

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

Petitioner's Demonstratives

Case No. IPR2023-01055
U.S. Patent No. 11,238,344
Before Barbara Parvis, Robert Weinschenk, and Russell Cass
Administrative Patent Judges

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Table of Contents


• The '344 Patent and Petition Overview	3
• IPR (Petition) Challenge	6
• Overview of Grotmol, Hickman, and the Combination	6
• Issue 1: Grotmol Discloses Modifying the Instruction Sets (Claim 7)	16
• Issue 2: Grotmol Discloses Modifying First and Second Circumstances (Claim 8)	21
• Issue 3**: Grotmol Discloses “connection[s]” (Claim 16)	25
• Issue 4*: Grotmol Renders Obvious Claim 10	30
• Issue 5**: Grotmol Renders Obvious Claim 11	35
• Issue 6**: GHC Renders Obvious Claims 6, 11, and 12	39
• Motion to Amend	43

'344 Patent and Petition Overview




DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

The 344 Patent



TESLA1001, cover.



TESLA1001 ('344 pat.), FIGS. 4A, 4B.

Asserted Grounds

Claims Challenged	35 U.S.C. §	Reference(s)/Basis
3–5, 7–9, 13–18	102(a)(2)	Grotmol
10, 11	103	Grotmol
6, 11, 12	103	Grotmol, Hickman

Claims 1-2, 19, and 20 were disclaimed. POPR, 1.

ID., 8; Pet., 3-4.

Invalidity Challenge


Overview of Grotmol, Hickman, and Combination

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Grotmol

TESLA1004



TESLA1004, FIG. 1.

(57)

ABSTRACT

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

TESLA1004, Abstract.

FISH.

TESLA1004, Abstract, FIG. 1; TESLA1002, ¶¶[47]-[51]; Pet., 7-9; I.D., 23.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Grotmol

TESLA1004




FIG. 20A

TESLA1004, FIG. 20A.

In one or more implementations, images 2000, 2010 may be obtained with a camera 166 mounted on a robotic vehicle 160 of FIG. 1. Images 2000, 2010 may comprise representations of visual field that may be obtained at times $t_1 < t_2$, respectively. The image 2000 may comprise representations one or more objects, e.g., a tree 2002, a rectangle, a ball 2008, and/or other objects. During motion of the vehicle 160 relative position and/or orientation of the camera and the visual scene may change. By way of an illustration, the vehicle 160 may execute a right turn. The image 2010 may be obtained subsequent to the turn execution. Representation of objects in the image 2010 may be shifted relative representations of the same objects in the frame obtained prior to the turn execution. Additional object representations (e.g., 2016 in FIG. 20A) may appear, while representations of some object that may be present in one image (e.g., 2000) may be absent from the subsequent image (e.g., 2010). Comparing image 2010 to image 2000, representations of the tree 2012 the rectangle, 2014, and the ball 2018 are shifted (horizontally) within the frame 2010 relative representations, 2002, 2004, 2008 in the frame 2010. The amount of the shift may be configured based on relative motion between the images (e.g., motion speed and time interval $t_2 - t_1$ between the images). It may be further noticed from

TESLA1004, 38:52-39:8.

Grotmol

TESLA1004




FIG. 20B

FIG. 20B illustrates output of the shift and trim operation. Frame 2020 depicts the trimmed version of the image 2000; frame 2040 depicts the shifted and trimmed version of the image 2010 in FIG. 20A. It is noteworthy that matching images may not necessarily be identical to one another. Variability between matching images may exist. By way of an illustration, the ball 2018 in the image 2010 may be disposed closer to the rectangle 2014 relative the ball 2008 position in the image 2000. Accordingly, the shifted/trimmed frames 2020, 2040 in FIG. 20B may comprise differences.

A similarity metric may be determined between the shifted/trimmed frames (e.g., 2020, 2040 in FIG. 20B). In some implementations, the similarity may be determined by performing pixel-wise difference between the overlapping regions (e.g., the frames 2020, 2040 in FIG. 20B) of the given image and the matching image. A norm of the image difference in the overlapping region may be determined for

TESLA1004, 40:64-41:15.

Grotmol

TESLA1004

Grotmol

In listen mode, the robot may be “listening” to or otherwise observing what the teacher teaches, and may not do anything on its own. However, the robot may learn an association. But if the teacher stops sending a command, the

TESLA1004, 23:3-12 Pet., 8.

Trained behaviors of the robotic device may be configured based on learning of associations between sensory context (e.g., presence of an obstacle in front of the robotic vehicle) and a respective action (e.g., right turn) during training

TESLA1004, 33:59-64; Pet., 16.

Dr. Janet’s Declaration

48. Grotmol’s training process first involves a “listen mode,” in which a robot observes a user’s commands in order to “learn an association.” Ex. 1004, 23:3-11. Grotmol discloses that the robot may then switch—either manually or automatically—into an “autonomous mode,” in which the robot performs tasks autonomously based on learned associations. Ex. 1004, 23:46-48, 24:1-16. Additionally, Grotmol describes “override” and “control” modes that use varying amounts of user operation to facilitate the learning process. In “override mode,” a robot will operate autonomously until it receives a user’s command, then apply the command, and “try to remember it for the next time the same situation occurs.” *Id.*,

TESLA1002, ¶¶[48]; see also ¶¶[77], [97], [130].

Overview of Hickman

Hickman discloses a cloud-based shared knowledge base.

FIG. 3 shows an example of multiple robots 301, 302, 303 interacting with a cloud computing system 304 having a shared robot knowledge base 306 according to some embodiments of the disclosed systems and methods.

their queries. The cloud computing system 304 may also receive data from the robots, and update information in the shared robot knowledge base 306 based on the data received from the robots.

The shared robot knowledge base 306 may include multiple component and/or sub-component knowledge bases. The shared robot knowledge base 306 may contain additional or fewer knowledge bases than the ones shown in FIG. 3. Similarly, the distribution of the information across the various component knowledge bases shown FIG. 3 is merely one example of one embodiment. The contents of the knowledge bases may be combined into a single common knowledge base or distributed across multiple different knowledge bases in a different fashion than described herein. Similarly, in




FIGURE 3

TESLA1005, FIG. 3 (annotated).

TESLA1005, 10:35-66; Pet. Rep., 21-22; Pet. Opp. To MTA, 13-14.

Pet. Rep., 22.

Overview of Hickman

Hickman's knowledge base can be shared by multiple robots.

Also, as described above, some applications in the application knowledge base 311 may build upon data from other knowledge bases. For example, the “Fetch” application may rely on a “grasp” task from the task knowledge base 309 and a “ball” object from the general object knowledge base 308. As individual robots become better at applying the “grasp” task to the “ball” object through either trial and error and/or human “coaching,” the individual robots can update data in the shared robot knowledge base 306 based on their individual experiences. As a result, the cloud computing system 304 can optimize the “Fetch” application based on the collective experiences of many different robots. Thus, each time a robot connects to the cloud computing system 304 to execute the “Fetch” application from the application knowledge base 311, that robot gains the benefit of the collective experiences of many other robots that have executed the “Fetch” application in the past.

TESLA1005, 24:8-25; TESLA1002, ¶¶[146]-[147].

The advantages of a shared robot knowledge base 306 multiply as robots are deployed around the world. If information about a particular object is “public” and therefore accessible by all robots, then information that a single robot learns about that object and uploads to the shared robot knowledge base 306 may be accessible by other robots all around the world. In operation, many millions or even many billions of robots may ultimately access the shared robot knowledge base 306. As a result, the shared robot knowledge base 306 enables robots to benefit from the collective information learned by many other robots.

TESLA1005, 11:31-41; TESLA1002, ¶[147].

Hickman's Structure is Similar to the '344 Patent's Structure

'344 Patent




FIG. 4B

TESLA1001, FIG. 4B.

TESLA1010, pg. 28; Pet. Rep, 25-26.

Hickman




FIGURE 3

TESLA1005, FIG. 3.

TESLA1010, pg. 29 ; Pet. Rep, 25-26.

A POSITA Would Have Combined Grotmol and Hickman

Grotmol

one of two states. The memory component may be configured to store training sets. A given training set may comprise an instance of the video frame, an instance of the binary signal, and an instance of a motor control indication configured to cause the robot to execute an action. The one or

TESLA1004, 3:35-39.; Pet., 12-13.

