

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

Case No. IPR2023-01054
U.S. Patent No. 10,452,974
Before Barbara Parvis, Robert Weinschenk, and Russell Cass
Administrative Patent Judges

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Table of Contents



• The '974 Patent and Petition Overview	3
• IPR (Petition) Challenge	8
• Overview of Grotmol, Hickman, and Combination	8
• Issue 1: Grotmol Discloses [4.3]	19
• Issue 2: Grotmol Discloses [5]	29
• Issue 3: Grotmol Discloses the Claimed “connection[s]” in [7]	33
• Issue 4: GHC Renders Claims 4 and 11 Obvious	39
• Motion to Amend	44

'974 Patent and Petition Overview

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

The 974 Patent



TESLA1001, cover.

TESLA1001 ('974 pat.), FIGS. 4A, 4B; Petitioner Reply (Pet. Rep.) 25.

Asserted Grounds

Claim(s) Challenged	35 U.S.C. §	Reference(s)/Basis
4–5, 7, 14–15	102(a)(2)	Grotmol
11	103	Grotmol in view of the knowledge of a person of ordinary skill
4, 11	103	Grotmol, Hickman

ID., 9; Pet., 3-4.

Claims 1-3, 6, 8-10, 12, 13, and 16-20 were disclaimed. POPR, 1.

Claims 4–5, 7, 11, and 14–15 remain pending. Institution Decision (I.D.), 1.

'974 Patent Claim 4

1. A method implemented using a computing system that includes one or more processor circuits, the method comprising:

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

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

anticipating the first one or more instruction sets for operating the first device learned in the learning process based on at least partial match between the third circumstance representation and the first circumstance representation; and

at least in response to the anticipating, executing the first one or more instruction sets for operating the first device learned in the learning process, 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 learned in the learning process.

4. The method of claim 1, wherein the learning process includes:

generating or receiving the first circumstance representation, and generating or receiving the second circumstance representation; and

obtaining or receiving the first one or more instruction sets for operating the first device, and obtaining or receiving the second one or more instruction sets for operating the first device, and wherein the knowledgebase further includes a fourth correlation including a fourth circumstance representation correlated with a fourth one or more instruction sets for operating the second device, and wherein the fourth circumstance representation represents a fourth circumstance detected at least in part by the one or more sensors of the second device, and wherein at least a portion of the fourth correlation is learned in another learning process that includes:

generating or receiving the fourth circumstance representation; and

obtaining or receiving the fourth one or more instruction sets for operating the second device.

TESLA1001, 169:36-170:37.

Additional Disputed Claims

5. The method of claim 1, wherein the anticipating the first one or more instruction sets for operating the first device learned in the learning process based on the at least partial match between the third circumstance representation and the first circumstance representation includes:

determining that a number of at least partially matching portions of the third circumstance representation and portions of the first circumstance representation exceeds a threshold number, or

determining that a percentage of at least partially matching portions of the third circumstance representation and portions of the first circumstance representation exceeds a threshold percentage.

TESLA1001, 170:38-50.

7. The method of claim 1, wherein the knowledgebase further includes a fourth correlation including a fourth circumstance representation correlated with a fourth one or more instruction sets for operating: the first device, the second device, or a third device, and wherein a first connection is generated to connect the first correlation with the second correlation, and wherein a second connection is generated to connect the second correlation with the fourth correlation.

TESLA1001, 170:59-67.

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; TESLA1003, ¶¶[47]-[51]; Pet., 7-9; I.D., 23.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Grotmol

TESLA1004



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 to 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. 20A

TESLA1004, FIG. 20A.

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.

The given image may be compared to images of the primary particle set. In some implementations, the primary particle set may comprise a previously used image I0 (e.g., 2000 in FIG. 20A), one or more (e.g., 2) images preceding the I0 image in time in the training buffer, and one or more (e.g., 5) images following the I0 image in time in the training buffer.

TESLA1004, 39:65-40:5.; MTA Sur-Reply, 3-4.

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.

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.*,

TESLA1003, ¶[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. 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 '974 Patent's Structure

'974 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

Petition

Hickman teaches a shared robot knowledge base that could advantageously be paired with Grotmol's autonomous robotic system in order to further facilitate the sharing of learned behaviors.

Pet., 69.

