

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent of: JASMIN COSIC
U.S. Patent No.: 10,452,974 Attorney Docket No. 49649-0026IP1
Issue Date: October 22, 2019
Appl. Serial No.: 15/340,991
Filing Date: November 2, 2016
Title: ARTIFICIALLY INTELLIGENT SYSTEMS, DEVICES,
AND METHODS FOR LEARNING AND/OR USING A
DEVICE'S CIRCUMSTANCES FOR AUTONOMOUS
DEVICE OPERATION

SUPPLEMENTAL DECLARATION OF DR. JASON JANET

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	GROUND 1: GROTMOL ANTICIPATES CLAIMS 4, 5, AND 7	4
A.	Claim 4	4
B.	Claim 5	10
C.	Claim 7	15
1.	PO’s Overly Narrow Construction of “Connection” Should Be Rejected	16
2.	Grotmol Discloses the Contested Features of Claim 7, Including the Claimed “Connection[s]”	18
III.	GROUND 2: GROTMOL RENDERS OBVIOUS CLAIM 11	21
IV.	GROUND 3: THE GROTMOL-HICKMAN COMBINATION RENDERS OBVIOUS CLAIMS 4 AND 11	24
A.	PO mischaracterizes Hickman’s organization of knowledge bases	24
B.	PO mischaracterizes Hickman’s disclosure of circumstance representation.....	27
C.	PO fails to rebut the rationale for combining Grotmol and Hickman.....	30
D.	Claims 4 and 11 are rendered obvious by Grotmol and Hickman	30
V.	ASSESSMENT OF THE SUBSTITUTE CLAIMS	31
VI.	GROTMOL, HICKMAN, AND TOLSTEDT RENDER THE SUBSTITUTE CLAIMS OBVIOUS	32
A.	Motivation to Combine.....	34
1.	Grotmol and Hickman	34
2.	Grotmol and Tolstedt.....	37
B.	The Amended Features were Known in the Art and Obvious	40
1.	Substitute Feature 1 (Element [Sub-1]).....	40
2.	Substitute Feature 2 (Element [Sub-2]).....	41
(a)	Hickman	41

(b)	Tolstedt (alternative to Hickman)	42
3.	Substitute Feature 3 (Element [Sub-3]).....	43
(a)	Hickman	44
(b)	Tolstedt (alternative to Hickman)	45
4.	Substitute Feature 4 (Element [Sub-4]).....	45
(a)	Hickman	45
(b)	Tolstedt (alternative to Hickman)	46
5.	Substitute Feature 5 (Element [Sub-5]).....	47
(a)	Hickman	47
(b)	Tolstedt (alternative to Hickman)	47
VII.	CONCLUSION.....	47

I, Jason Janet, PhD, declare as follows:

I. INTRODUCTION

1. I have been retained as a technical expert by counsel on behalf of Tesla, Inc. (“**Tesla**”/“**Petitioner**”) in Case No. IPR2023-01054 with respect to U.S. Patent No. 10,452,974 (“**the ’974 patent**”) (TESLA1001).

2. I understand that this declaration (“**Supplemental Declaration**”) will be submitted in support of Petitioner’s Reply to Patent Owner’s Response (“**POR**”) to the Petition for *inter partes* review of the ’974 patent. This Supplemental Declaration supplements certain portions of my prior declaration in support of the Petition for this IPR Proceeding (TESLA1002, “**my Prior Declaration**”). In my Prior Declaration, I address many topics, including (but not limited to) my background and qualifications, the level of skill in art, an overview of the ’974 patent, a POSITA’s background knowledge, certain legal standards explained to me by counsel, legal principles applied for the purposes of my analysis and this declaration, and a detailed analysis of the prior art against the claim elements of the ’974 patent.

3. I maintain all of my earlier analysis and continue to hold the opinions expressed in my Prior Declaration. This Supplemental Declaration does not replace my Prior Declaration. These declarations are complementary and should be read in tandem. Accordingly, this Supplemental Declaration does not cover every topic

addressed in my Prior Declaration.

4. In forming this Supplemental Declaration, I have considered, among various documents, the following: the declaration of Patent Owner (PO)'s declarant, Dr. Saber (EX2007), the Board's institution decision (Paper 20), my knowledge and experience in this field, including my research, and work experience in the field, my experience working with others involved in the field, and the third-party documents cited throughout my analysis. *See* TESLA1002, ¶¶5-21 (discussing qualifications) and ¶¶22-42 (discussing legal standards and the level of skill for a "POSITA"). I have also considered the following additional references:

- Transcript of the Deposition of Dr. Saber in IPR2023-01054 (TESLA1011)
- Transcript of the Deposition of Dr. Saber in IPR2023-01055 (TESLA1012)
- U.S. Pub. No. 2014/0089232 ("Buibas" or TESLA1013)
- U.S. Patent Apl. Pub. No. 2010/0063663 ("Tolstedt" or TESLA1015)
- U.S. Patent No. 9,435,653 ("Zeng" or TESLA1017)

5. Although I cite and quote to selected portions of various references in this Declaration, the reader should understand that these citations are representative examples. A POSITA would have viewed each reference in its entirety and in combination with other references. Accordingly, I intend the references identified in this Declaration to be viewed as incorporated in their entireties.

6. I am not, and never have been, an employee of Tesla. I received no

compensation for this Supplemental Declaration beyond my normal hourly compensation based on my time actually spent analyzing the '974 patent, the prior art publications cited below, and various issues related thereto. I will not receive any added compensation based on the outcome of any IPR or other proceeding involving the '974 patent.

