (cited in Reply, p. 19)

Substitute Claim Limitations 85.9-85.10

85.9-85.10: wherein the machine readable code, when executed by the one or more processors, further causes the one or more processors to perform at least: **generating a stream of circumstance representations** comprising:


Grotmol Discloses Machine Readable Code Executed by a Processor

Grotmol:

Dr. Janet:

56. Grotmol discloses that its mobile robot includes machine-readable code that, when executed by one or more processors, causes the one or more processors to perform operations. TESLA-1008, 2:15-18, 3:27-36, 7:66-8:4, 38:33-42, FIG. 8; Pet., 62-63 (Element [45.4]).

TESLA-1022, ¶56 (cited in Reply, p. 8)



*one or more non-transitory
machine readable media
storing machine readable code*

FIG. 8

TESLA-1008, FIG. 8 (annotated).

TESLA-1008, FIG. 8 (cited and annotated in Pet, p. 63)

Grotmol Discloses a Stream of Circumstance Representations

Grotmol:




FIG. 20A

TESLA-1008, FIG. 20A (cited and annotated in Opp, p. 46)

objects within the field-of-view **168**. In one or more implementations, such as object recognition, and/or obstacle avoidance, the output provided by the sensor component **166** may comprise a stream of pixel values associated with one or more digital images. In one or more implementations of


TESLA-1008, 8:57-9:3 (cited in TESLA-1022, ¶57)

Dr. Janet:

57. Grotmol's mobile robot includes a sensor component that "provide[s] information associated with objects within the field-of-view" and that includes a "a stream of pixel values associated with one or more digital images." TESLA-1008, 8:57-9:3, 44:46-55 (describing a "camera stream"). Grotmol further discloses that the robot, during autonomous operation, generates a stream of observations of the environment (e.g., images 2000 and 2010) that represent circumstances encountered by the robot over time. TESLA-1008, 38:52-39:28, 38:38-41 ("sequence of images"), 42:46-43:11 (describing autonomous operation including capturing sensory inputs to be correlated with motor commands at each "given time step"), 44:56 ("At individual time steps[] ..."), 44:34-46:10, FIGS. 8, 20A-B; Pet., 63-64 at [45.5] ("generating a circumstance representation"). In the combination, Zhu TESLA-1022, ¶57 (cited in Reply, p. 9)

Zhu Also Discloses a Stream of Circumstance Representations

In the combination, Zhu likewise discloses generating a stream of circumstance representations representative of the circumstances of an autonomous vehicle



TESLA-1023, FIG. 9 (cited and annotated in TESLA-1022, ¶41)

Zhu:

The data provided by the sensors 146 may be processed by the autonomous driving computer system 144 in real-time. In this context, the sensors 146 may continuously update their output to reflect the environment being sensed at or over a range of time, and continuously or as demanded. The sensors 146 may provide the updated output to the autonomous driving computer system 144 so that it can determine whether the autonomous vehicle's 102 then-current direction or speed

TESLA-1023, 10:20-28 (cited in TESLA-1022, ¶57)

Dr. Janet:

[45.5] ("generating a circumstance representation"). In the combination, Zhu likewise discloses generating a stream of circumstance representations representative of the circumstances of an autonomous vehicle| TESLA-1023, 10:20-28 ("[T]he sensors 146 may continuously update their output to reflect the environment being sensed at or over a range of time, and continuously or as demanded."), 8:25-46, 13:46-58, FIGS. 1, 7-10.

TESLA-1022, ¶57 (cited in Reply, p. 9)

Substitute Claim Limitation 85.11

85.11: the generated circumstance representation that (i) includes **multiple object representation data structures** that each comprises: **a label for object type, coordinates for object location, and a computer model for object size**, and (ii) represents **multiple objects detected at a first time**, and

Grotmol and Zhu Both Disclose Multiple Object Representation Data Structures

Grotmol:




FIG. 20A

TESLA-1008, FIG. 20A (cited in TESLA-1022, ¶59)

Zhu:




FIG. 9

TESLA-1023, FIG. 9 (cited in TESLA-1022, ¶59)

multiple object representation data structures that each comprises: a label for object type, coordinates for object location, and a computer model for object size and (ii) represents multiple objects detected at a first time

Dr. Janet:

[45.5]; see also TESLA-1023, 10:20-28, 8:25-46, 13:46-58, FIGS. 1, 7-10. For example, the robot may capture an image 2000 of its environment that includes representations of multiple objects (e.g., a tree, a box, and a ball) at a first time. TESLA-1008, 38:43-39:12 ("object representations"). According to Grotmol, the

60. Zhu further describes techniques for classifying objects detected in a vicinity of an autonomous vehicle at a given time and generating, for each object, a label for object type, coordinates for object location, and a computer model for object size. For example, Zhu discloses labeling the object type for each object by TESLA-1022, ¶¶59-60 (cited in Opp, p. 10-12)

Zhu Classifies Objects With a Type

Zhu:

In detecting vehicles in various driving environments, the data 112 may include vehicle data 116 that defines one or more parameters for classifying a vehicle. Classifications of vehicle may include such classifications as “passenger car,” “bicycle,” “motorcycle,” and other such classifications. The parameters defined by the vehicle data 116 may inform the autonomous driving computer system 144 as to the type of vehicle detected by a given sensor. For example, the vehicle data 116 may include parameters that define the type of vehicle when the vehicle is detected by one or more of the camera sensors, one or more of the laser sensors, and so forth.

TESLA-1023, 11:11-36 (cited in TESLA-1022, ¶60)

When the autonomous driving computer system 144 detects an object, the autonomous driving computer system 144 may apply an object label to the detected object based on one or more of the sensor parameters 148 previously discussed. The object label may be defined by a bounding box encompassing the object. In alternative embodiments, the object label may be defined by a bounding oval or other bounding shape. FIGS. 9-10 illustrate examples of object labels being applied to various objects (i.e., vehicles) detected in the camera image 702 and the laser point cloud image 802.

TESLA-1023, 14:43-52 (cited in TESLA-1022, ¶60)

multiple object representation data structures that each comprises: a label for object type, coordinates for object location, and a computer model for object size and (ii) represents multiple objects detected at a first time

Dr. Janet:

60. Zhu further describes techniques for classifying objects detected in a vicinity of an autonomous vehicle at a given time and generating, for each object, a label for object type, coordinates for object location, and a computer model for object size. For example, Zhu discloses labeling the object type for each object by generating “such classifications as ‘passenger car,’ ‘bicycle,’ ‘motorcycle,’ and other such classifications” that indicate “the *type* of vehicle detected by a given sensor.” TESLA-1024, 11:16-36, 14:43-52 (“may apply an object label to the detected object”), 10:58-64 (data 112 defining parameters for “the *type* of object detected by the sensor”). Zhu’s system is also configured to “define[] ... a bounding TESLA-1022, ¶60 (cited in Opp, p. 11-12)

Grotmol Explicitly Discloses that Objects have a Size and Position

Grotmol:

The RKNN process may partition available data into subsets comprising a given number of features from the lower-dimension/lower redundancy data. The given number of features associated with lower-dimension/lower redundancy data may be referred to as the mandatory feature(s). As used herein the term feature may be used to describe one or more integer or floating point values characterizing the input, e.g., the presence or absence of an edge, corner, shape, texture, color, object, at particular locations in the image, values of pixels in an image, patches of color texture, brightness in the image, and/or in the image as a whole; properties of the mentioned features, such as size, orientation, intensity, predominance with respect to the surround, of an edge, corner, shape, texture, color, object; the position of one of the features in the image or the relative position of two or more of the above mentioned features; changes in features across consecutive frames—changes in position (optic flow), intensity, size, orientation; the pitch, intensity, spectral energy in specific bands, formants of sounds, the temporal changes thereof, disparity measure between two or more images, input from proximity sensors (e.g., distance, proximity alarm, and/or other), motor feedback (e.g., encoders position), motion sensor input (e.g., gyroscope, compass, accelerometer), previous motor commands or switching commands, a binary/Boolean categorical variable, an enumerated type, a character/string, and/or practically any characteristic of the sensory input.