In the combination of Grotmol and Hickman, the shared knowledge based would include multiple training sets (i.e., circumstance representations correlated with instruction set) learned during the operation of multiple devices. Accordingly,

Pet., 63.

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 feature [4.3]

*[4.3] ... wherein the fourth circumstance representation represents a fourth circumstance detected at least in part by the one or more sensors of **the second device**, and wherein at least a portion of the fourth correlation is learned in **another learning process** ...*

Grotmol Discloses “a fourth 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., 27-28; TESLA1003, ¶¶[93]-[94]..

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, 27-28; TESLA1003, ¶¶[67]-[73], [93], [94]..

20

Grotmol Discloses “*the second device*” and “*another learning process*”

DR. JANET’S DECLARATION.

44. As I explained above in Section II.A, Grotmol teaches that a second device can provide training information and learned behaviors for loading onto the first device. TESLA1004, 33:21-35. 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, 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 trainers.

Pet., 27-28; Pet. Rep., 2; TESLA1010, ¶¶[44].

12. A POSITA would have understood Grotmol’s disclosure to mean 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 (first) robot can benefit from the learning of the transmitting (second) robot and also perform “learned behaviors.” See e.g., 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 second device. TESLA1004, 33:32-35 (“**match process configuration**”).

Pet., 27-28; Pet. Rep., 2; TESLA1010, ¶¶[12].

Grotmol Discloses “*the second device*” and “*another learning process*”

PET. REPLY

TESLA1010, ¶¶10-12, 43-45. The result is a knowledgebase on the first robot that includes both first and second correlations learned in a learning process on the first robot *and* a fourth correlation learned in another learning process on the second robot. *Id.*

Pet. Rep., 2..

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

Patent Owner Response

Grotmol only discloses *one* user operating *one* device in *one* learning process.

Ex. 2007, ¶¶ 74-78. The knowledge gained from this process in Grotmol, when

POR, 16.

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., 7.

GROTMOL




FIG. 10B

TESLA1004, FIG. 10B; Pet. Rep., 4, 7.

PO's Characterization of Grotmol is Incorrect On its Face

Grotmol explicitly teaches multiple users.

PO Sur-Reply

1. Reply, 7-8. Petitioner has failed to point to a single disclosure of multiple users using a single device in Grotmol (much less multiple users using multiple devices in multiple learning processes), and instead states without explanation that a POSA would have understood that user commands can be provided by multiple users. Reply, 7.

POSR, 9-10.

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.

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., 27-28; TESLA1003, ¶¶[93]-[94]; Pet. Rep., 4-5.



TESLA1004, FIG. 10B; Pet. Rep., 4, 7.

PO's Rationale is Inconsistent with the Law

Pet. Reply

additional training on the first robot. PO's and Dr. Saber's overly rigid analyses fail to account for the reasonable inferences that a POSITA would bring to bear in reviewing Grotmol, and they hence follow an approach that is inconsistent with well-established law of anticipation. See, e.g., *In re Predai*, 401 F.2d 825, 826 (CCPA 1968) ("it is proper to take into account not only specific teachings of the reference but also the inferences which one skilled in the art would reasonably be expected to draw therefrom"); see also *Eli Lilly v. Los Angeles Biomedical Research Inst.*, 849 F.3d 1073, 1074-1075 (Fed. Cir. 2017); *In re Graves*, 69 F.3d 1147, 1152 (Fed. Cir. 1995); *In re Paulsen*, 30 F.3d 1475, 1480 (Fed. Cir. 1994).

Pet., Rep., 9.

PO's Arguments

PO when arguing against Grotmol alone

could be done by different users.” Petition, 57-58. Petitioner is wrong since this argument describes a situation involving *two* trained configurations each stored and learned on a different robot, whereas claim 11 recites a *single knowledgebase* that includes *mixed knowledge of correlations* (i.e., first and second) learned in one learning process involving one device and one user and an *additional correlation* (i.e., fourth) learned in another learning process involving a second device and another user. Ex. 2007, ¶¶93-103. Grotmol also never discloses that the second

POR, 28.

PO when arguing against Grotmol and Hickman

and 11 and shown below). Ex. 2007, ¶¶104-105. 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*. *Id.* Further, Hickman teaches away from the claimed

POR, 35.

