

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent of: Jasmin Cosic
U.S. Patent No.: 11,055,583
Issue Date: July 6, 2021
Appl. Serial No.: 16/584,736
Filing Date: September 26, 2019
Title: MACHINE LEARNING FOR COMPUTING ENABLED
SYSTEMS AND/OR DEVICES

SUPPLEMENTAL DECLARATION OF DR. JASON JANET

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	GROUND 1: GROTMOL ANTICIPATES CLAIMS 1-9, 11-24, 26, 28-29, AND 31-34	4
A.	Claims 1-6, 8, 11-13, 19-22, 24, 26, 28, and 32-34	4
B.	Claim 7	4
C.	Claim 9	9
D.	Claim 14	13
E.	Claim 15	16
F.	Claim 16	22
G.	Claim 17	26
H.	Claim 18	29
I.	Claim 23	34
	1. PO's Construction of "Connection" is Unsupported by the '583 Patent and File History.....	34
	2. Grotmol Discloses Claim 23 and the Recited "Connections" ..	38
J.	Claim 29	39
K.	Claim 31	41
III.	GROUND 2: GROTMOL RENDERS OBVIOUS CLAIMS 10 AND 11	42
A.	Claim 10	42
	1. Grotmol's Multiple Learning Processes	43
	2. Grotmol's Multiple Users for Training the Robot in Different Learning Processes	47

B.	Claim 11	48
IV.	GROUND 3: THE GROTMOL-HICKMAN COMBINATION RENDERS OBVIOUS CLAIMS 25 AND 27	55
A.	Combining Grotmol and Hickman Would Have Been Obvious to a POSITA	55
B.	Claim 25	59
C.	Claim 27	60
V.	GROUND 4: THE GROTMOL-HOFFMAN COMBINATION RENDERS OBVIOUS CLAIM 30	63
A.	Combining Grotmol and Hoffman Would Have Been Obvious to a POSITA	63
B.	Claim 30	64
VI.	CONCLUSION	66

I, Jason Janet, Ph.D., 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 Nos. IPR2023-01172 and IPR2023-01173 (“**the IPRs**”) with respect to U.S. Patent No. 11,055,583 (“**the ’583 patent**”) (EX1001).

2. I understand that this declaration (“**Supplemental Declaration**”) will be submitted in support of the respective Petitioner’s Replies to Patent Owner’s Responses (“**PORs**”) in each of the IPRs on the ’583 patent. This Supplemental Declaration supplements certain portions of my prior declaration in support of the Petitions for each of the IPRs (EX1003, “**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 ’583 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 ’583 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 decisions (Paper 21 in each IPR), 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 EX100X, ¶¶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 documents:

- EX1001 U.S. Patent No. 11,055,583 (the "'583 patent")
- EX1002 Prosecution File History of U.S. Patent No. 11,055,583
- EX1005 U.S. Patent No. 9,604,359 ("Grotmol")
- EX1007 U.S. Patent No. 8,639,644 ("Hickman")
- EX1008 Service, Robert F., *Minds of Their Own*, Sci. Magazine. Oct. 10, 2014, 346(6206):182-183, available at, https://www.usna.edu/AcResearch/_files/documents/Science_Minds_of_Their_own_182.full1.pdf
- EX1009 U.S. Patent No. 9,283,674 ("Hoffman")
- EX1012 Patent Owner's Response, IPR2023-01055 (U.S. Patent No. 11,238,244), Paper 26 (March 7, 2024)
- EX1015 U.S. Patent No. 10,607,134 ("Cosic '134")
- EX1016 Transcript of Deposition of Eli Saber, Ph.D., IPR2023-01055 (U.S. Patent No. 11,238,344) (May 21, 2024)
- EX1018 Selected Excerpts from the New Oxford American Dictionary (3d Ed.) (2010)

- EX1019 Selected Excerpts from Webster's New World College
Dictionary (4th Ed.) (2012)
- EX1020 U.S. Pub. No. 2015/0258683 ("Izhikevich-683")
- EX1021 U.S. Pub. No. 2014/0277718 ("Izhikevich-718")
- EX1022 U.S. Pub. No. 2014/0089232 ("Buibas")

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 '583 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 '583 patent.

7. The analysis and opinions expressed in my Prior Declaration fully explain why [1] each and every feature of the '583 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 and its

expert, Dr. Saber, have considered my opinions and offered their own, some of which are inconsistent with my view. I address certain points of dispute below. The fact that I may have not addressed all of PO's arguments or Dr. Saber's opinions should not be interpreted as agreement with them.

II. GROUND 1: GROTMOL ANTICIPATES CLAIMS 1-9, 11-24, 26, 28-29, AND 31-34

A. Claims 1-6, 8, 11-13, 19-22, 24, 26, 28, and 32-34

8. In my Prior Declaration (EX1003), I explained in detail how each of claims 1-9, 11-24, 26, 28-29, and 31-34 of the '583 patent is anticipated by Grotmol. My opinion that Grotmol anticipates each of these claims remains unchanged. Below, I address certain errors in the arguments that PO and Dr. Saber raised against the challenges to claims 7, 9, 14-18, 23, 29, and 31. I understand that PO did not advance patentability positions or otherwise respond to the Petition's anticipation challenges for any of claims 1-6, 8, 11-13, 19-22, 24, 26, 28, or 34. Grotmol nonetheless anticipates each of these claims for the reasons addressed in my Prior Declaration.