7. The analysis and opinions expressed in my Prior Declaration fully explain why [1] each and every feature of the '974 patent's claims is provided in the prior art; and [2] a POSITA would have been motivated to combine the known prior art elements in a manner that reaches the claimed invention. Patent Owner's expert, Dr. Saber, has considered my opinions and offered his own, some of which are inconsistent with my view. I will address certain points of dispute below. The fact that I may have not addressed all of Dr. Saber's opinions should not be interpreted as agreement with them.

II. GROUND 1: GROTMOL ANTICIPATES CLAIMS 4, 5, AND 7

A. Claim 4

8. Claim 4¹ (TESLA1001, 170:16-37) recites:

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 knowledge-base 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.

9. PO argues that “Grotmol fails to disclose that the knowledgebase includes the claimed fourth correlation learned in another learning process involving the second device” because “Grotmol only discloses *one* user operating *one* device

¹ I have added emphasis in the claims using yellow highlighted boxes.

in *one* learning process.” POR, 15-19. I disagree. PO’s arguments fail because they are inconsistent with Grotmol’s teachings and the understanding a POSITA would have of Grotmol’s disclosure. The error in PO’s arguments is also corroborated by PO’s declarant, Dr. Saber.

10. In my Prior Declaration, I explained that “Grotmol discloses a learning process that results in a memory storing multiple correlations (training sets), where each correlation includes a “circumstance representation” (image of scene from a robot vehicle’s operation) correlated with an “instruction set[] for operating a first device” (motor control indication configured to cause the robot to execute an action).” TESLA1002, ¶¶67-77, 91-94; TESLA1004, 3:27-39, Abstract, 8:66-9:3, 10:23-26, 38:36-58, 39:13-22. The cited disclosure from Grotmol included the following teachings:

Training of the robot’s controller may be based on the user observing robot’s actions and sending one or more target control commands to the robot via the 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, 33:21-35².

11. Grotmol thus expressly discloses that a given robot may perform learning and store a training configuration that enables the robot to perform autonomous operations. *Id.* This stored training configuration can then be shared with another robot “*in order to provide learned behaviors.*” *Id.*

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**”).

² I have indicated emphasis through bolded and/or italicized text unless noted otherwise.

13. Like Grotmol's training configurations that include first and second circumstance representations learned on a first device and stored in the knowledgebase on the first device (which PO does not dispute), Grotmol is clear that training configurations containing circumstance representations learned on another device and loaded on the first device are also stored in the knowledgebase. Indeed, Grotmol discloses one memory component 814 that stores the knowledgebase containing all the trained configurations and controllers for the learned behaviors on a device. TESLA1004, FIG. 8, 38:38-42, 3:35-40, 48:9-13. A POSITA would have understood that all training configurations—whether learned locally on the first device or shared by another device—are stored in the same knowledgebase (e.g., in memory 814) such that the training sets (trained configurations) are all retrievable to enable performance of different trained behaviors on the first device. TESLA1004, FIG. 8, 38:38-42, 3:35-40, 48:9-13.

14. Moreover, Grotmol explicitly discloses a system in which multiple types of robots are trained. TESLA1004, 18:33-37 (“In some *implementations of training autonomous robotic devices* (e.g., rovers, bi-pedaling robots, wheeled vehicles, aerial drones, robotic limbs, and/or other robotic devices), the update periodicity Δt may be configured to be between 1 ms and 1000 ms”), title, 34:49-55. And because each robot has a trainer/user, it would have been understood by a POSITA that Grotmol is directed to multiple types of robots that have a trainer/user.

TESLA1004, 16:4-8, 34:49-39:22, FIG. 1.

15. A POSITA would also have understood that different types of robots would have different learning processes. In particular, given Grotmol's disclosure of using various robots and sharing a robot's learning with other robots, it would have been understood that the training configuration generated and stored by one device is different than the training configuration generated and stored by another (second) device. For instance, if one robot is an aerial device and another robot is a terrain-based device, the two robots will move along different paths and take different pictures. In another example, if the two robots are both wheeled vehicle robots, the two robots would necessarily occupy different spaces and therefore have different experiences (e.g., move along paths at different times, take different pictures or detect different environments). Consequently, the "fourth circumstance representation" and corresponding correlation in a second robot would be different from a first robots' circumstance representation, and the learning processes of the two robots would have been understood by a POSITA to be different. Thus, when the training configuration is loaded onto the first device from the second device, a POSITA would have understood that such a training configuration includes "another learning process" that is different than a learning process of the first device.

16. Grotmol also discloses that a robot can receive training from different users. For instance, as explained with respect to FIGS. 9, 10A, and 10B (reproduced

below), a user can provide inputs to manage robot behavior, e.g., Behavior 1 through Adaptive Controller 1 and Behavior 2 through Adaptive Controller 2. TESLA1004, 25:10-26:48. As acknowledged by Dr. Saber during his deposition, it would have been understood by a POSITA that the user commands can be provided by multiple users such that a robot could be trained by one or more users. TESLA1011, 166:12-18 (“Question: And so these adaptive controllers can receive commands *from different users* and then have the outputs; is that right? [Answer] That's correct. I mean, *you receive commands from different users* to change a certain behavior or control a certain behavior or adjust it.”) *See also, e.g.*, TESLA1013, [0070]-[0071], [0181]-[0184], [0196].




FIG. 10B

Grotmol (TESLA1004), FIG. 10B.

17. Indeed, Grotmol explicitly discloses that its “BrainOS software platform may be configured to *enable users* to selectively remove a learned behavior.” TESLA1004, 34:18-21. Thus, PO’s assertion that Grotmol is limited to “*one* user operating *one* device in *one* learning process” is explicitly contradicted by Grotmol and PO’s own expert.