Issue 2

Grotmol Discloses [5]

*[5] ... determining that **a number [or percentage] of at least partially matching portions** of the third circumstance representation and portions of the first circumstance representation exceeds a threshold number [or percentage] ,...*

PO Relies on Importing Unclaimed Features

Patent Owner Response

on, does not disclose the elements of claim 5. Petition, 29. Using the number of at least partially matching portions in *the matching process involves counting the portions and determining whether this counted number of portions exceeds a numerical threshold*, and Grotmol does not disclose any such counting of portions and numerical threshold. Ex. 2007, ¶¶79-83. Further, using the percentage of at least partially matching portions in *the matching process involves counting the portions to compute the percentage and determining whether this computed percentage of portions exceeds a percentage threshold*, and Grotmol does not disclose any such counting of portions, computation, and percentage threshold, either. *Id.*

POR, 20-21.

30

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Grotmol Discloses Matching Portions of First and Third Circumstances

GROTMOL

During autonomous operation, control process of the robot may be configured to compare a given (e.g., the most recent, current) image with one or more the images from the training buffer. In some implementations, the matching process may comprise comparing the given image to every image in the training buffer.

TESLA1004, 39:23-28; Pet., 29.

For computational efficiency reasons, it may not be desirable and/or feasible to compare each new camera image with every one of the stored images seen during training, according to some implementations. The robot may take advantage of the prior information about what are the likely regions of the path where it might be located, and only search those regions. The robot may search a random sample of other regions in case the prior information is inaccurate or invalidated for some reason.

TESLA1004, 39:29-37; Pet., 29.

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

TESLA1004, 41:8-14; Pet., 29; Pet., Rep., 10.

than the quantity of features within the set of features. The method may comprise, based on the comparison, determining a similarity measure for a given training set of the plurality of training feature sets. The similarity measure may characterize similarity between features of the subset and features of the given training set. The method may comprise, responsive to the similarity measure breaching a threshold, selecting one or more training sets from the plurality of training sets. The method may comprise determining one or

TESLA1004, 4:30-38; Pet., 29; Pet., Rep., 10.

Grotmol Discloses Matching Portions of Images

GROTMOL

In some implementations, history of the image matching process may be utilized in order to determine best match image. By way of an illustration, if a match search set image with the best match score (e.g., the lowest norm) belongs to the primary particle set than it may be selected as the best match. If the image with the best match score belongs to the secondary particle set than it may be selected based on an evaluation of image history parameter. In some implementations, image history parameter evaluation may be performed as follows:

- (i) a running window average match score may be determined by averaging over last N images within individual particle sets. In some implementations, the aver-
- ...
- (ii) the average match score for individual secondary particle sets may be compared to individual match scores from the match search set;

TESLA1004, 41:33-60; Pet., Rep., 10.

32

Issue 3

Grotmol Discloses the Claimed “connection[s]” in [7]

[7] ... wherein a *first connection* is generated to connect the first correlation with the second correlation, and wherein a *second connection* is generated to connect the second correlation with the fourth correlation.

PO's Construction is Unclear and Unsupported

PO Response

a “connection” should be construed as a “relationship that includes a link.” Ex. 2007, ¶¶55-62.

POR, 11.

Pet. Rep.

In addition to the lack of support, PO's construction is problematic because PO's construction includes the term “link”—the meaning of which is unclear. The '974 patent does not define a link, and during his deposition, PO's expert likewise failed to define or clarify what the meaning of “link” over almost 40 minutes of questioning. TESLA1012, 12:12-36:5; TESLA1010, ¶¶32-33.

“connection.” POR, 14. 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 '974 patent does not even use the word “link” to describe a “connection.” Even if a connection may include a link in some

Pet. Rep., 15.

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, 14; Pet. Rep., 15.

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, 14; Pet. Rep., 15.

PO's arguments contradict the POM

Response, a sequence of elements inherently does not need any connections since the sequence is an ordered group of elements of a single path (i.e., sequence),

POSR, 17.

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Examples and Embodiments of “Connection” in the ’974 patent are non-limiting

Pet. Rep.