B. Claim 7

9. Claim 7 of the '583 patent recites the following language:

7. The system of claim 4, wherein the determining the first one or more instruction sets for operating the first device based on the at least partial match between the new one or more digital pictures and the first one or more digital pictures includes determining at least partial match between one or more objects detected in the new one or more digital pictures and one or more objects detected in the first one or more digital pictures.

10. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that claim 7 is anticipated by Grotmol. See EX1003, ¶¶238-239.

11. PO argues in the POR that “Grotmol merely discloses comparing images or portions (i.e., features) thereof and does not disclose detecting objects and then comparing the detected objects.” POR, 12. I do not agree. PO’s argument appears to be rooted in a flawed characterization of Grotmol’s “image feature[s]” as merely “a portion of an image (i.e., each image feature is a collection of pixels of the image)” and “not a detected object.” POR, 13. Grotmol is in fact clear that its “image features” include indications of detected objects. When Grotmol compares features of images from a newly acquired image to features of a stored image to determine at least a partial match, the comparison includes a comparison of detected objects when the images’ features themselves include indications of detected objects. Instruction sets such as motor commands can then be determined by identifying the particular instruction sets that are correlated to a matching image.

12. To illustrate in further detail, Grotmol discloses “feature extractors” that detect and extract features from a sensory input such as an image of a robot’s surroundings:




FIG. 10A

EX1005, FIG. 10A (annotated); *see also id.*, 10:54-60 (“receive sensory context (e.g., **camera input** and/or input from other sensors, other context signals based on internal states or system history (which may be processed to **extract interesting features**), and/or other information associated with sensory context”), 39:13-17 (“During training, images (e.g., raw and/or pre-processed) may be stored in a memory buffer (training buffer). In one or more implementations, preprocessing operations may comprise resampling, cropping, light balancing, and/or **feature extraction**.”).¹

13. For autonomous operation, Grotmol’s robot can determine motor

¹ All emphasis added unless otherwise noted.

control commands by comparing features from a new sensory input (e.g., a new digital picture) to features of sensory inputs (e.g., earlier-captured pictures) that were previously acquired during training and stored in a training buffer. EX1005, 4:22-42 (“comparing individual features”), 39:13-28, 46:60-65. The robot determines at least a partial match between features from the new sensory input and features of one or more of the sensory inputs stored in the training buffer. *Id.*, 39:14-40:51 (“match process”), FIG. 15. The robot uses the learned motor control commands that are correlated with the matching sensory inputs to determine the operations it should perform autonomously. *Id.*

14. Grotmol expressly discloses that the “features” of a sensory input such as a digital picture can include “one or more integer or floating point values characterizing the input, e.g., **the presence or absence of an ... object**, at particular locations in the image.” EX1005, 28:19-24. Grotmol’s feature set can further characterize “properties,” “the position ... or the relative position,” and “changes ... across consecutive frames” of objects and other features detected in the image. *Id.*, 28:14-40. A POSITA would have recognized from these disclosures that Grotmol’s feature set thus includes indications of detected objects such as values characterizing the presence or absence of an object at particular locations in the image, and/or values characterizing properties, position, and changes in objects across consecutive frames. As Grotmol uses the image features for comparison and the features include

indications of detected objects, it is clear that Grotmol's system compares indications of detected objects. Grotmol even explains that the robot can use information about "object size, location" during autonomous operation—i.e., at least two characteristics of detected objects in the robot's environment. EX1005, 10:23-34.

15. Even if Grotmol's image features were limited to "a portion of an image" or "a collection of pixels of the image" that correspond to a depicted object as PO argues (incorrectly for the reasons noted above), Grotmol would still anticipate claim 7 in my opinion. A POSITA would have understood that isolating a portion of an image corresponding to a depicted object would involve detecting the object to segment it from the remainder of the image. Accordingly, the object is detected and the detected object (e.g., the portion of the image) would be used in Grotmol's matching process to determine a best matching image from the training buffer. Regardless, for the reasons addressed above, Grotmol's feature set is not limited to "a portion of an image" or "a collection of pixels of the image" that correspond to a depicted object. To the contrary, Grotmol itself lists "values of pixels in an image" as just one example of an image feature apart from a range of other features including indications of the presence or absence of an object at particular locations in the image, and properties, position, and changes in features (e.g., objects) across consecutive frames. EX1005, 28:19-40.

C. Claim 9

16. Claim 9 of the '583 patent recites the following language:

9. The system of claim 4, wherein the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device are included in a knowledgebase, and wherein the knowledgebase further includes a second one or more digital pictures correlated with a second one or more instruction sets for operating the first device, and wherein at least a portion of the first one or more digital pictures or at least a portion of the first one or more instruction sets for operating the first device is learned in a first learning process that includes operating the first device at least partially by a first user, and wherein at least a portion of the second one or more digital pictures or at least a portion of the second one or more instruction sets for operating the first device is learned in a second learning process that includes operating the first device at least partially by the first user.

17. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that claim 9 is anticipated by Grotmol. See EX1003, ¶¶243-246. I do not agree with PO's assertion that "Grotmol discloses only a **single learning process**" in which "multiple tasks are learned in the **single learning process**." POR, 15. As I described in my Prior Declaration, Grotmol discloses a "first learning process" and a "second learning process" as claimed, where each involves operating a first device at least partially by a first user as recited in claim 9. EX1003, ¶¶243-246.

18. For example, Grotmol discloses multiple learning processes for training different behaviors of the robot, including both "simple tasks" and "complex tasks." EX1005, 24:17-67, 26:1-20. Grotmol describes a "hierarchical" learning framework

in which the robot learns complex tasks/behaviors through the individual training of multiple simpler (sub-)tasks/behaviors involved in the complex task. EX1005, 10:35-44; *generally id.*, 24:17-27:17, FIGS. 9-11B. With this approach, different tasks/behaviors can be implemented by different ones of a collection of adaptive controllers (also referred to as “modules”) on the robot, each of which can be trained in a different learning process. FIG. 10D, which is reproduced below for example, shows a pair of “[h]igh-level” adaptive controllers 1 and 2 for more complex tasks that selectively activate different low-level tasks/behaviors controlled by low-level controllers 1-N:




FIG. 10D

EX1005, FIG. 10D (annotated); *see also id.*, 24:29-27:17, FIG. 11B.

19. A POSITA would have understood that the separate training of

different tasks/behaviors on different controllers provides separate “learning process[es]” in Grotmol. To this point, Grotmol explains that “[w]hile training a particular behavior, **the user may concentrate on a particular behavior and may not show any examples of other behaviors** thus making a prediction task for the classifier substantially easier.” EX1005, 24:57-61. Then, “**after [a] user is done training a particular behavior, he may create another module and activate it**” such that “[t]he system may be configured to **learn [a] second behavior** from the user.” *Id.*, 26:1-6. In an illustrative example, Gotmol discloses training the robot to fetch a target object (which is a relatively complex task/behavior) using a first controller trained in a first learning process to “control the robot ... to approach the target while avoiding [] obstacles (‘Behavior 1’)” while a second controller is trained in a second learning process to “control the robot ... to turn and grasp the target (‘Behavior 2’).” *Id.*, 25:30-26:34. A separate higher-level controller orchestrates the overall fetch behavior by selectively activating the controllers for the lower-level tasks of the fetch behavior under appropriate conditions. *Id.*

20. Further confirming the presence of multiple learning processes, Grotmol explains that “**[a]fter the user is done training** a single complex behavior, **he may train another** switcher to switch between different modules and therefore obtain [a] **second complex behavior**.” EX1005, 26:65-67. Grotmol also teaches that a user may individually “reset a module to naive state, revert a module to a

particular point in history, delete a module, save/load locally, upload, download, and/or other operations.” *Id.*, 25:54-59. These examples show that Grotmol’s controllers and their respective tasks/behaviors can be trained separately, at different times, and by different methodologies suitable to the respective task. *Id.* Grotmol discloses multiple learning processes as claimed. Indeed, Grotmol describes a separate RKNN “learning process” applicable to each of multiple controllers on a robot. EX1005, 38:39-42; *see also id.*, 15:24-26, FIG. 8.

21. For each of these reasons, it is clear that Grotmol’s memory stores a knowledgebase that includes at least a first controller/module containing a first one or more digital pictures correlated with a first one or more instructions sets learned in **a first learning process** and a second controller/module containing a second one or more digital pictures correlated with a second one or more instructions sets learned in **a second learning process**. EX1005, 38:39-42, FIG. 8. And while I have noted that PO characterizes the “knowledgebase” in claim 10 as storing “mixed knowledge,” in apparent reference to the inclusion of knowledge learned through multiple learning processes and/or users, Grotmol’s knowledgebase stored in memory of the robot (e.g., memory 814) also provides mixed knowledge at least because it stores training sets of digital pictures and instructions sets (e.g., motor commands) for multiple tasks/behaviors learned in different learning processes by one or more users. EX1005, 3:36-39, 38:39-42, FIG. 8.

D. Claim 14

22. Claim 14 of the '583 patent recites the following language:

14. The system of claim 4, wherein the new one or more digital pictures depict at least a portion of the second device's surrounding, and wherein the first one or more instruction sets for operating the first device or a copy of the first one or more instruction sets for operating the first device are modified and applied to the second device, and wherein the second device is caused to perform one or more operations defined by: the modified the first one or more instruction sets for operating the first device, or the modified the copy of the first one or more instruction sets for operating the first device.

23. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that claim 14 is anticipated by Grotmol. See EX1003, ¶¶256-260. PO argues that “Grotmol only discloses correcting device behavior by issuing new commands ... and does not disclose that any instruction sets are modified” or “applied to [a] second device, where the second device performs an operation defined by the modified instruction set.” POR, 16-18. I do not agree.