18. Dr. Saber’s testimony during his deposition highlighted the lack of support in his and PO’s arguments. For example, during his deposition, Dr. Saber was asked to identify disclosure in Grotmol that teaches that each device is operated by “*one* user” or a “single user” (as alleged by PO), but he was unable to do so. TESLA1011, 185:21-186:5; TESLA1012, 78:15-83:16.

19. In reviewing the POR, the PO seems to desire explicit disclosure of the claimed features in Grotmol, which, Petitioner’s counsel has advised me, is not necessary to sustain an anticipation ground. Here, for the reasons I have already explained, a POSITA would have understood that sharing of configuration information between robots would have been understood to be the result of “another learning process” and that the second robot provides a “fourth correlation including a fourth circumstance representation correlated with a fourth one or more instruction sets” to the first robot, as recited in claim 4.

B. Claim 5

20. Claim 5 (TESLA1001, 170:37-50) recites:

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.

21. According to PO, claim 5 requires “counting the [matched] portions and determining whether this counted number of portions exceeds a numerical threshold.” POR, 21. PO also argues that Grotmol’s “similarity measure is for measuring similarity of elements, not for counting their number or computing their percentage.” POR, 21. PO is wrong on both accounts.

22. To start, PO’s arguments rely on terms such as “counting” that are nowhere to be found in the language of claim 5. The plain language of claim 5 simply requires 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]—i.e., a determination that can be performed with or without counting. PO’s implicitly narrow construction is inconsistent with the claim language, is supported by neither

intrinsic nor extrinsic evidence, and should be rejected.

23. As I explained in my Prior Declaration, Grotmol discloses that “an acquired image may be compared with one or more images from the training buffer in order to determine the most likely match.” TESLA1004, Abstract, 39:23-64; TESLA1002, ¶¶47, 72. A similarity measure “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.” TESLA1004, 41:8-17, 3:50-52, 4:30-39, 47:25-27. Grotmol teaches that “[t]he similarity measure may characterize similarity between features of the subset and features of the given training set.” TESLA1004, 2:30-34. “[R]esponsive to the similarity measure breaching a threshold,” Grotmol’s system “select[s] one or more training sets from the plurality of training sets.” TESLA1004, 4:30-39, 47:4-6.

24. As can be appreciated from the foregoing, Grotmol teaches determining “a number of at least partially matching portions” by performing pixel-wise comparisons to determine a degree of similarity between a new image (e.g., a third circumstance representation) and an image in a training set (e.g., a first circumstance representation). TESLA1004, 41:8-17. The similarity measure is compared with a threshold, and when that threshold is exceeded by the similarity measure (“a number of at least partially matching portions ... exceeds a threshold number”), one or more training sets are selected. TESLA1004, 4:30-39, 47:4-6.

25. Even if PO's overly narrow claim interpretation applied, Grotmol would still disclose the disputed features in claim 5, including "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." Grotmol discloses that "shifted/trimmed" portions of a new image (corresponding to portions of a "third circumstance representation") are compared to "shifted/trimmed" portions of a subset of images from the training set (referred to in Grotmol as "the matching search set," which correspond to a "first circumstance representation") to determine a "best match" from the training set:

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 individual images from the matching search set. An image corresponding to the lowest norm may be referred to as the best match.*

TESLA1004, 41:8-17.

26. According to Grotmol, the “shifted/trimmed frames” for which a similarity metric is determined are “at least partially matching portions.” For example, Grotmol describes how an “image correlation operation” is performed between the new image and each candidate image in the matching search set to align the images before the pixel-wise difference is determined between the images. TESLA1004, 40:52-41:16. Through the image correlation operation and the resulting shifting/trimming of correlated frames, a POSITA would have understood that Grotmol only selectively performs the pixel-wise difference comparison between “portions” of a new image and images in a pre-determined “matching” search set, i.e., between “*at least partially matching portions* of the third circumstance representation and portions of the first circumstance representation,” as recited in claim 5. TESLA1004, 40:52-55; *see also id.*, 39:23-40:63, 40:67-41:7 (“Variability between matching images may exist.”).

27. Grotmol also teaches that determining a “best match” between a new image and an image in the “matching search set” involves determining that a number of shifted/trimmed portions of the new image (i.e., at least partially matching portions) exceeds a threshold. The threshold that must be exceeded is either (a) zero when a “best match” is determined with a portion of an image from a “primary particle” of the matching search set or (b) one or more depending on the window size when a “best match” is determined with a portion of an image from a “secondary

particle” set. TESLA1004, 31:43-42:2 (“[I]f 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.”). Grotmol thus plainly discloses “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],” as recited in claim 5.

C. Claim 7

28. PO seeks to distinguish the claimed features over Grotmol by advancing a narrow construction of the term “connection.” Specifically, PO contends that “a ‘connection’ should be construed as a ‘relationship that includes a link.’” POR, 10. I disagree.

29. Nothing in the specification, file history, or extrinsic evidence supports importing a requirement for a “connection” to include a “link” as proposed in the POR. And while it is clear that Grotmol discloses the claimed first and second connections under a plain meaning interpretation of “connection,” Grotmol still discloses the claimed features even under PO’s improper construction.

1. PO's Overly Narrow Construction of "Connection" Should Be Rejected

30. I have been informed by Petitioner's Counsel that the relevant inquiry in assessing whether to narrow claims and depart from plain meaning is whether the specification, with sufficient clarity and deliberateness, reflects lexicography or disavowal that warrants a narrowing of the plain and ordinary meaning of the disputed claim phrase to require the additional narrowing features. Here, the '344 Patent includes no lexicography or disavowal to support Patent Owner's construction of "connections" and the POR does not even attempt to demonstrate lexicography or disavowal.