TESLA-1008, 28:19-40 (cited in TESLA-1022, ¶59)

multiple object representation data structures that each comprises: a label for object type, coordinates for object location, and a computer model for object size and (ii) represents multiple objects detected at a first time

Dr. Janet:

TESLA-1008, 38:43-39:12 (“object representations”). According to Grotmol, the robot can detect objects from sensory signals such as image 2000 and can further extract features associated with the detected objects that enable comparison to stored features. TESLA-1008, 39:13-28 (“images (e.g., raw and/or pre-processed) ... may be stored in a memory buffer” and the “pre-processing operations may comprise ... feature extraction”), 28:19-40 (“the term feature may be used to describe ... values characterizing the input, e.g., the presence or absence of ... [an] object, at particular locations in the image” including “properties of the mentioned features such as size, ... position”), 21:41-43, 9:41-2 (“object recognition and/or classification”), 10:23-34 (“The tasks of the robot may be configured based on a context. ... [T]he context

TESLA-1022, ¶59 (cited in Opp, p. 10-11)

Patent Owner Fails to Contest Grotmol's Disclosure

Dr. Saber:

26. In my opinion, Grotmol in view of Zhu does not disclose or suggest "multiple object representation data structures that each comprises: ... coordinates for object location, and a computer model for object size" as recited in limitations [85.11]-[85.12]. Although the Opposition mentions certain disclosures from Grotmol (Opp., 10-11), it relies on obviousness in view of Zhu for [85.11]-[85.12]. Opp., 13 ("It would have been obvious to a POSITA to apply Zhu's suggestion for ... defining bounding boxes for detected objects that provides a computer model for object size coordinates for object location [sic] to Grotmol..."); *see also id.*, 11-12. Limitations [85.11] and [85.12] recite "object representation data structures" that comprise the "coordinates for object location" and "model for "object size." The object representation data structures are part of a circumstance representation, which represents a circumstance detected by one or more actual, real-world sensors (*see* Ex. 1001, Figs. 3A-3E, showing Camera 92a, Microphone 92b, Lidar 92c, Radar 92d, and Sonar 92e) for detecting objects and their properties, as recited in [85.2] and [85.5].

EX-2006, ¶26

Second Deposition of Dr. Saber:

Q. Okay. I'm about ready to move on.
It's just a simple question. I just wanted to confirm. Is there any portion of your declaration that rebuts the portions of Grotmol that were cited in the opposition with respect to elements 85.11 to 85.12? And if there's not, you can just let me know. I just want to make sure I'm not missing something.

(Witness reviews document.)

A. I mean, I guess I don't understand exactly the question. But I feel I have responded to -- you know, the opposition to motion to amend raised an obvious argument with regards to Grotmol and a POSITA's obvious argument with regard to Zhu. So I'm addressing directly that.

I didn't say something specifically about the Grotmol, Column 10 section, if that's what you're asking, specifically about those words. But I think I have responded to the argument in general.

TESLA-1039, 16:19-17:16 (cited in Sur-Reply p. 3)

Patent Owner Fails to Contest Grotmol's Disclosure (cont'd)

Second Deposition of Dr. Saber:

Q. Okay. Well, I'm giving you another opportunity here. You didn't address it in your declaration, but you have an opportunity to now.

So is there anything in Grotmol that precludes object location as disclosed in Grotmol from being coordinates for object location as recited in the claims?

MR. JAIRAM: Objection, form; objection, scope.

THE WITNESS: And again, as I mentioned, I did not provide direct opinion on that. I'm not going to provide an opinion today here in the declaration -- in the deposition without, you know, taking the proper time to think about it. But the opinions that I provided were directly in response to what -- the arguments that were raised in the opposition to motion to amend and those contained in the paragraphs that we've been talking about. I think Paragraph 26 to 31.

TESLA-1039, 35:19-36:17 (cited in Sur-Reply p. 3)

The '474 Patent Does Not Require “Actual” Size or Location

The '474 Patent:

elements or information. Object Representation 625 may include an electronic representation of an object (i.e. Object 615 [later described], etc.) detected in Device's 98 surrounding. In some aspects, Collection of Object Representations

TESLA-1001, 72:4-7 (cited in Sur-Reply p. 7)

thereof that can be detected. Examples of object properties include existence of an object, type of an object (i.e. person, cat, vehicle, building, street, tree, rock, etc.), identity of an object (i.e. name, identifier, etc.), distance of an object, bearing/angle of an object, location of an object (i.e. distance and bearing/angle from a known point, coordinates, etc.), shape/size of an object (i.e. height, width, depth, computer model, point cloud, etc.), activity of an object (i.e. motion, gestures, etc.), and/or other properties of an object. Type of