24. My Prior Declaration highlighted Grotmol's teachings related to compensating for “position mismatch” and “orientation mismatch” as an example of “modifying” instruction sets. EX1003, ¶¶258-260; EX1005, 42:46-58, Abstract. According to Grotmol, a robot autonomously navigating in its environment can adjust its heading or trajectory when the robot's current position or orientation deviates from a target position or orientation. *Id.* The deviation can be determined by comparing a new digital picture that shows the robot's current surroundings to a

stored picture in the training buffer:

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

... In some implementations the shift amount may be utilized as an error metric by the control process of the robot. The steering signal (which may be adjusted leftwards or rightwards) may be selected as the control variable for the process. A negative feedback loop may be used in order to reduce the error metric to (and/or maintain at) a target level operation of the robot. ...

A PID controller may be used in order to reduce/maintain the error metric during operation of the robot. **In some implementations, motor commands at a given time step may be obtained by taking the stored motor commands from the training buffer that may correspond to the best matching stored image. Those motor commands may be combined with the output from the PID controller in order to stabilize operation of the robot.**

EX1001, 42:52-43:11.

25. Grotmol thus compensates for position/orientation mismatch by

combining “stored motor commands from the training buffer” with a correction, e.g., from a PID controller or corrector. The stored motor commands from the training buffer (**first one or more instruction sets for operating the first device**) are “modified” when they are combined with the output of the PID controller or corrector, and the modified motor commands are “applied” at least when they are provided for execution by the robot. Based on my review, the concepts described at col. 42:52-43:11 relate to and would be read by a POSITA in the context of FIG. 4A, which illustrates how an output 418 (e.g., stored motor commands) of an adaptive predictor 422 can be combined by a combiner 414 with the output 408 of a corrector 412 (akin to a PID controller) to produce a control signal 420 as a corrected, modified version of the output 418:




FIG. 4A

EX1005, FIG. 4A (annotated); *see also id.*, 15:18-17:20, 16:24-31 (“[T]he control signal 408 may comprise a correction to the target trajectory. The signals may comprise a target ‘correction’ to the current velocity and/or steering angle of the platform 430.”). The transfer function applied by combiner 414 “modifi[es]” the output of predictor 422. EX1005, 17:35-49; *generally id.*, 17:35-18:65.

26. Grotmol also discloses applying a modified version, or a copy of a modified version, of instruction sets for operating a first device to a **second device** to cause the **second device** to perform operations defined by the modified instructions. Recall that Grotmol discloses that a “trained configuration may be loaded to one or more other robots in order to provide learned behaviors.” EX1005, 33:30-32, 33:13-15. Hence, when a trained configuration learned by operating a first robot is loaded on and run by a second robot, the second robot can modify the motor commands (instruction sets) from the first robot by combining them with the output of the second robot’s PID controller as described above. As also described above and in my Prior Declaration, the modified commands can then be applied to the second device by providing them for execution.

E. Claim 15

27. Claim 15 of the ’583 patent recites the following language:

15. The system of claim 4, wherein the machine readable code, when executed by the one or more processors, causes the one or more processors to further perform at least:
modifying:
the first one or more instruction sets for operating the first device, or
a copy of the first one or more instruction sets for operating the first device, and wherein the determining the first one or more instruction sets for operating the first device based on the at least partial match between the new one or more digital pictures and the first one or more digital pictures includes:
determining the modified the first one or more instruction sets for operating the first device based on the at least partial match between the new one or more digital pictures and the first one or more digital pictures, or
determining the modified the copy of the first one or more instruction sets for operating the first device based on the at least partial match between the new one or more digital pictures and the first one or more digital pictures, and wherein the causing the first device or the second device to perform the one or more operations defined by the first one or more instruction sets for operating the first device includes:
causing the first device or the second device to perform one or more operations defined by the modified the first one or more instruction sets for operating the first device, or
causing the first device or the second device to perform one or more operations defined by the modified the copy of the first one or more instruction sets for operating the first device.

28. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that claim 15 is anticipated by Grotmol. See EX1003, ¶¶261-267.

29. As an initial matter, I understand that PO argues that “Grotmol only discloses correcting device behavior by issuing new commands in certain instances and does not disclose that any instruction sets are modified.” This is not correct.

Grotmol discloses modifying instructions sets for each of the reasons that I addressed above with respect to claim 14.

30. Next, I do not agree with PO's argument that, when Grotmol "combin[es]" motor commands from the training buffer with the output of a PID controller, that those motor commands somehow are not modified. EX1005, 17:35-49; *generally id.*, 17:35-18:65. Grotmol is clear that its "combining" operation involves modifying the stored motor command or instruction set. The modified instruction sets manifest in trajectories or actions of the robot that are different from what otherwise would have been performed if the instruction sets had not been modified. EX1005, 42:64-65 ("The steering signal ... may be adjusted leftwards or rightwards"), 35:55-60 ("the controllable parameter (e.g., correction 408 in FIG. 4A) may be used to modify the existing state of the system (e.g., comprising motor torque and/or robot platform acceleration)"), 16:25-26 ("the control signal 408 may comprise a correction to the target trajectory"), 15:29-62, 42:46-43:11, FIG. 4A. Indeed, a POSITA would have understood that a PID controller can produce virtually countless outputs as it is a high-order function with both derivative and integral components.