PO instead 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. EX2007, ¶¶55-62; POR, 11. But even these examples are not specifically described in terms of “connections.” TESLA1010, ¶¶30-31. For example, PO turns to FIGS. 17 and 18A, which are described as “embodiments,” and which do not impose a limiting effect. TESLA1001, 56:17-23. The ’974 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, 103:42-51. 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, 106:11-14. The ’974 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, 15:58-60. 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., 14.

’974 Patent

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

TESLA1001, 56:17-23.

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

TESLA1001, 15:58-60.

Grotmol Discloses the Claimed “*connection[s]*”

'974 Patent

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, 153:66-154:6; Pet., 32.

Grotmol

The robot may comprise a controller configured to store a sequence of images obtained during training phase. In some implementations, the images may be stored in an ordered buffer in a non-volatile memory (e.g., memory 814 in FIG. 8).

TESLA1004, 38:38-42; Pet., 32.

ing in an ordered buffer. Motor commands associated with a given image may be stored. During autonomous operation,

TESLA1004, Abstract; Pet., 32.

FISH.


DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Even under PO's Construction, Grotmol Discloses the Claimed “*connection[s]*”

Pet. Reply

Further, even if a “connection” required a link as urged by PO (which it does not), a POSITA would have understood the connections provided by Grotmol’s ordered training buffer to link temporally proximate images (e.g., to a particle image). TESLA1010, ¶40. If the images in the buffer were not linked, temporally proximate images would not be identified as they are in Grotmol. As noted above,

Pet., Rep., 17.



Pet., Rep., 18.

Grotmol

In some implementations, the search space may be narrowed using a form of **particle filtering**, where the robot maintains a plurality of particles indicating the likely parts of the path. That is, individual particle points at a particular image from the training buffer. As a new camera image arrives, the robot may search those images in the training buffer which are close to the particles. **Individual particles may be moved to a nearby location in the training buffer where the stored image matches closely with the newly arrived image.** Particles with poor match with the new image may be deleted. New particles may be created, either in the vicinity of the other particles, or from randomly sampled locations in the training buffer, shown in FIG. 22.

TESLA1004, 39:42-54; Pet. Rep., 16-17.

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Issue 4

The Grotmol-Hickman Combination (GHC) Renders Claims 4 and 11 Obvious

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

PO's Characterization of Hickman is Incorrect On its Face

Hickman is not limited to separate knowledge bases.

PO Response

and 11 and shown below). Ex. 2007, ¶¶104-105. 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. *Id.* Further, Hickman teaches away from the claimed

POR, 35.

However, Hickman teaches away from the claimed "fourth correlation . . ." (with all of its associated features as claimed in claims 4 and 11) by disclosing images and instructions in separate knowledgebases that are not correlated. *Id.* Hickman also

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: ¶¶67].

40

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

PO’s Characterization of Hickman is Incorrect On its Face

Hickman is not limited to separate knowledge bases.

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

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: ¶[68].

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

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

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., 2-324.

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

42

PO's Characterization of Hickman is Incorrect On its Face

Hickman is not limited to single object information

PO Response

correlated and stored *together*. *Id.* Further, Hickman teaches away from the claimed “fourth circumstance representation” because Hickman focuses on manipulating a single object and uses single object information and such single object information is not a circumstance representation. *Id.* In addition, instances of single object

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: ¶[55].

Motion to Amend

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Table of Contents

• Motion To Amend	44
• Issue 5: The Claims Fail Under 35 USC 101	48
• Grotmal/Hickman (35 USC 103)	58
• Issue 6: First and Second Data Structures	60
• Issue 7: Data Structures Correlated with Instruction Sets	63
• Grotmol/Tolstedt (35 USC 103)	65
• Issue 8: First and Second Data Structures	68
• Issue 9: Data Structures Correlated with Instruction Sets	72

Summary of Amendments

The amendments rewrite each of claims 4/5/7/11/14/15 in independent form and further limit the method of each of these claims by requiring that:

1. the first circumstance representation includes a first collection of object representations,
2. the second circumstance representation includes a second collection of object representations;
3. the first collection of object representations represents multiple objects detected at a first time or during a first time period,
4. the second collection of object representations represents the multiple objects detected at a second time or during a second time period;
- 5. the first and second collections of object representations comprise first and second data structures, respectively, which each include labels indicating types of the multiple objects and also include coordinates indicating locations of the multiple objects**
- 6. the first and second data structures be correlated with the first one or more instruction sets and the second one or more instruction sets for operating the first device, respectively.**