31. 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. For example, PO turns to FIGS. 17 and 18A, which are described as "embodiments." TESLA1001, 56:17-23. Counsel for Petitioner has indicated to me that the word "embodiments" does not impose a limiting effect. 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 do not necessarily limit the scope of a claim term.

32. Notably, the textual description of FIGS. 17 and 18A does not once use the word “link,” which PO seeks to import into the construction of “connection.” The specification also broadly describes “connections”—both “direct” and otherwise—at col. 66:3-23 without ever referring to or mandating that a connection include a “link” as proposed by PO. TESLA1001, 66:3-23.

33. 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. For instance, the ’974 patent does not offer a definition of the word “link.” Compounding this lack of clarity was the refusal of PO’s declarant to explain what the word “link” meant. For instance, based on Dr. Saber’s deposition transcript, it appears that he was asked for almost 40 minutes what the term “link” means during his deposition, but Dr. Saber did not define or clarify the term. TESLA1012, 12:12-36:5. With no further clarity offered by the declarant, a reader is left to guess what the scope of the term “connections” is because it’s unclear what the word “link” encompasses.

34. PO cites two dictionary definitions (EX2012, EX2013) as allegedly supporting its construction, each of which provide numerous definitions for

“connection.” POR, 14. PO chose a definition that includes the word “link” but did not include any explanation for why PO chose this definition 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 contexts, there is no basis in the specification, file history, or otherwise for limiting a connection to a “relationship that includes a link” as PO contends.

35. Indeed, based on the broad description in the ’974 patent (*see e.g.*, TESLA1001, 15:58-60, 66:3-23) of the term “connection,” no construction of “connection” is necessary because the plain and ordinary meaning is used in the ’974 patent and should be applied for understanding the claims as well. As shown in the dictionaries provided by the PO, the plain and ordinary meaning of “connection” encompasses “an association or a relationship,” which is also consistent with the description of connection, explicitly provided by the ’974 patent. EX2012, 399-400, EX2013, 481; TESLA1001, 15:58-60, 66:3-23. Thus, when two nodes or points of data have a connection, like in FIGS. 17 and 18A of the ’974 patent, an “association or relationship” between the nodes or points of data provides a “connection.”

**2. Grotmol Discloses the Contested Features of Claim 7,
Including the Claimed “Connection[s]”**

36. Claim 7 (TESLA1001, 170:59-67) recites:

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.

37. As a preliminary matter, because PO's arguments on claim 7 are dependent on PO's faulty construction of "connection," PO's arguments with respect to claim 7 fail.

38. Yet, even if PO's construction is applied, Grotmol still discloses the claimed connections. As explained in my Prior Declaration, Grotmol describes storing correlations including sequences of images with their corresponding instruction sets (i.e., a sequence of knowledge cells), and that the training buffer may be ordered temporally, according to when the images were acquired. TESLA1002, ¶¶98-101, TESLA1004, Abstract, 38:21-42, 39:17-40:19. Because storing data in a sequential manner in a memory creates an association or relationship between the stored data, e.g., based at least on the sequential storage locations, Grotmol discloses the claimed "first connection" and "second connection."

39. The connections generated by storing sequences of images with their corresponding instruction sets in a temporally ordered training buffer allow

Grotmol's control process to efficiently identify an image stored in the training buffer as a "best match" to a new image captured during autonomous operation:

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.

In order to reduce computational requirements of the image match process, the given image may be compared to a subset of images from the training buffer using image match process described in detail below.

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.

TESLA1004, 39:23-48.

40. Thus, the sequence of images stored in Grotmol's ordered training buffer serve to connect images that are temporally proximate to a limited number of "particle" images. Grotmol's control process uses the temporal connections to narrow the search space and more efficiently identify a matching image from the training set. 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). There is no record basis for limiting how a link must be manifested under PO's construction.

III. GROUND 2: GROTMOL RENDERS OBVIOUS CLAIM 11

41. Claim 11 (TESLA1001,) recites:

11. 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 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 operating the second device at least partially by another user.

42. As a preliminary matter, PO's arguments regarding claim 11 are similar to PO's arguments for claim 4, and thus fail for similar reasons as noted above with

respect to claim 4. POR, 26-34.

43. PO contends that claim 11 adds “at least partially by another user” to the features recited in claim 4. POR, 26-27. But even the “another user” limitation is rendered obvious by Grotmol.

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.

45. For example, it would have been obvious that a trainer for an aerial drone would be different from a trainer for a submarine device or a robot having robotic limbs. TESLA1004, 18:33-37. In another example involving similar robots, one trainer using a joystick to control a car (first robot) could train the car to turn in a particular direction (e.g., turn left or right), whereas another trainer could train a second car (second robot) to take other action and maneuvers (e.g., move in a circle).

TESLA1004, 35:5-25. In this example, two users train the two robots in different ways. Again, the notion of allowing multiple users for training purposes is consistent with Grotmol's disclosure. TESLA1004, 34:18-21. Accordingly, for at least these reasons, Grotmol renders claim 11 obvious.

46. PO also contends that I did not perform a proper obviousness analysis. POR, 33-35. These attacks are based on misrepresentations of my testimony. For example, in my Prior Declaration, I had presented my opinion on the obviousness of claim 11 in view of Grotmol. TESLA1002, ¶[140]. During my deposition, PO's counsel asked me why claim 11 was not anticipated. EX2011, 61:7-62:15. I responded by indicating that at the time of drafting my Prior Declaration, I felt that Grotmol should "be supplemented by what a person of ordinary skill would know," e.g., the role a different user would play. EX2011, 61:7-62:15. "[T]o me it's [(claim 11)] comfortably in this second ground [(obviousness)], but I think plausibly in the first [(anticipation)]." EX2011, 63:1-7. Simply because I considered an anticipation ground as "plausible," does not mean that my analysis provided in my Prior Declaration is incorrect.