TESLA-1001, 72:59-67 (cited in Sur-Reply p. 7)

Zhu Additionally Discloses Object Size and Location

Zhu:

Where the sensor that captured the image is one or more of the lasers, the object label parameters may also include a third pixel coordinate that defines the physical height of the object or a particular laser point depicted in the captured image (e.g., a Z-coordinate). This third pixel coordinate should not be confused with the height parameter of the object label, because this third pixel coordinate may indicate the elevation of the detected object or of a given laser point (e.g., 3 meters above sea level, 2 meters above sea level, etc.) This third pixel coordinate may further indicate the height of the detected object or laser point relative to the autonomous vehicle **102**.

Where the sensor that provides data about the object is one or more of the radar detection units, the object label parameters may include a height parameter that defines the height of the height of the bounding box (e.g., in pixels), a width parameter that defines the width of the bounding box (e.g., in pixels), a first pixel coordinate that defines the latitudinal placement of the bounding box (e.g., an X-coordinate), and a second pixel coordinate that defines the longitudinal placement of the bounding box (e.g., a Y-coordinate).

TESLA-1023, 15:1-20 (cited in TESLA-1022, ¶60)

*multiple object representation data structures that each comprises: a label for object type, **coordinates for object location, and a computer model for object size** and (ii) represents multiple objects detected at a first time*

Dr. Janet:

60. Zhu further describes techniques for classifying objects detected in a vicinity of an autonomous vehicle at a given time and generating, for each object, a label for object type, coordinates for object location, and a computer model for object size. For example, Zhu discloses labeling the object type for each object by

TESLA-1022, ¶60 (cited in Opp, p. 11-12)

Zhu Renders Coordinate Locations Obvious

Dr. Janet cited multiple prior art references disclosing coordinate systems, for example, McClure

McClure

The vision calibration mode is illustrated in FIG. 4. The purpose of this mode is to find the relationship between the CCD camera x-y coordinate system and the gripper x-y coordinate system. The two x-y planes

TESLA-1009, 5:32-36 (cited in TESLA-1003, ¶145)

Dr. Janet (First Declaration):

145. Third, a POSITA would have been motivated to pursue this combination because doing so is merely the application of known techniques (operating a robotic apparatus in a coordinate system) to a known system (Grotmol's mobile robotic apparatus) to yield predictable results (a mobile robotic apparatus using a coordinate system for navigation). It was well known before the critical date to use coordinate systems to define the location of an object(s) to a robot. TESLA-1009, 1:54-57, 2:12-16, 5:27-6:19; TESLA-1018, ¶¶[0073]-[0077], [0080]; TESLA-1019, ¶¶[0005]-[0006], [0043], [0046], [0049], [0071], [0076]-[0078], [0091]-

TESLA-1003, ¶145 (cited in Pet. p. 75)

multiple object representation data structures that each comprises: a label for object type, coordinates for object location, and a computer model for object size and (ii) represents multiple objects detected at a first time

Dr. Janet (Second Declaration):

62. Zhu also discloses that object bounding boxes are defined by "coordinates" that define the horizontal and vertical placement of the bounding box ("coordinates for object location" – shown above in FIG. 9). TESLA-1024, 14:43-15:20, FIGS. 9-10. A POSITA would have recognized and found obvious that these bounding boxes would have additionally been assigned a coordinate location in space, because as Zhu discloses, the location of the object is relative to the location of the autonomous vehicle, and as discussed in my Prior Declaration, coordinate systems were commonly used before the Critical Date to define such relationships. TESLA-1024, 7:14-19, 8:56-63.