Petition

In the combination of Grotmol and Hickman, the training configuration learned by operation of a robot described in Grotmol would be uploaded to the server containing the “shared robot knowledge base” described in Hickman, and a second robot would receive the circumstance representations and instructions stored in the knowledge base in order to perform autonomous operation. Decl. ¶ 242-244. Thus,

Pet., 60.

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

14

A POSITA Would Have Combined Grotmol and Hickman

Dr. Janet's Declaration

142. A POSITA would have been motivated to combine the autonomous robotic system of Grotmol with the shared robot knowledge base of Hickman, which would allow for simultaneous learning across multiple devices with multiple users, increasing efficiency and conserving local system resources. As detailed in Section

TESLA1002, ¶142.

the training and autonomous operation of robots. Hickman is also cited on the face of Grotmol as a cited U.S. patent document. *Id.*, p. 2. A POSITA would certainly consider the references listed on the face of Grotmol when looking to expand upon its teachings.

TESLA1002, ¶144.

ISSUE 1

Grotmol discloses Modifying the Instruction Sets (Claim 7)

[7] ... *modifying: the first one or more instruction sets* for operating the first device, or a copy of the first one or more instruction sets for operating the first device, ... *anticipating the modified the first one or more instruction sets* for operating the first device ... *executing the modified the first one or more instruction sets*”

PO's Arguments Contradict the '344 Patent's Description of Modification

'344 Patent

example, Modification Interface 130 comprises the functionality to access, create, delete, modify, and/or perform other manipulations on functions, methods, procedures, routines, subroutines, and/or other elements of Application Program 18. In a further example, Modification Interface 130 comprises the functionality to access, create, delete, modify, and/or perform other manipulations on source code, bytecode, compiled, interpreted, or otherwise translated code, machine code, and/or other code. In a further example, Modification Interface 130 comprises the functionality to access, create, delete, modify, and/or perform other manipulations on values, variables, parameters, and/or other data or information. Modification Interface 130 comprises any features, functionalities, and embodiments of Acquisition Interface 120, and vice versa. Modification Interface 130 also comprises other disclosed functionalities.

TESLA1001, 140:37-52; Pet. Rep., 2-3; TESLA1010, ¶[15].

PO Response

Overriding an instruction set is not modifying the instruction set. Ex. 2007, ¶¶77-79.

POR, 13.

Ex. 1004, 34:49-55). But Grotmol only discloses that a trainer may correct or change the operation of a robot in a “combiner” situation *by providing additional or new instructions*, Ex. 1004, 34:49-35:34, and not that *the instruction sets that are included with the training sets are actually modified*. Ex. 2007, ¶¶77-87. The

POR, 12.

Grotmol only discloses *correcting device behavior* in certain instances and does not disclose *modifying an instruction* set and then anticipating a modified instruction set, executing the modified instruction set, and performing an operation defined by that modified instruction set. Ex. 2007, ¶¶77-87. Grotmol discloses that

POR, 12.

PO's Narrow Interpretation of "Modifying" is Not Supported

Pet. Rep.

Additionally, although Grotmol's modifications are not limited to inserting or adding instruction sets, PO's assertions that such inserting/adding is not "modifying" contradicts the '344 patent's description of "modifying" and the plain meaning of the term. POR, 12-17; TESLA1014, 703. In the '344 patent, "modifying" encompasses operations including additions, deletions, or other alterations. TESLA1001, 140:8-52 ("Modification Interface 130 comprises the functionality to access, create, delete, modify[] ..."), 140:53-141:40 ("inserting additional code, overwriting or rewriting existing code, and/or branching to a separate segment of code"), FIGS. 2, 12, 30, 31A, 31B; TESLA1010, ¶¶10-13.

Pet. Rep., 3; TESLA1010, ¶¶15-16.

18

FISH

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Grotmol Explicitly Teaches Modifying Instruction Sets

GROTMOL

accordance with one or more implementations. In some implementations of supervised training of robots, control instructions (also referred to as corrections) produced by the trainer (e.g., human) may be combined with control instructions produced by the robot controller instructions (predictions).

TESLA1004, 34:50-55; Pet. 27; TESLA1002, ¶¶[191]-[193].

the spot which the robot is headed towards. Thus, if this spot is to the left of the spot where the robot is supposed to be headed towards (as defined by the camera image seen during training), then the robot may need to adjust its heading rightwards.

TESLA1004, 42:54-58; Pet. 27; TESLA1002, ¶¶[191]-[193].

The user UNDO indication may be configured based on one or more of the user activating a user interface element (e.g., a physical and/or virtual touch-screen button), a voice command, a gesture, and/or other actions, in one or more implementations. One or more undo indications may be utilized in order to remove multiple behaviors (and/or multiple versions of a given behavior). By a way of an illustra-

TESLA1004, 34:28-38; Pet. Rep., 3-4; TESLA1010, ¶[15].

In the override mode, the robot may execute what it has learned, unless a command is sent by the teacher. As soon as the teacher starts sending commands, the robot may stop taking the initiative and may let the teacher control it. For example, if the robot is turning left but the teacher wants the robot to turn right and provides a right turn command, then the robot may head the teacher's command, perform the action, and try to remember it for the next time the same situation occurs. Once a behavior only needs fine tuning, the correct mode may be used.

TESLA1004, 23:13-23; Pet. Rep., 2-3; TESLA1010, ¶[15].

19

Grotmol Explicitly Teaches Modifying Instruction Sets

GROTMOL

TABLE 3

Mode	Action of the robot in presence of training input	Action of the robot in absence of training input
Control	Executes the action in accordance with the training input; Does not learn the association	Idle
Listen	Executes the action in accordance with the training input; Learns the association	Idle
Override	Executes the action in accordance with the training input; Learns the association	Executes the action in accordance with the prediction
Correct	Combine the teaching input and the prediction and execute the action in accordance with the combined signal. Learn the resulting association	Executes the action in accordance with the prediction
Autonomous	Ignore the teaching input and execute the action in accordance with the prediction	Executes the action in accordance with the prediction

TESLA1004, 22:38-40, Table 3; Pet. Rep., 2.

Issue 2

Grotmol Discloses Modifying First and Second Circumstances (Claim 8)

[8] ... *modifying at least one of: the first circumstance representation*, a copy of the first circumstance representation, the second circumstance representation, or a copy of the second circumstance representation, and wherein the at least partial match between the second circumstance representation and the first circumstance representation includes: (i) *at least partial match between the modified the second circumstance representation and the first circumstance representation*, ...

Grotmol Describes Modifying Circumstances

Petition

Grotmol discloses that pixels from individual image instances are selected, scaled, normalized, and/or equalized. *See* Ex. 1004, 4:18-22 (“The method may comprise applying a transformation to an instance of the sensory input. The transformation may be configured to produce scaled input based on analysis of the history.”); 30:47-31:2 (discussing “normalizing” sensory input by “shifting and/or scaling input features”); 31:61-32:5 (discussing equalizing sensory data); 3:42-46; 4:58-63; 31:3-8. Thus, Grotmol discloses a system with processors configured to modifying the first and second circumstance representations. Decl. ¶ 196.

Pet., 30-31.

Grotmol

The method may comprise applying a transformation to an instance of the sensory input. The transformation may be configured to produce scaled input based on analysis of the history

In order to facilitate contributions from different types of signals for determining a distance measure between features in a metric space (e.g., Euclidian distance), data from highly redundant input (e.g., the RGB pixel values) may be normalized. Various other distance measures (metrics) may be utilized, e.g., Mahalanobis, Manhattan, Hamming, Chebyshev, Minkowski, and/or other metrics.

In some implementations, the normalization may comprise shifting and/or scaling input features to a given value range (e.g., A1=64 to A2=196 for an 8-bit pixel value, 0 to 1 range, and/or other range). In one or more implementations, the normalization may be configured

TESLA1002, ¶[196]; TESLA1004, 4:18-22, 30:47-31:2, 31:61-32:5, 3:39-48; 31:3-8.