47. I also explained to PO's counsel that "even if it [(Grotmol)] has to be elaborated on by what a person of ordinary skill in the art would know, yes, I believe that Grotmol discloses the elements of limitations of claim 11." EX2011, 63:18-23. Thus, I was attempting to explain to PO's counsel that I believed Grotmol was a

strong reference and that in combination with the knowledge that a POSITA has, would render claim 11 obvious by disclosing all its claimed features. *See e.g.*, EX2011, 69:13-20. At no point did I change my testimony regarding the obviousness challenge to claim 11 or the mapping of Grotmol applied to claim 11. EX2011, 72:22-73:2 (“I’m relying on obviousness in Ground 2, yes.”)

IV. GROUND 3: THE GROTMOL-HICKMAN COMBINATION RENDERS OBVIOUS CLAIMS 4 AND 11

A. PO mischaracterizes Hickman’s organization of knowledge bases

48. PO argues that Hickman is incompatible with Grotmol because Hickman “stor[es] 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 POS[IT]A would have found to be incompatible with Grotmol’s disclosure of training sets or correlations with images and instruction sets correlated and stored *together*.” POR, 35. I disagree.

49. As a preliminary matter, the organization of knowledgebases is not relevant to the combination of Grotmol and Hickman, which incorporates Hickman’s disclosure of sharing knowledge between different robots through a cloud computing system, and not the organization of various knowledgebases in Hickman. TESLA1002, ¶¶142-149.

50. To the extent PO's argument is considered, PO incorrectly describes Hickman. As shown below, Hickman's FIG. 3 discloses a datastructure 314 for a general object knowledge base 308 (both highlighted in yellow). TESLA1005, 15:28-37. The datastructure 314 includes, among other data, object identifier (GID), image data (IMG), and task data (TASKS). TESLA1005, 15:28-37.




FIGURE 3

Hickman (TESLA1005), FIG. 3.

51. The tasks data (TASKS) “may correspond to tasks that can be performed with the object corresponding to the general object identifier (GID);” “defines the types of interactions that a robot may execute with the physical object corresponding to the general object identifier (GID);” and may “include instruction code for controlling and operating” a robot.” TESLA1005, 18:34-19:34. Hickman explains that “[i]n 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*,” but, in other embodiments, the task data “can be looked up in a separate task knowledge base 309.” Thus, Hickman clearly discloses different implementations in which the task data can optionally be included or excluded in the same knowledge base as the images. PO’s argument that Hickman is somehow limited to “storing instruction sets and images separately in separate knowledgebases” contradicts Hickman’s disclosure.

52. To the extent PO argues that only the unique object knowledge base 307 should be considered in the Grotmol-Hickman combination (which I did not previously present and the Grotmol-Hickman combination is not limited to), Hickman teaches that “[v]arious 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 15 than the categories

shown in conceptual data structure 313.” TESLA1005, 12: 5-18. Thus, Hickman’s disclosure, even if limited to the use of the unique object knowledge base 307, would not have taught against the combination of Grotmol and Hickman because Hickman contemplates the addition of other fields to the unique object knowledge base 307 and does not exclude the addition of task data.

B. PO mischaracterizes Hickman’s disclosure of circumstance representation

53. PO also argues that “Hickman must disclose the circumstance representations and correlations in order to be combinable with Grotmol” and that Hickman teaches away from circumstance representations “by disclosing single object information (the information about the single object being manipulated in Hickman), which is not a circumstance representation.” POR, 35-37.


54. Again, PO’s argument is irrelevant because I did not rely on Hickman’s disclosure of circumstance representation.

55. To the extent PO’s argument is considered, PO’s assertion that Hickman is focused on manipulating a single object is unsupported and unclear³. Hickman, for instance, discloses that “the shared robot knowledge base includes data about a *vast number of various objects that robots may encounter and/or interact*

³ Neither PO nor Dr. Saber explain what “single object” information and manipulation entails.

with in their everyday environments. The shared robot knowledge base *may also include robot instructions for interacting with the objects.*” TESLA1005, 2:48-64, 6:53-57, 4:66-5:2, 8:36-44. Thus, Hickman clearly discloses robot interactions with multiple objects, not just one.

56. Simply because Hickman organizes its knowledge base by listing objects individually does not mean that Hickman is somehow limited to only interacting with a single object. In fact, Hickman’s organization of identified objects is similar to the organization of objects in the ’974 patent, which lists individual objects in the Collection of Object Representations 525 (horizontally as opposed to vertically), as shown below in FIG. 4B of the ’974 patent. TESLA1001, 71:38-73:12.



’974 Patent (TESLA1001), FIG. 4B.




FIGURE 3

Hickman (TESLA1005), FIG. 3.

57. PO makes blanket assertions (without evidentiary support) that Hickman does not teach circumstance representations, but fails to explain how Hickman's various knowledge bases that include object identification information, images, tasks, and more, as explained above and shown in FIG. 3, do not disclose circumstance representation (notwithstanding Grotmol's disclosure of circumstance

representation). POR, 35-37.