TESLA-1022, ¶62 (cited in Opp, p. 12-13)

Grotmol and Zhu Represents Objects with Respect to Time

Grotmol:




FIG. 20A

TESLA-1008, FIG. 20A (cited in TESLA-1022, ¶59)




FIG. 20A

TESLA-1008, FIG. 20A (cited in TESLA-1022, ¶64)

multiple object representation data structures that each comprises: a label for object type, coordinates for object location, and a computer model for object size and (ii) represents multiple objects detected at a first time

Dr. Janet:

[45.5]; see also TESLA-1023, 10:20-28, 8:25-46, 13:46-58, FIGS. 1, 7-10. For example, the robot may capture an image 2000 of its environment that includes representations of multiple objects (e.g., a tree, a box, and a ball) at a first time. TESLA-1008, 38:43-39:12 (“object representations”). According to Grotmol, the

60. Zhu further describes techniques for classifying objects detected in a vicinity of an autonomous vehicle at a given time and generating, for each object, a label for object type, coordinates for object location, and a computer model for object size. For example, Zhu discloses labeling the object type for each object by

TESLA-1022, ¶¶59-60 (cited in Opp, p. 10-12)

TESLA-1022, ¶59 (cited in Opp, p. 10-11)

Substitute Claim Limitation 85.12

85.12: another circumstance representation that (i) includes **another multiple object representation data structures** that each includes: **a label for object type, coordinates for object location, and a computer model for object size**, and (ii) represents the **multiple objects detected at a second time**

Grotmol and Zhu's Stream Includes Multiple Object Representation Data Structures at a First and Second Time

Grotmol:

Image 2000 including multiple objects detected at a first time




FIG. 20A

TESLA-1008, FIG. 20A (cited in TESLA-1022, ¶59)




FIG. 20A

Image 2010 including multiple objects detected at a second time

TESLA-1008, FIG. 20A (cited in TESLA-1022, ¶64)

another multiple object representation data structures that each includes: a label for object type, coordinates for object location, and a computer model for object size, and (ii) represents the multiple objects detected at a second time

Dr. Janet:

FIGS. 1, 7-10. In the combination, and for the reasons addressed in [85.11], the circumstance representation generated at the later (second) time would provide another circumstance representation that (i) includes another multiple object representation data structures that each includes: a label for object type, coordinates for object location, and a computer model for object size, and (ii) represents the multiple objects detected at a second time.

TESLA-1022, ¶64 (cited in Opp, p. 13-14)

Grotmol and Zhu Represent Objects with Respect to Time

Grotmol:

Image 2000 including multiple objects detected at a first time




FIG. 20A

TESLA-1008, FIG. 20A (cited in TESLA-1022, ¶59)




FIG. 20A

Image 2010 including multiple objects detected at a second time

TESLA-1008, FIG. 20A (cited in TESLA-1022, ¶64)

another multiple object representation data structures that each includes: a label for object type, coordinates for object location, and a computer model for object size, and (ii) represents the multiple objects detected at a second time

Dr. Janet:

[45.5]; see also TESLA-1023, 10:20-28, 8:25-46, 13:46-58, FIGS. 1, 7-10. For example, the robot may capture an image 2000 of its environment that includes representations of multiple objects (e.g., a tree, a box, and a ball) at a first time. TESLA-1008, 38:43-39:12 (“object representations”). According to Grotmol, the

60. Zhu further describes techniques for classifying objects detected in a vicinity of an autonomous vehicle at a given time and generating, for each object, a label for object type, coordinates for object location, and a computer model for object size. For example, Zhu discloses labeling the object type for each object by TESLA-1022, ¶¶59-60 (cited in Opp, p. 10-12)

Substitute Claim Limitations 85.13-85.14

85.13: wherein the determining the one or more instruction sets for operating the second device includes: **inputting the multiple object representation data structures of the stream into the one or more input neurons**, and **using a correlation between the one or more input neurons and the one or more instruction sets** for operating the second device; and

85.14: determining another one or more instruction sets for operating a third device at least by: **inputting the another multiple object representation data structures of the stream into the one or more input neurons**, and **using a correlation between the one or more input neurons and the another one or more instruction sets** for operating the third device.


Inputting Grotmol-Zhu's Object Representation Data Structures into Input Neurons Would Have Been Obvious to a POSITA

Dr. Janet:

25:30-47-42, 17:28-44, 28:41-44, FIGS. 10-11. Inputting such object representation data structures into the input neurons would have been obvious especially because it involves no more than using the controllers/neural networks in Grotmol according to their ordinary and intended purpose—namely, to process information characterizing sensory inputs including features of detected objects to produce controls for autonomously operating a robot. TESLA-1008, FIGS. 10A-11B. It is TESLA-1022, ¶¶67 (cited in Opp, p. 15-16)

circumstances. TESLA-1037, ¶¶[0004]-[0006], [0043]-[0048]. The '474 Patent likewise acknowledges that a POSITA would have possessed the skill necessary to implement such features on a neural network. TESLA-1001, 103:40-60 ("The disclosed systems, devices, and methods are independent of the artificial intelligence model and/or technique used and any model and/or technique can be utilized to facilitate the functionalities described herein."), 104:2-27, 169:30-54, FIG. 17.