Grotmol Describes Modifying Circumstances

Dr. Janet Declaration

in the training buffer and *move individual particles* so that stored images match closely with newly arrived images. TESLA1004, 39:42-54. In the image match process, a windowing function can also be applied to individual images “to *produce [a] windowed image and reduce edge effects.*” TESLA1004, 40:20-51. Grotmol also teaches that an image may be *shifted and trimmed*. TESLA1004, 40:52-41:18; FIGS. 20A, 20B (reproduced below).

TESLA1010, ¶[22].

Grotmol




FIG. 20A

Images are shifted
and/or trimmed.




FIG. 20B

Grotmol (TESLA1004), FIGS. 20A, 20B.

TESLA1010, ¶[2].2

Modified Circumstances are used in Comparisons

Dr. Janet Declaration

21. As noted in my Prior Declaration (and not disputed by PO), “Grotmol discloses that pixels from individual image instances are selected, scaled, normalized, and/or equalized.” TESLA1002, ¶¶196-197; *see also* TESLA1004, 30:47-31:2, 4:18-22, 31:61-32:5, 3:42-46, 4:58-63; 31:3-8. As part of the ...
TESLA1004, 30:54-31:2. Since the adjustment in parameters to perform normalization is determined, at least in part, on the total duration of the training, a POSITA would have understood that the normalization of the images occurs after the training is complete. That is, images (circumstance representations) are correlated with instructions, and then normalized. These images are used for matching purposes. *See e.g.*, TESLA1002, ¶¶196-197, 177-178. For at least this reason, PO’s assertion that the images are not used in matching, is incorrect.

TESLA1010, ¶[21].

Grotmol

motor control indication by at least: determining a subset of features comprising the binary signal and a set of pixels randomly selected from the plurality of pixels; scaling individual pixels of the set of pixels by a scaling factor; scaling features of the subset by a scaling factor; comparing individual scaled features of the subset to corresponding features of individual ones of the training sets; based on the comparison, determining a similarity measure for a given training set, the similarity measure characterizing similarity

TESLA1004, 3:42-52; TESLA1010, ¶[21].

Issue 3

Grotmol Discloses “connections” (Claim 16)

[16] ... wherein the first circumstance representation correlated with the first one or more instruction sets for operating the first device include at least a portion of the first one or more object representations *connected, using at least one or more connections*, with at least a portion of the first one or more instruction sets for operating the first device.

PO's Construction is Unclear and Unsupported

PO Response

a “connection” should be construed as a “relationship that includes a link.” Ex. 2007, ¶¶56-63.

POR, 6.

Pet. Rep.

Federal Circuit has emphasized that “[t]here are only two exceptions” to the “general rule” that “[t]he words of a claim are generally given their ordinary and customary meaning”: namely, “(1) when a patentee sets out a definition and acts as his own lexicographer,” and “(2) when the patentee disavows the full scope of a claim term either in the specification or during prosecution.” *Thorner v. Sony Comput. Entm't Am. LLC*, 669 F.3d 1362, 1365 (Fed. Cir. 2012). Neither condition is satisfied here. Tellingly, PO never argues that the specification or file history contains any relevant lexicography or disavowals. They do not.

Pet. Rep., 8.

“connection.” POR, 10. But PO offers no explanation as to why it cherry-picked one definition involving the word link from the various other available definitions, particularly when the description in the '344 patent does not even use the word “link” to describe a “connection.” TESLA1010, ¶28.

Pet. Rep., 10.

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

The POM of “connection” is not limited to a “link”

EX. 2013

con·nec·tion \kə'nekshən\ *n* -s [L. *connexio*-, *connexio*, fr. *connexus* (past part. of *connectere* to connect) + *-ion*-, *-io* -ion — more at **CONNECT**] 1 **a** : the act of connecting : a coming into or being put in contact (~ with the island was made by a causeway) **b** : sexual relation or intercourse (had had no ~ with any other woman —John Abernethy) 2 : the state of being connected or linked : ALLIANCE, UNION (Canada's political ~ with England) (~ between church and state) 3 **a** : relationship or association in thought (as of cause and effect, **logical sequence**, mutual dependence or involvement) (the ~ of intelligence and success) **b** : CONTEXT, REFERENCE, OCCASION (in this ~ the word has a different meaning) **c** : COHERENCE, CONTINUITY (a confused multitude without order or ~ —John Locke) 4 : CONNECTIVE **b, d** 5 **a** : something that connects : COUPLING, **LINK** (plumbing ~) **b** : a means of communication (telephone ~) or transport (the train makes ~ with the steamer) (to miss a ~) 6 : a person connected with others by marriage, remote blood relationship, or such a tie as a common interest (he has powerful ~s in high places) 7 : a social, professional, or commercial relationship in a practical or active way : as **a** : POSITION, JOB **b** : a permanent or continuing arrangement to execute orders or advance interests esp. at a distance (a firm's foreign ~s) **c** *slang* : a source of contraband (as a narcotic drug) 8 : a set or group of persons connected or associated together in a common interest : **a** : DENOMINATION, SECT **b** : a political faction **c** *Brit* : the owner of a racehorse and his associates **d** : a large family : CLAN **e** : a clientele esp. of a doctor or lawyer 9 : a religious association practicing connexionalism — **con·nec·tion·al** \-kshən'l, -kshnəl\ *adj*

EX2013, 481; POR, 10; Pet. Rep., 10.

EX. 2012

con·nec·tion (kə-nēk'shən) *n*. Abbr. **con.** 1. **a**. The act of connecting. **b**. The state of being connected. 2. One that connects; **a link**: a faulty connection in the circuit. 3. An association or a relationship: There appeared to be no connection between the two crimes. 4. The logical or intelligible ordering of words or ideas; coherence. 5. Reference or relation to something else; context: In this connection, the agreement can be seen as a step toward peace. 6. A person, especially one of influence or position, with whom one is associated, as by kinship, interests in common, or marriage: used her connections to land a job. 7. A conveyance or scheduled run providing continuing service between means of transportation: missed my connection in Atlanta. 8. A line of communication between two points in a telephone or similar wired system. 9. *Slang*. **a**. A drug dealer. **b**. A purchase of illegal drugs. — **con·nec'tion·al** *adj*.

EX2012, 399-400; POR, 10; Pet. Rep., 10.

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Examples and Embodiments of “Connection” in the ’344 patent are non-limiting

Pet. Rep.

PO attempts to import a requirement for a “connection” to include a “link” by mere reference to *examples* in the specification where a “link” is disclosed. POR, 6, 21. For example, PO turns to FIGS. 17 and 18A, which are described as “embodiments.” TESLA1001, 56:29-34; POR, 6-9. But these “embodiments” do not impose a limiting effect. TESLA1010, ¶25. The ’344 patent explains that Model A in FIG. 17 “*may* include a network of Nodes 852 ... and Connections 853 similar to that of a brain.” TESLA1001, 103:49-67. In FIG. 18A, “[e]ach of Connection 853z1 and Connection 853z2 *may* include or be associated with occurrence count, weight, and/or other parameter or data.” TESLA1001, 106:19-22. The ’344 patent also states that a “connection *may include or be associated* with at least one of: an occurrence count, a weight, a parameter, or a data.” TESLA1001, 16:5-6. These descriptions include a permissive “may,” and are precisely the type of non-limiting language that does not limit the scope of a claim term. *Aventis Pharma S.A. v Hospira, Inc.*, 675 F.3d 1324, 1330 (Fed Cir. 2012). Notably, the textual description of FIGS. 17 and 18A does not once use the word “link,” which PO seeks to import

Pet. Rep., 9.

’344 Patent

FIG. 17 illustrates various artificial intelligence methods, systems, and/or models that can be utilized in DCADO Unit 100 embodiments.

TESLA1001, 56:29-31; Pet. Rep., 9.

The connection may include or be associated with at least one of: an occurrence count, a weight, a parameter, or a data.

TESLA1001, 16:5-6; Pet. Rep., 10.

various properties. Knowledge cells may be connected, interrelated, or interlinked into knowledge structures using statistical, artificial intelligence, machine learning, and/or other models or techniques. Such interconnected or interrelated knowledge cells can be used for enabling autonomous device operation. The interconnected or interrelated knowledge cells may be stored or organized into a knowledgebase (i.e. Knowledgebase 530, etc.). In some embodiments,

TESLA1001, 154:10-17; POR, 9.