C. PO fails to rebut the rationale for combining Grotmol and Hickman

58. In my Prior Declaration, I provided several reasons why a POSITA would have combined Grotmol and Hickman, and point to explanations in Hickman and a 2014 Science Article that explain the benefits of the Grotmol-Hickman combination. TESLA1002, ¶¶142-149. For instance, I explained that shared robot learning would “allow far-flung robots to share those lessons and skills” and distribute computer processing across a cloud computing system, among various other benefits. TESLA1002, ¶¶142-149. *These benefits are entirely unrebutted by PO.* Indeed, the combination of Grotmol and Hickman is the predictable result of combining prior art elements according to the known methods disclosed in Grotmol and Hickman.

D. Claims 4 and 11 are rendered obvious by Grotmol and Hickman

59. PO’s arguments challenging the obviousness of claims 4 and 11 rise and fall with the arguments that have already been addressed above. For instance, PO incorrectly argues that Grotmol is limited to “one user operating one device in one learning process” and that Hickman “discloses only a single object information (not a circumstance representation) and separate (i.e., uncorrelated) instructions that do not disclose a correlation.” POR, 41-56. These two and other arguments that PO

relies on to argue against the obviousness of claims 4 and 11 have been rebutted above.

V. ASSESSMENT OF THE SUBSTITUTE CLAIMS

60. PO has also submitted a Motion to Amend (MTA) in which substitute claims 21-26 were submitted. Substitute claims 21-24 were submitted to contingently amend claims 4, 5, 7, and 11, respectively, and claims 14 and 15 were canceled and substituted with claims 25 and 26, respectively, without any contingency.

61. While I am not a lawyer, my review of the substitute claims revealed that the substitute claims do not recite any innovative subject matter and are generally directed to computerizing a mental process. Specifically, the substitute claims recite features related to determining a manner of operating a device in a current circumstance based on how that device, or another device, operated in similar circumstance(s) in the past.

62. This concept is implemented by human beings every day and has been for a long time. No feature recited in the substitute claims captures an operation or process differently than what a human mind would do in learning from their past experiences and impressions (whether through images or otherwise).

63. For example, 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. Likewise, a driver performs the abstract idea whenever he or she slows for a pedestrian who appears poised to jaywalk based on experience with other pedestrians jaywalking under similar conditions in the past. The fact that people not only can, but commonly do, base their present actions on past experiences is indisputable.

64. PO's substitute claims simply take the above-noted basic human intuition and apply it to a computer. But even the aspects of a computer system that are claimed, such as a memory and processor circuits, are well-known, generic computer components. Nothing in the claims strikes me as being part of an innovative or new computer, system, or method. The substitute claims, at best, are just a mental process implemented on a generic computer.

VI. GROTMOL, HICKMAN, AND TOLSTEDT RENDER THE SUBSTITUTE CLAIMS OBVIOUS

65. The substitute claims do not recite features that are unreachable by the prior art. In fact, the earlier prior art of record discloses most of the features. One additional reference, Tolstedt, can be added to the existing grounds to render all the substitute claims obvious. A table including the earlier grounds and grounds challenging the substitute claims is provided below, followed by my assessment of how the prior art references would apply to the substitute claims.

Original Claim	Original Invalidity Ground(s)	Substitute Claim	Substitute Invalidity Grounds
4	102: Grotmol 103: Grotmol and Hickman	21	103: Grotmol and Hickman 103: Grotmol and Tolstedt
5	102: Grotmol	22	103: Grotmol and Hickman 103: Grotmol and Tolstedt
7	102: Grotmol	23	103: Grotmol and Hickman 103: Grotmol and Tolstedt
11	103: Grotmol in view of POSITA knowledge 103: Grotmol and Hickman	24	103: Grotmol, POSITA knowledge, and Hickman 103: Grotmol, POSITA knowledge, and Tolstedt 103: Grotmol and Hickman
14	102: Grotmol	25	103: Grotmol and Hickman 103: Grotmol and Tolstedt
15	102: Grotmol	26	103: Grotmol and Hickman 103: Grotmol and Tolstedt

A. Motivation to Combine

1. Grotmol and Hickman

66. I previously explained why a POSITA would have been motivated and found it obvious to supplement the teachings of Grotmol with a POSITA's knowledge (Ground 2) and to combine Grotmol with Hickman (Ground 3) in a manner that would have rendered the issued '974 patent claims obvious. TESLA1002, ¶¶139-157. For example, I explained that a POSITA would have applied, to Grotmol, Hickman's suggested use of a cloud computing system with client devices (e.g., robots) to facilitate the sharing of knowledge between the various client devices. TESLA1002, ¶¶142-149.

67. Like Grotmol, Hickman discloses the storage of training data, including images and instructions, in a memory. TESLA1004, FIG. 8, 33:13-35, 38:33-41; TESLA1005, FIG. 3, 10:37-24:64. Hickman's shared robot knowledge base 306 includes several knowledge bases 307-311 that include a unique object knowledge base 207 and a general object knowledge base 308. TESLA1005, FIG. 3. The general object knowledge base 308 includes one or more data structures 314 with data fields, including, but not limited to, "an individual general object identifier (GID) [that] may have the following corresponding data: (i) language data (LNG); (ii) image data (IMG); (iii) manufacturer data (MFG) (if applicable); (iv) model number data (MOD)" and "may contain more or less information." TESLA1005,

15:28-20:20, FIG. 3. The unique object knowledge base 207 includes one or more data structures 313 with another set of data fields. TESLA1005, 11:47-15:27, FIG. 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.

68. In addition to the reasons I noted in my Prior Declaration, 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. 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.