Substitute Claim 21 (Amended claim 4)

Claim 21 (Proposed substitute for claim 4, if found unpatentable) A method implemented using a computing system that includes one or more processor circuits, the method comprising:

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

generating or receiving a third circumstance representation, wherein the third circumstance representation represents a third circumstance detected at least in part by the one or more sensors of the first device or at least in part by one or more sensors

of a second device;

anticipating the first one or more instruction sets for operating the first device learned in the learning process based on at least partial match between the third circumstance representation and the first circumstance representation; and

at least in response to the anticipating, executing the first one or more instruction sets for operating the first device learned in the learning process, 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 learned in the learning process;

wherein the learning process includes:

generating or receiving the first circumstance representation, and

generating or receiving the second circumstance representation; and

obtaining or receiving the first one or more instruction sets for operating the first device, and

obtaining or receiving the second one or more instruction sets for operating the first device, and

wherein the knowledgebase further includes a fourth correlation including a fourth circumstance representation correlated with a fourth one or more instruction sets for operating the second device, and wherein the fourth circumstance representation represents a fourth circumstance detected at least in part by the one

or more sensors of the second device, and wherein at least a portion of the fourth correlation is learned in another learning process that includes:

generating or receiving the fourth circumstance representation; and

obtaining or receiving the fourth one or more instruction sets for operating the second device;

wherein the first circumstance representation includes a first collection of object representations, and the second circumstance representation includes a second collection of object representations;

wherein the first collection of object representations represents multiple objects detected at a first time or during a first time period, and comprises a first data structure that includes: labels indicating types of the multiple objects, and coordinates indicating locations of the multiple objects;

wherein the second collection of object representations represents the multiple objects detected at a second time or during a second time period, and comprises a second data structure that includes: labels indicating types of the multiple objects, and coordinates indicating locations of the multiple objects;

wherein the first data structure is correlated with the first one or more instruction sets for operating the first device; and

wherein the second data structure is correlated with the second one or more instruction sets for operating the first device.

Issue 5

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 determining a manner 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, ¶¶60-61. This familiar concept 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

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, ¶¶62-63. 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

1. Elements [1.1]-[1.2] refer to “a computing system that includes one or more processor circuits” and a “memory” that stores the knowledgebase, but recitation of generic computer components and mere instructions to implement the abstract idea on a computer are insufficient to confer eligibility under § 101. *Alice Corp. v. CLS Bank*, 573 U.S. 208, 221, (2014) (“Stating an abstract idea while adding the words ‘apply it with a computer’ ... [yields] the same deficient result” and does not “provide[] any practical assurance that the process is more than a drafting effort designed to monopolize the abstract idea itself.”) (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)

at least in response to the anticipating, executing the first one or more instruction sets for operating the first device learned in the learning process, 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 learned in the learning process;

MTA App. A, 2.

Opp. To MTA

Finally, the “executing” step (Element [1.6]) 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 or second 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

an insignificant post-solution executing step. See, e.g., Elements [1.1] (“a computing system that includes one or more processor circuits”), [1.2] (“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,

Paper 31 (Opp. To MTA), 8.

54

Alice Step 2

Autonomously Executing Operations is Insufficient

PO Response

2. Advantages Provided by the '974 Patent

The '974 patent is directed to systems, methods, and computer readable media that store a knowledgebase that correlates device operating instruction sets with representations of objects encountered by a device, obtain new object representations using device sensors, match the new objects with objects in the knowledgebase, and cause a device that detected the new objects to autonomously execute an instruction set correlated to the matched objects. Ex. 1001, 1:39-2:9. This

POR, 3.

Opp. To MTA

performed autonomously.” *Id.* But automating conventional processes or 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

Paper 31 (Opp. To MTA), 9.

2024 AI Guidance Also Indicate Patent Ineligibility

Pet. SurReply

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. SurReply), 11.

56

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

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, ¶68. 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. *Id.* 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. *Id.*; TESLA1005, 10:57-11:7. The use of the data structures

Paper 31 (Opp. To MTA), 13.

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].

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, ¶[68].