FISH

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

A “connection” is an “association or a relationship”

Pet. Reply

Based on the broad description of the term “connection” in the ’344 patent (see e.g., TESLA1001, 16:5-6, 66:13-33), Petitioner submits that no construction is necessary and that the plain and ordinary meaning should apply. The meaning that most closely aligns with the breadth of this term in the specification is “an association or a relationship.” EX2012, 399-400, EX2013, 481; TESLA1001, 16:5-6, 66:13-33; Pet. 39 (citing TESLA1002, ¶208); TESLA1010, ¶29.

Pet. Rep., 10.

Dr. Janet Declaration and Grotmol

Trained behaviors of the robotic device may be configured based on learning of associations between sensory context (e.g., presence of an obstacle in front of the robotic vehicle) and a respective action (e.g., right turn) during training.

Ex. 1004, 33:24-26, 33:59-62. In my opinion, a POSITA would have understood the association between an object representation and an action to be a “connection” between the two. At least the data structures described in Grotmol linking object representations and instruction sets together in memory, for example, would comprise “one or more connections” as recited in the claim. *Id.*, 3:35-39 (“The TESLA1002, ¶[208].

TESLA1002, ¶[208].

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Issue 4

Grotmol Renders Claim 10 Obvious

[10] ... wherein at least a portion of the third circumstance representation or at least a portion of the second one or more instruction sets for operating the first device *is learned in another learning process that includes operating the first device at least partially by another user.*

PO has no support for “*one* user operating *one* device in *one* learning process” theory

Patent Owner Response

process.” Grotmol only discloses *one* user operating *one* device in *one* learning process. Ex. 2007, ¶¶103-112. Grotmol fails to disclose that at least a portion of the

POR, 24.

No citations to Grotmol. Only citations to PO's expert's repetitive statements

Even PO's Expert Contradicts PO's "one user, *one* device, *one* learning process" Theory

DR. SABER DEPO. TRANSCRIPT

3 Q. So Figure 10B, the adaptive controller,
4 there's a line that goes diagonally across the
5 adaptive controller, it says "user commands."
6 Do you see that?
7 A. Yes.
8 Q. And there's another one for adaptive
9 controller 2, also have user commands; is that
10 right?
11 A. Yes.
12 Q. And so these adaptive controllers can
13 receive commands from different users and then
14 have the outputs; is that right?
15 A. That's correct. I mean, you receive
16 commands from different users to change a certain
17 behavior or control a certain behavior or adjust
18 it.

TESLA1011, 166:3-18; Pet. Rep., 13.

GROTMOL




FIG. 10B

TESLA1004, FIG. 10B; Pet. Rep., 13.

PO's Characterization of Grotmol is Incorrect On its Face

Grotmol explicitly teaches multiple users.

PO Response

users Grotmol actually discloses. Since Grotmol discloses *only a single user*, *this single user performs the “training of both simple and complex tasks.” Nothing in Grotmol suggests that one user trains the robot for simple tasks and another user trains the robot for complex tasks.* Petitioner's contention that Grotmol in view of

POR, 25.

GROTMOL

The BrainOS software platform may be configured to enable users to selectively remove a learned behavior (and/or a portion thereof) via an undo and/or time machine operation.

TESLA1004, 34:18-21; Pet., Rep., 8.

DR. JANET DECLARATION

39. Because learning is user dependent and Grotmol explains that each robot can have multiple users, it would have been obvious to a POSITA that multiple users could operate a single robot and train the robot differently according to the users' instructions. Therefore, Grotmol renders obvious “a portion of the second one or more instruction sets for operating the first device is learned in *another learning process* that includes *operating the first device* at least partially *by another user.*”

TESLA1010, ¶[39].

33

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Multiple Users can Provide Training to Grotmol's Robot

The user may have an ability to perform managing operations on "Behavior 1". For example, the user may perform operations including one or more of change a name of a module, reset a module to naive state, revert a module to a particular point in history, delete a module, save/load locally, upload, download, and/or other operations. The user may change outputs and FE1. Such changes may depend on whether it is compatible with a trained behavior (e.g., dimensionality of features and outputs is preserved), in some implementations. **The user may train a target approach.** The user may change input from a target tracker to a base tracker so that robot will approach to base. A robot with multiple grippers may be trained to grasp with gripper 1 and user wired output from the adaptive module form gripper 1 to gripper 2.

TESLA1004, 25:53-67; Pet. Rep., 13.

The system may be configured to learn the second behavior from the user. Figure above shows that the user may select a different set of input features and a different set of outputs for "Behavior 2". The user may select a different type of the adaptive controller ("2"). For example, the user

TESLA1004, 26:5-10; Pet. Rep., 13.

GROTMOL




FIG. 10B

TESLA1004, FIG. 10B; Pet. Rep., 13.

Issue 5

Grotmol Renders Claim 11 Obvious

[11] ... wherein at least a portion of the *third circumstance* representation or at least a portion of the second 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: the user, or another user.*

Grotmol Discloses “a third circumstance”

GROTMOL

training handset. The trained controller of the robot may comprise a trained configuration configured to enable autonomous operation (e.g., without user input) by the robotic device. The trained configuration may be stored. A saved configuration may be loaded into the robot being trained thereby providing one or more trained behaviors to the robot. In some implementations, the trained configuration may be loaded to one or more other robots in order to provide learned behaviors. Subsequent to loading of the saved configuration, the controller learning process may match process configuration being present during saving of the configuration.

TESLA1004, 33:24-35; Pet., 54; TESLA1002, ¶[233]..

ured to store training sets. A given training set may comprise an instance of the video frame, an instance of the binary signal, and an instance of a motor control indication configured to cause the robot to execute an action. The one or

TESLA1004, 3:36-39; Pet., 12; TESLA1002, ¶¶[165]-[172]..

36

Grotmol Renders Obvious “*the third device,*” “*another user,*” and “*another learning process*”

DR. JANET’S DECLARATION.

47. In addition to teaching that a third device can provide training information and learned behaviors for loading onto the first device, Grotmol discloses that different types of robots (e.g., rovers, bi-pedaling robots, wheeled vehicles, aerial drones, robotic limbs, and/or other robotic devices) can be used, and that each robot has a trainer/user. TESLA1004, title, 33:21-35, 34:49-39:22, 16:4-8, 18:33-37, FIG. 1. Grotmol also discloses that a “trainer may be provided with the control of the robot during training.” TESLA1004, 34:49-62, 23:3-11, 39:13-22. It would therefore have been obvious to a POSITA that when there are multiple robots, particularly different types of robots, there would also have been different users.

Pet. Rep., 19; TESLA1010, ¶¶[47]..

45. A POSITA would have understood that the robots do not just share training information (e.g., correlation information) and then do nothing with the received information. Rather, this information is used so that the receiving robot (first device) can benefit from the learning of the transmitting robot (third device) and also perform “learned behaviors.” See e.g., TESLA1005, 11:30-41. Such sharing of learning processes was, in fact, well-known and understood in the art, e.g., as corroborated by Hickman, at least because it allowed “robots to benefit from the collective information learned by many other robots.” TESLA1005, 11:30-41. Consistent with this understanding, Grotmol teaches that matching is performed by the first device after performing the loading of the training information received from the third device. TESLA1004, 33:21-35 (“match process configuration”).

Pet. Rep., 18-19; TESLA1010, ¶¶[45]..

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE


37

Grotmol's Explicit Disclosure Contradicts PO's "one user, one device, one learning process" Theory

GROTMOL

training handset. The trained controller of the robot may comprise a trained configuration configured to enable autonomous operation (e.g., without user input) by the robotic device. The trained configuration may be stored. A saved configuration may be loaded into the robot being trained thereby providing one or more trained behaviors to the robot. In some implementations, the trained configuration may be loaded to one or more other robots in order to provide learned behaviors. Subsequent to loading of the saved configuration, the controller learning process may match process configuration being present during saving of the configuration.

TESLA1004, 33:24-35; Pet., 54; Pet. Rep., 15-18.



TESLA1004, FIG. 10B; Pet. Rep., 17.

Issue 6

The Grotmol-Hickman Combination (GHC) Renders Claims 6, 11, and 12 Obvious

A POSITA would have combined Grotmol and Hickman in the manner detailed in the Petition.