69. Moreover, a POSITA would have had a reasonable expectation of success implementing the Grotmol-Hickman combination ("GHC"). For instance,

the choice to store data in one knowledge base or multiple knowledge bases, was well-known, as disclosed by Hickman, whose filing date predates Grotmol's filing date by at least two years. TESLA1005, 10:57-11:7. Grotmol and Hickman belong to the same field of endeavor related to enhancing the autonomous operation of robotic devices, and both describe the training of robots in an analogous manner making their combination predictable. In addition, the use of various types of data structures to store data would have been entirely predictable particularly when extending Grotmol's system to a cloud system with various types of robots and devices connected to the cloud system. Applying Hickman's data structures in Grotmol's system would not have required extraordinary skill and would have been within the ordinary capabilities of a POSITA.

70. The GHC storage system would better allow for robot queries in the cloud system to be received and processed, and for relevant data to be retrieved from the shared knowledge base. TESLA1005, 10:41-56, 11:9-30. For example, if information obtained from an image includes a very specific type of car, then the information stored in data structure 313 of the unique object knowledge base 307 could be used for comparison purposes (and subsequent retrieval of instructions) because such a knowledge base includes information that can identify unique objects, such as a specific car model, as taught by Hickman. However, if the information obtained from an image does not provide information that could identify

a specific type of car, then the information stored in data structure 314 of the general object knowledge base 308 could be used for comparison purposes (and subsequent retrieval of instructions) because such a knowledge base includes information that can identify general objects, e.g., a car. In some implementations of the combination, multiple data fields could be combined if two or more knowledge bases are combined and a POSITA would have understood that the combination would have provided greater information for any one object than if individual knowledge bases were used. As can be appreciated from the foregoing, using these different knowledge bases, data structures, and combinations thereof would have improved knowledge sharing and object matching capabilities, thereby improving performance of Grotmol's system.

71. The motivation to combine Hickman with Grotmol in view of POSITA knowledge would have involved the same reasons noted above.

2. Grotmol and Tolstedt

72. Tolstedt, like Grotmol, is directed to robotic vehicle navigation, and more specifically to methods “for controlling the autonomous operation of a vehicle.” TESLA1015, ¶[2]; TESLA1004, Abstract, 16:15-23. Tolstedt discloses classifying data received from sensors, such as cameras, that obtain images. TESLA1015, ¶¶[68], [30]-[31], [70], [73], [82], [87], FIG. 4. Tolstedt also teaches that:

This classification may include identifying objects that have been detected in the environment” such 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.

73. TESLA1015, ¶[83]. This classification information is subsequently used when Tolstedt’s system receives environmental data (e.g., images) and determines how to instruct a vehicle to move based on matches with the classification data stored in a knowledge base 800. TESLA1015, ¶¶[84]-[87], FIG. 8.

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.

75. Moreover, the Grotmol and Tolstedt combination would have been predictable and a POSITA would have had a reasonable expectation of success implementing the combination. Grotmol and Tolstedt both describe techniques directed to improving autonomous operation of a vehicle, and both use object imagery to train the system and facilitate vehicle navigation. A POSITA looking to improve Grotmol's system would have turned to Tolstedt to realize the above-noted benefits and would have had a reasonable expectation of success, especially since incorporating Tolstedt's teachings in Grotmol's system would have been well within the capabilities of a POSITA and would have merely involved additional conventional programming and routine system modifications to realize Tolstedt's classification algorithm and map generation feature.

76. The motivation to combine Tolstedt with Grotmol in view of POSITA knowledge would have involved the same reasons noted above.

B. The Amended Features were Known in the Art and Obvious

1. Substitute Feature 1 (Element [Sub-1])

77. Substitute feature 1 added to each of the substitute claims is:

“wherein the first circumstance representation includes a first collection of object representations, and the second circumstance representation includes a second collection of object representations.”

78. As I noted in my Prior Declaration, “circumstance representation” is not defined in the ’974 patent specification, but would have been understood to include representations of objects. TESLA1002, ¶69. In Grotmol, the “*circumstance representation*” includes various types of data including, but not limited to, a collection of object representations, such as images of objects in a training set detected by one or more robot sensors, and data structures related to the objects and images. TESLA1004, FIGS. 20A, 20B; 38:43-39:12. For example, as illustrated in Grotmol’s FIGS. 20A, 20B, image 2000 is an example part of a “*first circumstance representation*” and “may comprise representations one or more objects, e.g., a tree 2002, a rectangle, a ball 2008, and/or other objects.” TESLA1004, FIG. 20A, 38:43-39:12. A subsequent image 2010 is an example part of a “*second circumstance representation*,” And may comprise representations of objects from image 2010. TESLA1004, FIG. 20B, 38:43-39:12. Thus, images 2000 and 2010 respectively include first and second collections of object representations.

Thus, Grotmol discloses substitute feature 1.

2. Substitute Feature 2 (Element [Sub-2])

79. Substitute feature 2 added to each of the substitute claims is:

“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.”

80. Grotmol’s FIGS. 20A and 20B illustrate how images 2000 and 2010 may be shifted “based on relative motion between the images (e.g., motion speed and time interval t_2-t_1 between the images).” TESLA1004, FIGS. 20A, 20B, 38:43-39:12. A POSITA would have understood that time t_1 corresponds to the time at which the first image (e.g., image 2000) was taken. Accordingly, Grotmol teaches at least *“wherein the first collection of object representations represents multiple objects detected at a first time or during a first time period.”*

(a) Hickman

81. Hickman discloses that data associated with a particular object that is stored in a knowledge base can include one or more of the object’s physical dimensions, weight, material composition, object identifier, serial number, image data, and location of the object (LOC). TESLA1005, 2:65-3:7, 12:5-18; FIG. 3. “The location data (LOC) may include GPS coordinates and/or other location-based

information, e.g., an association with (or a link to) one or more maps stored in the map knowledge base 310.” TESLA1005, 14:3-29. In some cases, when a general object knowledge base is being used, the stored data associated with an object may include “data about different ‘types’ of objects that a robot may encounter or interact with in its environment, including different ‘types’ of ‘unique’ and ‘non-unique’ objects.” TESLA1005, 14:61-19:22. Such object data is stored in data structures, e.g., data structures 313, 314, or 315, which along with object image data corresponds to the “*first collection of object representations*” in “*the first circumstance representation*.” Accordingly, GHC renders obvious “*a first data structure that includes: labels indicating types of the multiple objects, and coordinates indicating locations of the multiple objects.*”