TESLA-1022, ¶¶67 (cited in Opp, p. 15-16)



TESLA-1001, FIG. 17

'474 Patent:

Referring to FIG. 17, the disclosed artificially intelligent devices, systems, and methods for learning and/or using a device's circumstances for autonomous device operation may include various artificial intelligence models and/or techniques. The disclosed systems, devices, and methods are independent of the artificial intelligence model and/or technique used and any model and/or technique can be utilized to facilitate the functionalities described herein. Examples of these models and/or techniques include deep learning, supervised learning, unsupervised learning, neural networks (i.e. convolutional neural network, recurrent neural network, deep neural network, etc.), search-based, logic and/or fuzzy logic-based, optimization-based, tree/graph/other data structure-based, hierarchical, symbolic and/or sub-symbolic, evolutionary, genetic, multi-agent, deterministic, probabilistic, statistical, and/or other models and/or techniques.

TESLA-1001, 103:40-60 (cited in TESLA-1022, ¶¶67)

Multiple Prior Art References Corroborate that it was Known to Input Data Representations into Input Neurons

Dr. Janet:

receive inputs. *See, e.g.*, TESLA-1029, [0227]-[0234], FIG. 10A. A POSITA would have reasonably expected success inputting object representation data structures into the input neurons since neural networks were known to be capable of processing a wide range of data including features of detected objects as of the Critical Date. For example, Buibas corroborates that neurons can be used to detect “complex features and objects.” TESLA-1005, ¶¶[0140]-[0142], [0228]. Likewise, Yufik also discloses that neurons can receive and process semantic understanding of objects and circumstances. TESLA-1037, ¶¶[0004]-[0006], [0043]-[0048]. The '474 Patent likewise acknowledges that a POSITA would have possessed the skill necessary to implement such features on a neural network. TESLA-1001, 103:40-60 (“The TESLA-1022, ¶67 (cited in Opp, p. 15-16)

Yufik:

[0006] A computer based system capable of constructing situational models receives input from different sensors and utilizes several types of artificial neurons for input processing, as follows: sensory neurons (pixels) respond to different elements (features) in the incoming streams, temporal neurons respond to various temporal relations in the activation of sensory neurons (e.g., activation of neuron A preceded acti-


TESLA-1037, [0004] (cited in TESLA-1022, ¶67)

Buibas:

[0140] In one implementation, the detection signals are delivered to a next layer of the detectors 1012 (comprising detectors 1012_1, 1012_m, 1012_k) for recognition of complex object features and objects, similar to the exemplary implementation described in commonly owned and co-pending U.S. patent application Ser. No. 13/152,084, filed Jun. 2, 2011, entitled “APPARATUS AND METHODS FOR PULSE-CODE INVARIANT OBJECT RECOGNITION”, incorporated herein by reference in its entirety. In this implementation, each subsequent layer of detectors is configured to receive signals from the previous detector layer, and to detect more complex features and objects (as compared to the features detected by the preceding detector layer). For example, a bank of edge detectors is followed by a bank of bar detectors, followed by a bank of corner detectors and so on, thereby enabling, inter alia, alphabet recognition by the apparatus.

TESLA-1005, [0140] (cited in TESLA-1022, ¶67)

Buibas Discloses that Multiple Brain Images Can be Merged



TESLA-1005, FIG. 4 (cited in Pet., p. 43)

[0188] To this end, in the exemplary implementations, brain images may be uploaded/stored as full or partial images. Full images may be loaded on to a neuromorphic apparatus 1150 and run as a standalone control. Partial images may lack the full functions necessary to run certain features of the robotic device 610. Thus, partial images may be used to augment or upgrade (downgrade) a pre-loaded brain image or a stored brain image. It will be appreciated that a full brain image for a first device may serve as a partial brain image for second device with all of the functionality of the first plus additional features. In some implementations, two or more partial brain images may be combined to form full brain images.