59

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Issue 6

GHC Renders Obvious the First and Second Data Structures

(all amended claims)

“wherein the first collection of object representations represents multiple objects detected at a first time or during a first time period, and comprises a first data structure that includes: labels indicating types of the multiple objects, and coordinates indicating locations of the multiple objects;

wherein the second collection of object representations represents the multiple objects detected at a second time or during a second time period, and comprises a second data structure that includes: labels indicating types of the multiple objects, and coordinates indicating locations of the multiple objects.”

FISH 

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

GHC Renders Obvious the Two Data Structures

The SAME objects are included in both data structures.





FIG. 20A

TESLA1004, FIG. 20A.



GHC Renders Obvious the Two Data Structures

The SAME objects are included in both data structures.

It would have been obvious to a POSITA that, in the GHC, a row or set of rows in Hickman's structure 314 (or 313) that include Grotmol's image 2000 and objects therein may correspond to the "first data structure," and a row or set of rows in Hickman's structure 314 (or 313) that include Grotmol's image 2010 and objects therein may correspond to the "second data structure." The depiction of objects in the images together with the first and second data structures provide the claimed "first collection of object representations" and "second collection of object representations," respectively.

TESLA1010, ¶[87]; Paper 38 (Pet. SurReply), 5.

Issue 7

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

(all amended claims)

“wherein the first data structure *is correlated with the first one or more instruction sets* for operating the first device; and;
wherein the second data structure *is correlated with the second one or more instruction sets* for operating the first device. ”

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 / Pet. SurReply



Paper 38 (Pet. SurReply), 3-5.

Grotmol and Tolstedt Render the Substitute Claims Obvious

(all amended claims)

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

The Grotmol and Tolstedt Combination

A POSITA would have combined Grotmol and Tolstedt

74. 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 800 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 incorporating the object information 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, ¶[74]; Paper 31 (Opp. To MTA), 17.

67

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Issue 8

GTC Renders Obvious the First and Second Data Structures

(all amended claims)

“wherein the first collection of object representations represents multiple objects detected at a first time or during a first time period, and comprises a first data structure that includes: labels indicating types of the multiple objects, and coordinates indicating locations of the multiple objects;

wherein the second collection of object representations represents the multiple objects detected at a second time or during a second time period, and comprises a second data structure that includes: labels indicating types of the multiple objects, and coordinates indicating locations of the multiple objects.”

FISH 

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

GTC Renders Obvious the Two Data Structures

The SAME objects are included in both data structures.




FIG. 20A

TESLA1004, FIG. 20A.

Paper 38 (Pet. SurReply), 3-5.

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.

69

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

GTC Renders Obvious the Two Data Structures

The data structures may be a knowledgebase or map with information on the SAME objects.

Opp. To MTA

performing object avoidance). TESLA1015, ¶¶83-88. This information is stored in a data structure in a knowledge base. *Id.*; TESLA1010, ¶82.

In some cases, this data structure may also be a map that the vehicle uses for navigation purposes and includes labels for the types and locations of objects.

TESLA1015, ¶¶83-88 (“[t]his information may be used to generate a thematic map,

Paper 31 (Opp. To MTA), 21.

Pet. SurReply

Returning to Grotmol’s images 2000 and 2010 in FIG. 20A, as explained above, the same objects are explicitly apparent because between times t1 and t2, the same tree, rectangle and ball can still be seen. TESLA1004, FIG. 20A. A POSITA would have understood that when thematic maps based on these two images are created and the location and type of information is provided in the maps, the two data structures would have included information on the same multiple objects.

Paper 38 (Pet. SurReply), 8.

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.

70

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

GTC's "Thematic Map" is A Data Structure

PO's argument that Tolstedt's Map is Not a "Data Structure" fails in light of the '974 Patent's Broad Description of a "Data Structure."

'974 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. SurReply), 9-10.

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. SurReply), 9-10.

71

FISH.

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Issue 9

GTC Renders Obvious the First and Second Data Structures being Correlated with the First and Second Instruction Sets (all amended claims)

“wherein the first data structure *is correlated with the first one or more instruction sets* for operating the first device; and;
wherein the second data structure *is correlated with the second one or more instruction sets* for operating the first device. ”

GHC 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 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 would have

Paper 31 (Opp. To MTA), 24.

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.