31. PO further argues that claim 15 requires "that the modified instruction set must be the correlated instruction set" such that the instruction set is necessarily modified before the modified instruction sets are determined. *See* POR, 21-22. I do

not agree. Neither the plain language of claim 15 nor the '583 patent's specification require the instruction sets to be modified before the determination involving at least a partial match between the new and stored digital pictures. The explicit recitation of a "modifying" step ensures that the "modifying" must be performed by the claimed one or more processors (e.g., as opposed to the modification being performed by an unclaimed device if the modifying step were recited inferentially). But the specification makes clear that the instruction sets can be modified even as part of the "determining" step itself and after the matching operation:

Confirmation Unit 550 comprises the functionality for confirming, **modifying**, evaluating (i.e. rating, etc.), and/or canceling one or more anticipatory Instruction Sets 526, and/or other functionalities. ... For example, Decision-making Unit 540 may determine one or more anticipatory Instruction Sets 526 and provide them to User 50 for confirmation. User 50 may be provided with an interface (i.e. graphical user interface, selectable list of anticipatory Instruction Sets 526, etc.) to approve or confirm execution of the anticipatory Instruction Sets 526. In some aspects, Confirmation Unit 550 can automate User 50 confirmation. In one example, if one or more incoming Digital Pictures 525 from Picture Capturing Apparatus 90 and one or more Digital Pictures 525 from a Knowledge Cell 800 were found to be a perfect or highly

similar match, anticipatory Instruction Sets 526 correlated with the one or more Digital Pictures 525 from the Knowledge Cell 800 can be automatically executed without User's 50 confirmation. Conversely, if one or more incoming Digital Pictures 525 from Picture Capturing Apparatus 90 and one or more Digital Pictures 525 from a Knowledge Cell 800 were found to be less than a highly similar match, anticipatory Instruction Sets 526 correlated with the one or more Digital Pictures 525 from the Knowledge Cell 800 can be presented to User 50 for confirmation and/or **modifying**. ... In other embodiments, Confirmation Unit 550 can serve as a means of **modifying** or editing anticipatory Instruction Sets 526. **For example, Decision-making Unit 540 may determine one or more anticipatory Instruction Sets 526 and provide them to User 50 for modification. User 50 may be provided with an interface (i.e. graphical user interface, etc.) to modify the anticipatory Instruction Sets 526 before their execution.**

EX1001, 89:22-62; *see also id.*, 160:21-39 (“do not require the particular order shown ... other embodiments are within the scope of the following claims”), 68:45-47 (“Any of the operations of the described elements can be performed repeatedly and/or in different orders in alternate embodiments.”), 11:59-61 (“modifying one or more instruction sets at a source code write time, a compile time, an interpretation

time, a translation time, a linking time, a loading time, or a runtime”).

32. I also note that claim 15 does not require that the modified instruction set be correlated with the first one or more digital pictures or otherwise stored in a knowledgebase with the first one or more digital pictures. Claim 15 recites significantly different language than claim 18, for example, which does refer to “learning” involving modified instruction sets. The embodiments described at col. 89:22-62, for example, are thus relevant to Claim 15.

33. Properly interpreted, Grotmol anticipates claim 15 for at least the reasons described in my Prior Declaration and above with respect to claim 14. For example, Grotmol discloses that the motor commands at a given time step may be obtained by taking the stored motor commands from the training buffer that correspond to the best matching stored image. The motor commands are modified to compensate for, e.g., position mismatch and orientation mismatch. EX1005, 38:21-28, 42:46-43:11, 15:18-53, FIG. 4A. The modified motor commands are then executed to cause a first or second robot to perform one or more operations defined by the modified motor commands. Grotmol’s process thus aligns with the ’583 patent’s preferred embodiments, which similarly involve determining (unmodified) instruction sets based on an at least partial match and then modifying the instructions before they are executed. *See* EX1001, 89:22-62, 68:45-47, 160:21-39, 11:59-61.

34. Even if PO’s erroneous claim interpretation applied, Grotmol would

still anticipate. This would occur, for instance, by taking Grotmol's "combining" operation as the claimed "modifying" step that produces, e.g., a combined/modified control command 420, and taking Grotmol's subsequent processing and decoding of the combined/modified control command prior to execution of the combined/modified control command as the claimed step of "determining" the modified first one or more instruction sets. EX1005, 15:18-62, 42:46-43:11, FIG. 4A. For example, Grotmol discloses "determining the modified first one or more instruction sets" after the first one or more instruction sets are modified by translating a format of the combined/modified motor control command, rescaling the combined/modified motor control command, or both. *Id.*, 15:26-52.

F. Claim 16

35. Claim 16 of the '583 patent recites the following language:

16. The system of claim 4, wherein the machine readable code, when executed by the one or more processors, causes the one or more processors to further perform at least:
modifying at least one of: the first one or more digital pictures, a copy of the first one or more digital pictures, the new one or more digital pictures, or a copy of the new one or more digital pictures, and wherein the at least partial match between the new one or more digital pictures and the first one or more digital pictures includes: (i) at least partial match between the modified the new one or more digital pictures and the first one or more digital pictures, (ii) at least partial match between the modified the copy of the new one or more digital pictures and the first one or more digital pictures, (iii) at least partial match between the new one or more digital pictures and the modified the first one or more digital pictures, (iv) at least partial match between the new one or more digital pictures and the modified the copy of the first one or more digital pictures, (v) at least partial match between the modified the new one or more digital pictures and the modified the first one or more digital pictures, (vi) at least partial match between the modified the copy of the new one or more digital pictures and the modified the copy of the first one or more digital pictures, (vii) at least partial match between the modified the new one or more digital pictures and the modified the copy of the first one or more digital pictures, or (viii) at least partial match between the modified the copy of the new one or more digital pictures and the modified the first one or more digital pictures.

36. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that claim 16 is anticipated by Grotmol. See EX1003, ¶¶268-270.

37. To start, Grotmol discloses “modifying at least one of: the first one or more digital pictures, a copy of the first one or more digital pictures, the new one or more digital pictures, or a copy of the new one or more digital pictures,” as recited in claim 16. For example, my Prior Declaration cited Grotmol’s disclosure that

“[d]uring training, images (e.g., raw and/or pre-processed) may be stored in a memory buffer (training buffer)” and the “preprocessing operations may comprise resampling, cropping, light balancing, and/or feature extraction.” EX1005, 39:13-22. A POSITA would have understood that several of Grotmol’s pre-processing operations such as resampling, cropping, and light balancing involve “modifying” the raw digital picture. In an example, the raw pictures that Grotmol stores in the training buffer provide the claimed “first one or more digital pictures correlated with the first one or more instruction sets” and the additional pre-processed digital pictures stored in the training buffer provided the claimed “modified the first one or more digital pictures.”

38. PO argues that Grotmol fails to disclose “using [a] modified digital picture in a matching. POR, 24. I disagree. Grotmol discloses determining whether there is a partial match between a newly acquired image and a pre-processed digital image stored in the training buffer. For example, as I described in my Prior Declaration with respect to limitation [4.3], Grotmol’s matching process involves calculating a similarity measure between a new image and an old image and determining whether that similarity measure breaches a threshold. Pet., 20; EX1005, 4:30-38 (“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.”), Abstract (“an acquired image may be compared with one or more images from the training buffer in order to determine the most likely match”), 28:19-24 (“the term feature” includes “values of pixels in an image”), claim 12 (“training set” comprises “an instance of the video frame”). Thus, for example, when a pre-processed image stored in the training buffer—i.e., an example of a modified version of the claimed first one or more digital pictures—is compared to a newly acquired image to determine a best matching picture, an at least partial match is determined between the new digital picture and the modified the first one or more digital pictures as claimed.

39. Grotmol also discloses that “[a] similarity metric ... 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.” EX1005, 41:8-13. This disclosure provides another example of Grotmol modifying digital pictures and using the modified digital pictures in a matching process. In particular, Grotmol explains that before the pixel-wise difference is calculated between a newly acquired image and images from the training buffer in the matching search set to determine a best matching image, the images are shifted or trimmed (i.e., modified) to compensate for any misalignment between the images. EX1005,

40:17-41:18, FIGS. 20A-20B. “By way of an illustration, image 2010 may be shifted to the right” and “[t]he shifted image may be trimmed from the edge towards which the shift is performed.” *Id.*, 40:57-63. “The matching image (e.g., 2000 in FIG. 20A) may be trimmed on the left by the amount [of] the shift.” *Id.*, 40:61-63. “A similarity metric may be determined between the shifted/trimmed frames (e.g., 2020, 2040 in FIG. 20B).” *Id.*, 41:8-9.

40. Based on my review, PO also mischaracterizes claim 16 in alleging that “claim 16 requires that the correlated digital picture is modified.” PO, 25. The examples described above do involve modifying the correlated digital picture, but the breadth of claim 16 is not so limited. Claim 16 also encompasses modifying the new one or more digital pictures (e.g., a newly acquired picture) rather than or in addition to modifying the first one or more digital pictures that are correlated. EX1001, Claim 16. Grotmol’s disclosures of scaling, normalizing, and/or equalizing an instance of the sensory input (e.g., a new digital picture) and using the scaled/normalized/equalized sensory input in a matching process further meet these alternative aspects of claim 16. EX1005, 4:18-22, 30:47-31:2, 31:61-32:5, 3:41-46, 4:58-63, 31:3-8.

G. Claim 17

41. Claim 17 of the ’583 patent recites the following language:

17. The system of claim 4, wherein the machine readable code, when executed by the one or more processors, causes the one or more processors to further perform at least: modifying:
the first one or more digital pictures, or
a copy of the first one or more digital pictures, and
wherein the learning the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device includes:
learning the modified the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device, or
learning the modified the copy of the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device, and
wherein the determining the first one or more instruction sets for operating the first device based on the at least partial match between the new one or more digital pictures and the first one or more digital pictures includes:
determining the first one or more instruction sets for operating the first device based on at least partial match between the new one or more digital pictures and the modified the first one or more digital pictures, or
determining the first one or more instruction sets for operating the first device based on at least partial match between the new one or more digital pictures and the modified the copy of the first one or more digital pictures.

42. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that claim 17 is anticipated by Grotmol. See EX1003, ¶¶271-280. Much of my analysis from claim 16 applies to claim 17 as well. Above, Claim 16.