TESLA-1005, [0188] (cited in TESLA-1022, ¶72)

Dr. Janet:

72. For example, Buibas discloses techniques by which many different users can independent train their respective robotic devices to autonomously perform different tasks/behaviors. Users can then make the trained configurations of their robots available to others (e.g., by uploading trained brain images to the cloud). TESLA-1005, [0171]-[0180] (describing details of the training application provided to each user), [0181]-[0193] (describing details of sharing trained brain images through the cloud), [0194]-[0207], [0211]-[0213], [0222]-[0225], [0070], FIGS. 16-19. Users can then select one or more brain images to download and load on their own robots to take advantage of the training performed by others, and can even merge multiple brain images learned on other devices to obtain a combined set of capabilities. TESLA-1005, [0090]-[0127], [0188]-[0189], [0225], FIGS. 3-9. The

TESLA-1022, ¶72 (cited in Reply., p. 20-21)

The Substitute Claims Lack Written Description Support

The Substitute Claims Lack Written Description Support

Dr. Janet:

79. Substitute claim 85 requires inputting a very specifically-recited set of multiple object representation data structures into input neuron(s) of a neural network and using correlations between the input neurons and instruction sets for operating a second device and a third device to determine instruction sets from the inputted data structures. MTA, 1-3. Based on my review, Patent Owner has not shown, however, that the inventor was in possession of these concepts, and hence has not shown that the substitute claims are supported by the '474 Patent or the '991 Priority Application. The MTA includes lengthy string citations to various portions of the lengthy specification of the '991 Priority Application, but the citations to alleged support for processing the recited data structures with a neural network are quite limited.

TESLA-1022, ¶79 (cited in Opp., p. 23)

80. For example, the MTA cites to FIG. 17 and the feedforward neural network shown in Model A as alleged support for these features. MTA, 15-16. But none of the disclosures cited in the MTA describe inputting the multiple object representation data structures into the input neurons of a neural network for processing as claimed—and particularly not multiple object representation data structures that include the recited label for object type, coordinates for object location, and computer model for object size. There is simply no evidence that the inventor possessed the use of a neural network for this specific scenario. Patent Owner instead improperly treats the '991 Application's boilerplate that "[t]he disclosed systems, devices, and methods are independent of the artificial intelligence model and/or technique used and any model and/or technique can be utilized to facilitate the functionalities described herein," as license for asserting long after the filing date that the inventor possessed the specifically claimed combination of features recited in claim 85, in the context of a neural network. EX2004, 93:24-32.

TESLA-1022, ¶79 (cited in Opp., pp. 23-24)

The Substitute Claims Lack Written Description Support (cont'd)

Dr. Janet:

81. The MTA likewise fails to demonstrate support for a number of specific nuances recited in Substitute Claim 85, including that the instruction sets are determined using correlations between input neurons and two *different* devices (i.e., the second and third devices) for object representation data structures pertaining to the *same* “multiple objects” detected at different times. MTA, [85.11]-[85.14]. For example, Patent Owner points to FIG. 38 as depicting a “second device” (“Loader 98a”), even though this disclosure does not address the claimed “third device,” let alone demonstrate that instruction sets can be determined for operating a third device based on the representations of the same objects that were used to determine instructions for operating the second device (“Loader 98a”). MTA, 16-17; TESLA-1001, FIG. 38. Citing to an entirely different figure, Patent Owner then alleges that “Boat 98b” is the claimed third device. *Id.*; TESLA-1001, FIG. 39. The Patent Owner’s theory would appear to require that (1) the Loader 98a and Boat 98b learned instructions for operating their respective devices based on representations of the same objects, and (2) the first device would need to be capable of executing instructions from both Loader 98a (second device) and Boat 98b (third device). *Id.* TESLA-1022, ¶81 (cited in Opp., pp. 24-25)

Dr. Saber:

Q. How is it that a neural network can have correlations to instruction sets for operating a second device and a third device?

A. I don't understand what you're asking.

Q. Well, that's my fear, sir.

A. Well, I don't think you're -- I don't think you're asking a question that makes any sense. I mean, neural networks have multiple outputs. So output 1 could be move the chair. Output 2 could be have a drink of coffee. Those are two different instruction sets.

TESLA-1028, 124:7-17 (cited in Opp., p. 24)

Thank You