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

PO's Characterization of Hickman is Incorrect On its Face

Hickman is not limited to separate knowledge bases.

PO Response

(as recited in claims 6, 11 and 12 and shown below). Ex. 2007, ¶¶121-123. Hickman discloses storing instruction sets and images *separately in separate knowledgebases* (images in Unique Object Knowledge Base 307 and instructions in Task Knowledge Base 309, *see* Ex. 1005, Fig. 3), *which a POSA would have found to be incompatible* with Grotmol's disclosure of training sets or correlations with images and instruction sets correlated and stored *together*. Ex. 2007, ¶¶121-123. Further, Hickman teaches

POR, 35.

However, Hickman teaches away from the correlations in Grotmol by disclosing images and instructions *in separate knowledgebases that are not correlated*. Ex.

POR, 36.

Hickman

example of one embodiment. The contents of the knowledge bases may be combined into a single common knowledge base or distributed across multiple different knowledge bases in a different fashion than described herein. Similarly, in some embodiments, the functionality of the cloud processing engine 305 may be combined with the functionality of the shared robot knowledge base 306 and/or other components of the cloud computing system 304. Additional systems and

TESLA1005, 10:63-11:3; TESLA1010: ¶¶67].

herein. In some embodiments where the task data (TASKS) includes robot instruction code for performing tasks, the shared robot knowledge base 304 may not require a separate task knowledge base 309.

TESLA1005, 19:30-33; Pet. Rep.: 23.

Dr. Janet's Declaration

3. Hickman explains that the knowledge base has flexibility in design and the various knowledge bases can be combined into one knowledge base or distributed over the computer system. TESLA1005, 10:57-11:7.

TESLA1010: ¶¶75].

40

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

PO's Characterization of Hickman is Incorrect On its Face

Hickman explains that additional data fields would have been part of the knowledge base.

Dr. Janet's Declaration

predictable results. For instance, applying Hickman's use of knowledge base(s) to facilitate sharing across a cloud computing system would have involved the modification of Grotmol's storage structure to include one or more knowledge bases with Hickman's data structures. The choice of whether to include one knowledge base, as in Grotmol, or various knowledge bases, as in Hickman, would be a designer's choice that was well-within the capabilities of a POSITA, as described in Hickman. TESLA1005, 10:57-11:7. The use of the data structures like those disclosed by Hickman (e.g., data structures 313-316) in Grotmol's system would have been accomplished by programming and memory storage techniques well within a POSITA's capabilities.

TESLA1010: ¶[76].

Hickman

data (SER). Various embodiments of the unique object knowledge base 307 may contain more or less data corresponding to a particular unique object identifier (UID). For example, in some embodiments, the unique object knowledge base 307 may have many more additional categories of data than the categories shown in conceptual data structure 313. Similarly, the unique object knowledge base 307 may store more or less data for some unique objects as compared to the data stored for other unique objects.

TESLA1005, 12:10-18; Pet. Rep., 23-24.

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

41

PO's Characterization of Hickman is Incorrect On its Face

Hickman is not limited to single object information

PO Response

sets correlated and stored *together*. Ex. 2007, ¶¶121-123. Further, Hickman teaches away from the claimed “circumstance representations” because Hickman focuses on manipulating a single object and single object information and such information is not a circumstance representation. Ex. 2007, ¶¶121-123. In addition, instances of

POR, 35.

Hickman

to determine the identity of the object. The shared robot knowledge base contains data related to an expansive inventory of objects, robot tasks, maps, robot applications, etc. In some embodiments, the shared robot knowledge base includes data about a vast number of various objects that robots may encounter and/or interact with in their everyday environments. The shared robot knowledge base may also include robot instructions for interacting with the objects. For example, the shared robot knowledge base may have data related to various types of furniture, appliances, household objects, utensils, tools, electronic devices, food, supplies, etc. Data about any object that that a robot might encounter or interact with may be included in the shared robot knowledge base.

TESLA1005, 2:51-64; TESLA1010: ¶¶[59], [97].

Motion to Amend

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Table of Contents

• Motion To Amend	43
• Issue 7: The Claims Fail Under 35 USC 101	47
• The Claims Fail Under 35 USC 103	56
• Issue 8: GHC Renders Obvious Data Structures with Detected Location and Size Information	59
• Issue 9: GHC Renders Obvious Computer Models Indicating the Detected Sizes of First and Second Objects	62
• Issue 10: GHC Renders Obvious Data Structures Correlated with Instruction Sets	64
• Issue 11: GTC Renders Obvious Data Structures with Detected Location and Size Information	69
• Issue 12: GTC Renders Obvious Computer Models Indicating the Detected Sizes of First and Second Objects	75
• Issue 13: GTC Renders Obvious Data Structures Correlated with Instruction Sets	77

Summary of Amendments

The amendments rewrite each of claims 4/5/6/7/9/10/11/12/16 in independent form and further limit the method of each of these claims by requiring that:

1. the first one or more object representations of the first circumstance representation comprise first and second object representations that include first and second data structures, respectively, which are each generated at least by processing one or more digital pictures to detect a type, a location, and a size of the first (or second) object,
2. and in response to the processing, generating the first (or second) data structure that includes a label indicating the detected type of the first (or second) object, a coordinate indicating the detected location of the first (or second) object, and a computer model indicating the detected size of the first (or second) object,
3. wherein the first and second data structures are correlated with the first one or more instruction sets for operating the first device.

Substitute Claim 21 (Amended claim 4)

Claim 21 (Proposed substitute for claim 4, if found unpatentable) [21.pre]A system comprising:

[21.1]one or more processors configured to perform at least:

[21.2]accessing a memory that stores at least a knowledgebase that includes a first circumstance representation correlated with a first one or more instruction sets for operating a first device, wherein the first circumstance representation represents a first circumstance detected at least in part by one or more sensors of the first device, and wherein at least a portion of the first circumstance representation or at least a portion of the first one or more instruction sets for operating the first device is learned in a learning process that includes operating the first device at least partially by a user;

[21.3]generating or receiving a second circumstance representation, wherein the second circumstance representation represents a second circumstance detected at least in part by: the one or more sensors of the first device, or one or more sensors of a second device;

[21.4]anticipating the first one or more instruction sets for operating the first device based on at least partial match between the second circumstance representation and the first circumstance representation; and

[21.5]at least in response to the anticipating, executing the first one or more instruction sets for operating the first device, wherein the first device or the second device autonomously performs one or more operations defined by the first one or more instruction sets for operating the first device;

[21.6]wherein the first circumstance representation includes a first one or more object representations, and wherein the second circumstance representation includes a second one or more object representations;

[21.7]wherein the second circumstance representation represents the second circumstance detected at least in part by the one or more sensors of the first device, and wherein the first device autonomously performs the one or more operations defined by the first one or more instruction sets for operating the first device;

[21.8]wherein the first one or more object representations of the first circumstance representation comprise a first object representation and a second object representation;

[21.9]wherein the first object representation includes a first data structure generated at least by:

[21.10]processing one or more digital pictures to detect: a type of a first object, a location of the first object, and a size of the first object; and

[21.11]in response to the processing, generating the first data structure that includes: a label indicating the detected type of the first object, a coordinate

indicating the detected location of the first object, and a computer model indicating the detected size of the first object;

[21.12]wherein the second object representation includes a second data structure generated at least by:

[21.13]processing one or more digital pictures to detect: a type of a second object, a location of the second object, and a size of the second object; and

[21.14]in response to the processing, generating the second data structure that includes: a label indicating the detected type of the second object, a coordinate indicating the detected location of the second object, and a computer model indicating the detected size of the second object;

[21.15]wherein the first data structure and the second data structure are correlated with the first one or more instruction sets for operating the first device.

MTA App. A, 1-3.

46

Issue 7

The Claims Fail Under 35 USC 101

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Alice Step 1

The Claims Are Directed to an Abstract Idea

PO's substitute claims are directed to an abstract idea of operating a device in a current circumstance based on how that that device, or another device, operated in similar circumstance(s) in the past. TESLA1010, ¶69. This familiar concept is a mental process that continues to be performed in the mind by humans in countless contexts each day. *CyberSource Corp. v. Retail Decisions, Inc.*, 654 F.3d 1366, 1373 (Fed. Cir. 2011) (“[C]omputational methods which can be performed entirely in the human mind are the types of methods that embody the ‘basic tools of scientific and technological work’ that are free to all men and reserved exclusively to none.”) (quoting *Gottschalk v. Benson*, 409 U.S. 63, 67 (1972)); *Elec. Power Grp., LLC v.*

Paper 31 (Opp. To MTA), 2.