43. I do not agree with PO's assertion that Grotmol fails to disclose "learning [a] modified digital picture" and "using this modified digital picture in a matching." POR, 26. Grotmol teaches that "[d]uring training, images (e.g., raw

and/or pre-processed) may be stored in a memory buffer (training buffer)” and the “preprocessing operations may comprise resampling, cropping, light balancing, and/or feature extraction.” EX1005, 39:13-22. Several of Grotmol’s pre-processing operations, including resampling, cropping, and light balancing involve “modifying” a digital picture at least because they alter the raw digital picture in various ways. Further, the pre-processed (modified) digital pictures are stored in the training buffer correlated with “[m]otor commands issued by a trainer corresponding to time instances when the images are acquired.” EX1005, 39:13-26, 3:35-38. Grotmol’s system thus learns the pre-processed (modified) digital pictures correlated with motor commands (e.g., instruction sets) for operating a first robot consistent with the features recited in claim 17. EX1005, 33:59-62 (“Trained behaviors of the robot 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.”), 22:4-5, 43:5-9.

44. Grotmol also discloses use of pre-processed (modified) digital pictures in a matching process to determine motor commands (e.g., instruction sets) for operating a first device. EX1005, Abstract, 2:30-37, 4:30-38, 28:19-40, 41:8-13. In fact, Grotmol explains that “motor commands at a given time step may be obtained by taking the stored motor commands from the training buffer that may correspond to the best matching stored image.” EX1005, 43:5-9; *see also id.*, 39:13-22. As the

“stored image” can be a pre-processed (modified) image, the “best matching” image can be a pre-processed (modified) image. EX1005, 39:13-22

H. Claim 18

45. Claim 18 of the '583 patent recites the following language:

18. The system of claim 4, wherein the machine readable code, when executed by the one or more processors, causes the one or more processors to further perform at least:

- modifying:
 - the first one or more instruction sets for operating the first device, or
 - a copy of the first one or more instruction sets for operating the first device, and wherein the learning the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device includes:
 - learning the first one or more digital pictures correlated with the modified the first one or more instruction sets for operating the first device, or
 - learning the first one or more digital pictures correlated with the modified the copy of the first one or more instruction sets for operating the first device, and wherein the determining the first one or more instruction sets for operating the first device based on the at least partial match between the new one or more digital pictures and the first one or more digital pictures includes:
 - determining the modified the first one or more instruction sets for operating the first device based on the at least partial match between the new one or more digital pictures and the first one or more digital pictures, or
 - determining the modified the copy of the first one or more instruction sets for operating the first device based on the at least partial match between the new one or more digital pictures and the first one or more digital pictures, and wherein the causing the first device or the second device to perform the one or more operations defined by the first one or more instruction sets for operating the first device includes:
 - causing the first device or the second device to perform one or more operations defined by the modified the first one or more instruction sets for operating the first device, or
 - causing the first device or the second device to perform one or more operations defined by the modified the copy of the first one or more instruction sets for operating the first device.

46. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that claim 18 is anticipated by Grotmol. See EX1003,

¶¶281-291. For example, PO argues that “Grotmol does not disclose learning any modified instruction set.” POR, 31. I do not agree. As I described in my Prior Declaration, Grotmol discloses an “override mode” for training a robot. EX1005, 23:4-22. “Override” is one of at least four training modes available to Grotmol’s robot, as shown in Table 3:

TABLE 3

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

EX1005, Table 3 (22:27-46) (annotated).

47. Grotmol explains that “[i]n the override mode, the robot may execute what it has learned, unless a command is sent by the teacher.” EX1005, 23:13-14. “As soon as the teacher starts sending commands, the robot may stop taking the initiative and may let the teacher control it.” *Id.*, 23:14-16. “For example, if the

robot is turning left but the teacher wants the robot to turn right and provides a right turn command, then the robot may head [sic] the teacher's command, perform the action, **and try to remember it for the next time the same situation occurs.**" *Id.*, 23:16-21; *see also id.*, 23:59-63, 35:16-24, 35:35-50, FIGS. 18-19. When the teacher issues command(s) in override mode, Grotmol's system modifies previously learned command(s) by replacing them with the teacher's new command(s) that the robot learns and then applies. *Id.*

48. I also do not agree with PO's argument that the claimed "modifying" step does not include "replacing" instructions as allegedly occurs in Grotmol's override mode. POR, 30-31 ("Replacing an element with another element is fundamentally different than modifying the original element."). A POSITA would have understood that replacing a motor command with another motor command in an instruction set modifies the set of learned motor commands, much like the roster of a sports team changes when players are traded to other teams. PO's position is also inconsistent with the '583 patent's own descriptions of "modifying" as encompassing "replacing." EX1001, 147:36-42 ("Executing may include **modifying** an element of a logic circuit with one or more **alternate instruction sets**"), 136:9-11 ("Ptrace can be used to modify a running application such as **modifying** an application with **alternate instruction sets**"), 140:15-18 ("**Modifying** execution and/or functionality of Logic Circuit 250 can be implemented by

replacing its input values with anticipatory input values”), 140:31-34, 141:2-6.