(b) Tolstedt (alternative to Hickman)

82. Tolstedt also discloses that sensor data (one or more images) is processed to classify objects detected in the data and identify attributes of the detected objects such as the type, dimensional information, location, height, width, and color of the object, and actions for the vehicle to take in association with the object (e.g., performing object avoidance). TESLA1015, ¶¶83-88. This information is stored in a data structure in a knowledge base. *Id.*

83. 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, which may contain a spatial pattern of attributes”), 73, 103, 123, 134; TESLA1001, 65:45-66:2. This map can be stored in the knowledge base. TESLA1015, ¶¶84-87.

84. Tolstedt also teaches that the location information is obtained using the Global Positioning System. TESLA1015, ¶¶101, 120, 130, 31, 68-69, 73. Thus, it would have been obvious to a POSITA that the location information of the objects would have been “*coordinates indicating locations of the multiple objects.*” TESLA1017, 1:8-2:6 (indicating that GPS systems use coordinates). Accordingly, the Grotmol-Tolstedt system renders obvious “*a first data structure that includes: labels indicating types of the multiple objects, and coordinates indicating locations of the multiple objects.*” This first data structure along with images of detected objects are included in the “*first collection of object representations*” of “*the first circumstance representation.*”

3. Substitute Feature 3 (Element [Sub-3])

85. Substitute feature 3 added to each of the substitute claims is:

“*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.*”

86. In the example of FIGS. 20A and 20B, Grotmol teaches that an amount of shift between images 2000 and 2010 may be shifted “based on relative motion between the images (e.g., motion speed and **time interval t2-t1 between the images**).” TESLA1004, FIGS. 20A, 20B, 38:43-39:12. A POSITA would have understood that time t2 corresponds to the time at which the second image (e.g., image 2010) was taken. Accordingly, Grotmol teaches at least “*wherein the second collection of object representations represents multiple objects detected at a second time or during a second time period.*”

(a) Hickman

87. For the reasons explained above with respect to feature 2, Hickman discloses data structures (“second data structure”) that store “labels indicating types of the multiple objects, and coordinates indicating locations of the multiple objects.” 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.

(b) Tolstedt (alternative to Hickman)

88. Tolstedt also renders substitute feature 3 obvious for the same or similar reasons noted above in Section III.B.2(b). It would have been obvious to a POSITA that Tolstedt discloses a “*second data structure*.” This “*second data structure*” would have included images at a time (and/or location) and a set of objects from the objects that may be different than the images (and/or location) associated with the first data structure in Section III.B.2(b). *Id.* In some cases, the second data structure can be a second map of a second location different from the first location and on which multiple objects and their coordinates are added. TESLA1015, ¶¶123, 134, 83-88. In other cases, the knowledge base includes the same type of information as added to the second map. 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. For these reasons, Grotmol-Tolstedt renders obvious substitute feature 3.

4. Substitute Feature 4 (Element [Sub-4])

89. Substitute feature 4 added to each of the substitute claims is:

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

(a) Hickman

90. 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 POSITA would have understood that because the tasks for the robot (“*the first one or more instruction sets for operating the first device*”) are part of the first data structure and provide instructions for the robot, “*the first data structure is correlated with the first one or more instruction sets.*”

(b) Tolstedt (alternative to Hickman)

91. 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 been obvious to a POSITA that a portion of the knowledge base storing information on an object or a map showing the object (“*first data structure*”) is correlated with instructions/actions for operating the vehicle/robot such as object avoidance (“*correlated with the first one or more instruction sets for operating the first device*”).

5. Substitute Feature 5 (Element [Sub-5])

92. Substitute feature 5 added to each of the substitute claims is:

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

(a) Hickman

93. For reasons similar to those noted above with respect to substitute feature 4, it would have been obvious to a POSITA that in GHC, tasks (“***second one or more instruction sets for operating the first device***”) for operating the robot when objects in image 2010 are detected, are stored in a second data structure. Thus, GHC renders obvious substitute feature 5.

(b) Tolstedt (alternative to Hickman)

94. For reasons similar to those noted above with respect to substitute feature 4, it would have been obvious to a POSITA that the Grotmol-Tolstedt system renders obvious substitute feature 5. *Id.* Here, it would have been obvious to a POSITA that the second data structure could include object data and correlated instructions for objects detected at a different time and/or location than objects for which information is stored in the first data structure. *Id.*

VII. CONCLUSION

95. After reviewing the POR and Patent Owner’s expert declaration, I was not persuaded by any of the arguments therein. As such, I still conclude that the

Remaining Challenged Claims (claims 4, 5, 7, and 11) are obvious.

96. Moreover, while the proposed substitute claims appear to be directed to well-known mental processes implemented on a generic computer, the substitute claims are also obvious over prior art, as I noted above.

97. I declare that all statements made in this Declaration of my own knowledge are true and that all statements made on information and belief are believed to be true. I further declare that these statements are made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both (under Section 1001 of Title 18 of the United States Code).

Executed this 30th day of May, 2024.



A handwritten signature in black ink, consisting of a large loop followed by a series of smaller loops and a final upward stroke, positioned above a horizontal line.

Jason Janet