Alice Step 1

Taking Action When Encountering Similar Images (e.g., traffic signs) is a Mental Process

For well over a century, people have used their mental faculties to make real-time decisions driving automobiles informed by their past experiences. TESLA1010, ¶¶70-71. Each time a driver decides to take a detour based on traffic conditions that resemble conditions of the past (e.g., a circumstance where the detour was previously taken), the driver mentally performs the claimed abstract idea. *Id.*

Paper 31 (Opp. To MTA), 2.

Alice Step 1

Claiming General Functions is Insufficient

Elec. Power Grp., 830 F.3d at 1353-54 (“[W]e have treated collecting information, including when limited to particular content (which does not change its character as information), as within the realm of abstract ideas.”); *Ericsson Inc. v. TCL Comm. Tech. Holdings Ltd.*, 955 F.3d 1317, 1328 (Fed. Cir. 2020) (“Merely claiming those functions in general terms, without limiting them to technical means for performing the functions that are arguably an advance, does not make a claim eligible at step one.”) (citations and internal quotations omitted).

Paper 31 (Opp. To MTA), 4.

Alice Step 1

Applying Generic Computer Elements to the Claims is Insufficient

The substitute claims' additional elements also fail to integrate the abstract idea into a practical application or otherwise cure their deficiencies under *Alice* Step 1. *TESLA1010*, ¶¶69-72. Elements [21.pre], [21.1], and [21.2] refer to generic computer components—including a “system,” “one or more processors” and a “memory”—on which the abstract idea is to be implemented. But merely applying an abstract idea to a computer does not confer eligibility under § 101. *Alice Corp. v. CLS Bank*, 573 U.S. 208, 221, (2014) (quoting *Mayo Collaborative Servs. v. Prometheus Labs., Inc.*, 566 U.S. 66, 72 (2012)); MPEP § 2106. That the

Paper 31 (Opp. To MTA), 6.

Alice Step 1

The Claims Recite Insignificant Post-solution Activity

MTA (new claim)

[21.5]at least in response to the anticipating, executing the first one or more instruction sets for operating the first device, wherein the first device or the second device autonomously performs one or more operations defined by the first one or more instruction sets for operating the first device;

MTA App. A, 2.

Opp. To MTA

Finally, the “executing” step (Elements [21.5], [21.7]) is merely incidental to the abstract idea itself and amounts to no more than insignificant post-solution activity. By the terms of claim 1, the “first one or more instructions sets” are explicitly instructions “for operating” a device. Simply executing the “first one or more instruction sets” to operate the first device according to the instructions’ stated function does nothing to meaningfully limit the application of the abstract idea or otherwise integrate it into a practical application. MPEP § 2106.05(g).

Paper 31 (Opp. To MTA), 7.

Alice Step 2

The Claims Lack an Inventive Concept

For example, as discussed above, the additional elements in PO's substitute claims consist of generic computer components, conventional sensors, and an insignificant post-solution executing step. *See, e.g.*, Elements [21.pre]-[21.1] ("A system comprising: one or more processors ..."), [21.2] ("a memory ..."; "detected at least in part by one or more sensors of the first device"). None of these elements, whether individually or in combination, can supply an "inventive concept" as required by *Alice* and its progeny. *Alice*, 573 U.S. at 223 ("[I]f a patent's recitation of a computer amounts to a mere instruction to 'implemen[t]' an abstract idea 'on . . . a computer,' that addition cannot impart patent eligibility."); *Elec. Power Grp.*, 830 F.3d at 1355 ("We have repeatedly held that such invocations of computers and networks that are not even arguably inventive are 'insufficient to pass the test of an inventive concept in the application' of an abstract idea."), *id.* ("do not even require a new source or type of information, or new techniques for analyzing it"); *Parker v. Flook*, 437 U.S. 584, 588-89 ("The notion that post-solution activity, no matter how

Paper 31 (Opp. To MTA), 8-9.

53

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Alice Step 2

Autonomously Executing Operations is Insufficient

PO Response

circumstances. *Id.*, 57:58-66; *see also* Ex. 2007, ¶¶47-48. The circumstance and instruction set information is stored together in a knowledgebase and used to operate the device autonomously. Ex. 1001, 57:58-66; *see also* Ex. 2007, ¶¶47-55 (Dr. Saber

POR, 2.

Dr. Saber Declaration

invention. Further, both references fail to contemplate the use of sophisticated data structures including the variously claimed correlations to enable autonomous operation.

EX2007, ¶[129].

Opp. To MTA

Dr. Saber contends that the claimed subject matter “use[s] ... sophisticated data structures including [] variously claimed correlations to enable autonomous operation.” *Id.*, ¶129. But automating conventional processes and instructing a computing system to apply an abstract idea are not the types of improvements that courts have deemed permissible under § 101. Like the claims at issue in *Electric Power Group*, here, “the focus of the claims is not on such an improvement in computers as tools, but on certain independently abstract ideas that use computers as tools.” *Elec. Power Grp.*, 830 F.3d at 1354. PO’s substitute claims in this

Paper 31 (Opp. To MTA), 10.

2024 AI Guidance Also Indicate Patent Ineligibility

Pet. Sur-Reply

to the examples of abstract ideas. For instance, the 2024 AI Guidance provides the following two examples of a method of organizing a human activity:

- Claims to “collect[ing] information on a user's movements and location history [and] electronically record[ing] that data;” and
- A claim to “monitoring the location of a mobile thing and notifying a party in advance of arrival of that mobile thing [].” 2024 Guidance, 58135.

The proposed substitute claims recite similar features that are directed to collecting circumstance and object representations to determine which instruction sets to apply when a user encounters similar circumstances. The “notifying a party

Paper 38 (Pet. Sur-Reply), 10.

The Claims Fail Under 35 USC 103

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Grotmol and Hickman Render the Substitute Claims Obvious

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

The Grotmol and Hickman Combination

A POSITA would have combined Grotmol and Hickman

Opp. To MTA

In addition to the reasons noted in the Petition, it would have been obvious to combine Grotmol and Hickman because the combination merely entails the simple substitution of one known element for another to obtain predictable results. TESLA1010, ¶76. For instance, applying Hickman's use of knowledge base(s) to facilitate sharing across a cloud computing system would have involved the modification of Grotmol's storage structure to include one or more knowledge bases with Hickman's data structures. TESLA1010, ¶76. The choice of whether to include one knowledge base, as in Grotmol, or various knowledge bases, as in Hickman, would be a designer's choice that was well-within the capabilities of a POSITA, as described in Hickman. TESLA1010, ¶76; TESLA1005, 10:57-11:7. The use of the

Paper 31 (Opp. To MTA), 14.

Hickman

example of one embodiment. The contents of the knowledge bases may be combined into a single common knowledge base or distributed across multiple different knowledge bases in a different fashion than described herein. Similarly, in some embodiments, the functionality of the cloud processing engine 305 may be combined with the functionality of the shared robot knowledge base 306 and/or other components of the cloud computing system 304. Additional systems and

TESLA1005, 10:63-11:3; TESLA1010: ¶[76].

Dr. Janet's Declaration

Hickman. TESLA1005, 10:57-11:7. The use of the data structures like those disclosed by Hickman (e.g., data structures 313-316) in Grotmol's system would have been accomplished by programming and memory storage techniques well within a POSITA's capabilities.

TESLA1010, ¶[76].

58

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Issue 8

GHC Renders Obvious Data Structures with Detected Location and Size Information

(all amended claims)

“ [21.10] (a *first data structure* generated at least by) ... processing one or more digital pictures to *detect*: a type of a first object, *a location of the first object, and a size of the first object*;

[21.13] (a *second data structure* generated at least by) processing one or more digital pictures to *detect*: a type of a second object, *a location of the second object, and a size of the second object*;

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

GHC Detects and Stores Location and Size Info

Hickman

The data associated with a particular object may vary from object to object. For example, most objects in the shared robot knowledge base have at least a baseline set of data that may include one or more of the object's (i) physical dimensions, weight, material composition, (ii) manufacturer (and possibly model number), (iii) ordering information (if the object is consumable), (iv) relevant tasks that a robot may perform with the object, (v) location of the object, etc. However, the shared object knowledge base may have additional types of data about the objects as well.

TESLA1005, 2:65-3:7; Paper 38 (Pet. Sur-Reply), 3.

For example, specification data (SPC) may include many different physical attributes of a particular "type" of object, e.g., physical dimensions (height, width, depth, shape, volume, area, etc.), weight, mass, color, texture, material composition, etc.

For example, a particular "type" of coffee mug may have a height (6 inches), width (4 inches), shape (cylindrical), volume (16 ounces), weight (1/4 pound), color (silver), texture (smooth), and material composition (stainless steel), etc.

TESLA1005, 18:10-18; Paper 38 (Pet. Sur-Reply), 3.

object knowledge base 307 shown in FIG. 3, an individual unique object identifier (UID) may have the following corresponding data: (i) a general object identifier (GID); (ii) image data (IMG); (iii) location data (LOC); and (iv) serial number data (SER). Various embodiments of the unique object knowledge base 307 may contain more or less data corresponding to a particular unique object identifier (UID). For

TESLA1005, 12:6-10; Paper 38 (Pet. Sur-Reply), 3.


GHC Detects and Stores Location and Size Info

The data structures include location and size information.

Pet. Sur-Reply

The fields in Hickman's knowledge bases 307-309 having corresponding data structures 313-315 include location data (LOC), model data (MOD), specification data (SPC), and task data (TASK) for objects. TESLA1005, FIG. 3 (reproduced below), 14:61-20:20. "The specification data (SPC) may include physical attributes associated with the particular 'type' of object corresponding to the general object identifier (GID)" including "e.g., physical dimensions (height, width, depth, shape, volume, area, etc.)." TESLA1005, 18:4-14, 2:65-3:21, 12:5-18, 14:1-30.

Paper 38 (Pet. Sur-Reply), 3.



Derived from TESLA1005 (Hickman), FIG. 5.

Paper 38 (Pet. Sur-Reply), 2.

Issue 9

GHC Renders Obvious Computer Models Indicating the Detected Sizes of the First and Second Objects

(all amended claims)

“[21.11]in response to the processing, generating the first data structure that includes: ... *a computer model indicating the detected size of the first object;*

[21.14]in response to the processing, generating the second data structure that includes: ... *a computer model indicating the detected size of the second object;*”

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

GHC Renders Obvious Computer Models Indicating Size Information

Pet. Sur-Reply

Hickman's SPC field (in conjunction with data stored in the model data (MOD) field) provides a computer model indicating the size of an object. *Id.*; Opposition, 20-22. PO is mistaken, however, because as described in the preceding section,

Paper 38 (Pet. Sur-Reply), 5.

Opp. to MTA

generated in knowledge base 306 (including a combination of knowledge bases) as part of the "*the first object representation*" and would have included "*a first data structure*" that includes the type, location (coordinates), size of the first object, along with a label indicating the detected type of the first object and a computer model (combination of Hickman's SPC and MOD) indicating the detected size of the first object, as recited in [21.9]-[21.11]. TESLA1010, ¶90.

Paper 38 (Opp. to MTA), 21-22.

Hickman

For example, specification data (SPC) may include many different physical attributes of a particular "type" of object, e.g., physical dimensions (height, width, depth, shape, volume, area, etc.), weight, mass, color, texture, material composition, etc.

TESLA1005, 18:10-14; Opp. to MTA, 21.

Similarly, the model data (MOD) may include a model number, variant, or other useful designation assigned to a particular type of object by the object's manufacturer, maker, builder, producer, etc.

TESLA1005, 17:30-33; Opp. to MTA, 21.

313					314				
UID ₁	GID	IMG	LOC	SER	GID ₁	LNG	IMG	MFG	MOD SPC TASKS
UID ₂	GID	IMG	LOC	SER	GID ₂	LNG	IMG	MFG	MOD SPC TASKS
UID ₃	GID	IMG	LOC	SER	GID ₃	LNG	IMG	MFG	MOD SPC TASKS
...					...				
UID _n	GID	IMG	LOC	SER	GID _n	LNG	IMG	MFG	MOD SPC TASKS

UID ₁	GID ₁	IMG	LOC	SER	LNG	IMG	MFG	MOD SPC TASKS
UID ₂	GID ₂	IMG	LOC	SER	LNG	IMG	MFG	MOD SPC TASKS
UID ₃	GID ₃	IMG	LOC	SER	LNG	IMG	MFG	MOD SPC TASKS
...					...			
UID _n	GID _n	IMG	LOC	SER	LNG	IMG	MFG	MOD SPC TASKS

Derived from TESLA1005 (Hickman), FIG. 5.

Paper 38 (Pet. Sur-Reply), 2.

63

FISH

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Issue 10

GHC Renders Obvious First and Second Data Structures Correlated with Instruction Sets

(all amended claims)

“[21.15] wherein the first data structure and the second data structure *are correlated with* the first one or more instruction sets for operating the first device.”

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

GHC Renders Obvious Correlated Data Structures

Both data structures have correlated instruction sets (Tasks).

Opp. To MTA

Hickman teaches that for each detected object, the knowledge base may store “relevant tasks that a robot may perform with the object” in the task data (TASK) field. TESLA1005, 2:65-3:8, 15:28-43, 18:34-20:20; FIG. 3. Accordingly, a


Paper 31 (Opp. To MTA), 24.

Hickman

The data associated with a particular object may vary from object to object. For example, most objects in the shared robot knowledge base have at least a baseline set of data that may include one or more of the object’s (i) physical dimensions, weight, material composition, (ii) manufacturer (and possibly model number), (iii) ordering information (if the object is consumable), (iv) relevant tasks that a robot may perform with the object, (v) location of the object, etc. However, the shared object knowledge base may have additional types of data about the objects as well.

TESLA1005, 2:65-3:8; Paper 31 (Opp. To MTA), 24.

Hickman



Derived from TESLA1005 (Hickman), FIG. 5.

Paper 38 (Pet. Sur-Reply), 2.

Grotmol and Tolstedt Render the Substitute Claims Obvious

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Overview of Tolstedt

[0083] Sensor processing algorithms 604 receives sensor data from sensor system 400 and classifies the sensor data into thematic features. This classification may include identifying objects that have been detected in the environment. For example, sensor processing algorithms 604 may classify an object as a person, curb, tree, waste container, light pole, driveway, or some other type of object. The classification may be performed to provide information about objects in the environment. This information may be used to generate a thematic map, which may contain a spatial pattern of attributes. The attributes may include classified objects. The classified objects may include dimensional information, such as, for example, location, height, width, color, and other suitable information. This map may be used to plan actions for the vehicle. The action may be, for example, planning paths to follow an operator in a side following mode or performing object avoidance.




FIG. 4

TESLA1015, FIG. 4, ¶[83]; Paper 31 (Opp. To MTA), 17.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

The Grotmol and Tolstedt Combination

A POSITA would have combined Grotmol and Tolstedt

82. A POSITA would have combined Grotmol and Tolstedt as doing so would have involved combining prior art elements according to known methods to yield predictable results. Here, applying Tolstedt's known classification method involving the use of (1) a knowledge base that provides information on the type and location of objects, along with instructions associated with these objects, and (2) the generation of a thematic map that includes the object information into Grotmol's system would have yielded the predictable result of improving Grotmol's object recognition and matching capabilities through the use of a more detailed classification system, and the rendering of a thematic map that would better facilitate object avoidance when the robot is a vehicle.

TESLA1010, ¶[82]; Paper 31 (Opp. To MTA), 18.

68

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Issue 11

GTC Renders Obvious First and Second Data Structures with Detected Location and Size Information

(all amended claims)

“ [21.10] (a *first data structure* generated at least by) ... processing one or more digital pictures to *detect*: a type of a first object, *a location of the first object, and a size of the first object*;

[21.13] (a *second data structure* generated at least by) processing one or more digital pictures to *detect*: a type of a second object, *a location of the second object, and a size of the second object*;

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

GTC Renders Obvious Detecting Size and Location of Objects

Pet. Sur-Reply

Guidance reached an incorrect conclusion with respect to the GTC grounds. And in adopting the analysis from the Preliminary Guidance, PO's Reply fails to appreciate that identifying attributes such as location, height, width, and dimension information when processing images would necessarily involve detecting these attributes. TESLA1015, ¶¶83-88. The Preliminary Guidance appears to require Petitioner to

Paper 38 (Pet. Sur-Reply), 7.

Dr. Janet Declaration

object (e.g., performing object avoidance). TESLA1015, ¶¶83-88. A POSITA would have understood that this information is stored in data fields (which are a type of data structure, as explained in the '344 patent) in a knowledge base. TESLA1001, 65:55-66:6. *Id.* Each detected object in image would have its own data structure.

TESLA1010, ¶[91].

Tolstedt

[0083] Sensor processing algorithms 604 receives sensor data from sensor system 400 and classifies the sensor data into thematic features. This classification may include identifying objects that have been detected in the environment. For example, sensor processing algorithms 604 may classify an object as a person, curb, tree, waste container, light pole, driveway, or some other type of object. The classification may be performed to provide information about objects in the environment. This information may be used to generate a thematic map, which may contain a spatial pattern of attributes. The attributes may include classified objects. The classified objects may include dimensional information, such as, for example, location, height, width, color, and other suitable information. This map may be used to plan actions for the vehicle. The action may be, for example, planning paths to follow an operator in a side following mode or performing object avoidance.

TESLA1015, ¶[83]; Paper 31 (Opp. To MTA), 20-21.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

70

GTC Renders Obvious Detecting the Location of Objects

Opp. To MTA

Tolstedt also teaches that the **Global Positioning System** is used to obtain **location information**. TESLA1015, ¶¶101, 120, 130, 31, 68-69, 73. Thus, it would

Paper 31 (Opp. To MTA), 22.

Tolstedt

914, machine control process sends a command to a sensor component, **such as the global positioning system**, to detect the location of the vehicle. Sensor information **918** is received from the **global positioning system**, and **sensor processing algorithms 916** processes the information to identify the **location of the vehicle on a map**, such as a map of the operating environment stored in knowledge base **910**. Path planning module **904** then records the location of the vehicle at each waypoint and the waypoints are associated with the path and stored in knowledge base **910**. At a future time, user input **914**

TESLA1015, ¶¶101; Paper 31 (Opp. To MTA), 22.

PO's Characterization of Tolstedt Lacks Credibility

PO Reply to Opp. To MTA

the recited first data structure analysis to be even possible. Thus, *Tolstedt does not disclose any data structure*, let alone the first data structures with all of its recited limitations.

PO Reply to Opp. To MTA), 10.

GTC Renders Obvious “Digital Pictures”

Pet. Sur-Reply

To the extent, PO questions whether Tolstedt discloses digital images, Tolstedt teaches that its images are acquired through a machine that controls the camera. TESLA1015, ¶82. Acquired camera images are processed for object classification (e.g., using sensor processing algorithms) in a computer system. *Id.* The images processed in Tolstedt’s system would have been digital pictures because such processing and classification would only be possible if the images were digital pictures.

Paper 38 (Pet. Sur-Reply), 8.

GTC Detects Objects When Classifying The Objects

Pet. Sur-Reply

Moreover, Tolstedt explains that different techniques can be used to classify objects in an image. For example, an autonomous classification or input from a user can be received to detect the size of an object. TESLA1015, ¶84. A user can input a size, and the computer detects the size upon processing the user input. Location can be detected using lidar, sensors, and other algorithms. TESLA1015, ¶84. Thus, to the extent how detection is performed is required, Tolstedt teaches several possible methods to detect a size and location of an object

Paper 38 (Pet. Sur-Reply), 8.

Tolstedt

[0084] The classification may be done autonomously or with the aid of user input through user interface 606. User interface 606 may be, in one illustrative embodiment, presented on a display monitor mounted on a side of a vehicle and viewable by an operator. User interface 606 may display ... in an illustrative embodiment, sensor processing algorithms 604 receives data from a laser range finder, such as two dimensional/three dimensional lidar 406 in FIG. 4, identifying points in the environment. The information processed by sensor processing algorithms 604 is displayed to an operator through user interface 606. User input may be received to associate a data classifier with the points in the environment, such as, for example, a data classifier of “curb” associated with one point, and “street” with another point. Curb and street are examples of thematic features in an environment. Sensor processing algorithms 604 then interacts with knowledge base 608 to locate the classified thematic features on a

TESLA1015, ¶[84]; Paper 38 (Pet. Sur-Reply), 7-8.

Issue 12

GTC Renders Obvious Computer Models Indicating the Detected Sizes of the First and Second Objects

(all amended claims)

“[21.11]in response to the processing, generating the first data structure that includes: ... *a computer model indicating the detected size of the first object*;

[21.14]in response to the processing, generating the second data structure that includes: ... *a computer model indicating the detected size of the second object*,”

GTC's "Thematic Map" Corresponds to a Computer Model in A Data Structure

PO's argument that Tolstedt's Map is Not a "Data Structure" fails in light of the '344 Patent's Broad Description of a "Data Structure."

'344 Patent

Where a reference to a data structure is used herein, it should be understood that any variety of data structures can be used such as, for example, array, list, linked list, doubly linked list, queue, tree, heap, graph, map, grid, matrix, multi-dimensional matrix, table, database, database management system (DBMS), file, neural network, and/or any other type or form of a data structure including a custom one. A data structure may include one or more fields or data fields that are part of or associated with the data structure. A field or data field may include a data, an object, a data structure, and/or any other element or a reference/pointer thereto. A data structure can be stored in one or more memories, files, or other repositories. A data structure and/or elements thereof, when stored in a memory, file, or other repository, may be stored in a different arrangement than the arrangement of the data structure and/or elements thereof. For

TESLA1001, 65:45-57; Paper 38 (Pet. Sur-Reply), 8-10.

Tolstedt

[0083] Sensor processing algorithms 604 receives sensor data from sensor system 400 and classifies the sensor data into thematic features. This classification may include identifying objects that have been detected in the environment. For example, sensor processing algorithms 604 may classify an object as a person, curb, tree, waste container, light pole, driveway, or some other type of object. The classification may be performed to provide information about objects in the environment. This information may be used to generate a thematic map, which may contain a spatial pattern of attributes. The attributes may include classified objects. The classified objects may include dimensional information, such as, for example, location, height, width, color, and other suitable information. This map may be used to plan actions for the vehicle. The action may be, for example, planning paths to follow an operator in a side following mode or performing object avoidance.

TESLA1015, ¶[83]; Paper 31 (Opp. To MTA), 20-21 ; Paper 38 (Pet. Sur-Reply), 8-10.

76

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Issue 13

GTC Renders Obvious First and Second Data Structures Correlated with Instruction Sets

(all amended claims)

“[21.15] wherein the first data structure and the second data structure *are correlated with* the first one or more instruction sets for operating the first device.”

GTC Renders Obvious Correlated Data Structures

Tolstedt discloses providing correlated instructions for object avoidance.

Opp. To MTA

Tolstedt’s system also determines actions (“the first one or more instruction sets for operating the first device”) for a vehicle to take in association with a detected object (e.g., performing object avoidance). TESLA1015, ¶¶83-88. In addition to storing classification data on detected objects in the sensor data, Tolstedt’s knowledge base 800 also provides information to a machine control process 602 to “determine which commands [“instruction sets”] to send to the vehicle’s steering, braking, and propulsion components.” TESLA1015, ¶84. It

Paper 31 (Opp. To MTA), 24-25.

Tolstedt

able information. This map may be used to plan actions for the vehicle. The action may be, for example, planning paths to follow an operator in a side following mode or performing object avoidance.

TESLA1015, ¶[83]; Paper 31 (Opp. To MTA), 24.

with the landmark localization. Machine control process 602 receives the environmental data from sensor processing algorithms 604, and interacts with knowledge base 608 and behavior library 610 in order to determine which commands to send to the vehicle’s steering, braking, and propulsion components.

TESLA1015, ¶[84]; Paper 31 (Opp. To MTA), 24.