49. Grotmol’s “correct mode” also anticipates the disputed claim features. “In the correct mode, the robot may integrate what the teacher commands with what the robot already knows,” e.g., by “sum[ming] the teacher’s command with what the robot already knows to get a final motor command.” EX1005, 23:23-29. The robot then “[l]earns the resulting association.” *Id.*, 22:40-43, Table3, 23:30-35. Rather than replacing instructions with entirely new instructions, correct mode involves determining a modified instruction based on the robot’s known instruction and the teacher’s new instruction, which is then learned, determined, and applied in a similar manner to the instructions in override mode. *Id.*

50. PO’s argument that Grotmol’s override mode does not involve “learning” is also unfounded. POR, 30, EX1005, 23:20-22. PO overlooks express disclosure from Table 3 that the robot “**[l]earns the association**” when a teacher’s input is received in override mode. EX1005, 22:37-39. The teacher’s input is explicitly learned in Grotmol, and for the teacher’s input to be accounted for “the next time the same situation occurs,” that input (e.g., a modified instruction set) is stored in the training buffer and correlated with digital pictures. EX1005, 39:13-22, 43:5-11, 3:27-39, Abstract, 33:59-62, 42:37-40, 42:46-58; Pet., p. 34 and analysis of [1.6], [4.3]-[4.4], [18.2]-[18.3]. The modified motor commands learned during override mode are then determined and applied to cause the robot to perform the

operations defined by the motor commands for the reasons described in my Prior Declaration.

I. Claim 23

51. Claim 23 of the '583 patent recites the following language:

23. The system of claim 4, wherein the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device include the first one or more digital pictures connected, using at least one or more connections, with the first one or more instruction sets for operating the first device.
--

1. PO's Construction of "Connection" is Unsupported by the '583 Patent and File History

52. 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 '583 Patent includes no lexicography or disavowal to support Patent Owner's construction of "connection" as a relationship that includes a "link" and the POR does not even attempt to demonstrate lexicography or disavowal.

53. 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. POR, 7-11. For example, PO turns to FIGS. 16 and 17A, which are

described as mere “embodiments.” EX1001, 57:43-48. Counsel for Petitioner has indicated to me that the word “embodiments” does not impose a limiting effect. The ’583 Patent explains that Model A in FIG. 16 “**may** include a network of Nodes 852 ... and Connections 853 similar to that of a brain.” EX1001, 95:65-96:36. In FIG. 17A, “[e]ach of Connection 853z1 and Connection 853z2 **may** include or be associated with occurrence count, weight, **and/or other parameter or data.**” EX1001, 106:19-22. The ’583 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.” EX1001, 15:58-60. These descriptions include a permissive “may” and do not necessarily limit the scope of a claim term.

54. Notably, the textual description of FIGS. 16 and 17A do not once use the term “link,” which PO seeks to import into the construction of “connection.” The specification also broadly describes “connections”—both “direct” and otherwise—at col. 67:21-31 without every referring to or mandating that a connection include a “link” as proposed by PO. EX1001, 67:20-31.

55. In addition to the lack of support, PO’s construction is problematic because it includes the term “link”—the meaning of which is unclear. The ’583 Patent provides no explanation as to what a “link” necessarily is and what it includes or excludes. 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 could not clearly explain what the term meant. EX1016, 12:12-36:5. With no further clarity, a reader is left to guess what the scope of the term “connections” is under PO’s construction because it’s unclear what the word “link” encompasses or how it manifests in practice. This is exemplified by the fact that the connections 853 that PO points to in the Reply (from FIGS. 16 and 17A) as “links” are merely abstractions that simply provide visual depictions of the association or relationship between the nodes 852 or knowledge cells 800 in a graph. Given the wide range of models disclosed in the ’583 patent that may include connections, any construction of the term (even if construed to include a “link”) must account for the breadth of examples provided in the patent. *See* EX1001, 95:49-64 (“Referring to FIG. 16, the disclosed artificially intelligent systems, devices, and methods for learning and/or using visual surrounding for autonomous device operation may include various artificial intelligence models and/or techniques. ... Examples of these models and/or techniques include deep learning, supervised learning, unsupervised learning, neural networks (i.e. convolutional neural network, recurrent neural network, deep neural network, etc.), search-based, logic and/or fuzzy logic-based, optimization-based, tree/graph/other data structure-based, hierarchical, symbolic and/or sub-symbolic, evolutionary, genetic, multi-agent, deterministic, probabilistic, statistical, and/or other models

and/or techniques.”).

56. PO cites two dictionary definitions (EX2012, EX2013) as allegedly supporting its construction, each of which provides numerous definitions for “connection.” POR, 14. PO chose a definition that includes the word “link” but did not include any explanation as to why PO chose this definition as opposed to the various other available definitions, *particularly when the description in the ’583 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.


57. Considering the broad descriptions of the term “connection” in the ’583 patent, a re-definition or disavowal would be inappropriate and the term’s plain and ordinary meaning should apply. EX1001, 15:58-60, 66:3-23. As shown in the dictionaries relied on 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 ’583 patent. EX2012, 399-400, EX2013, 481; EX1001, 15:58-60, 67:21-31, 66:3-23. Thus, when two nodes or points of data have a connection, like in FIGS. 16 and 17A of the ’583 patent, an “association or relationship” between the nodes or points of data provides a “connection.”

2. Grotmol Discloses Claim 23 and the Recited “Connections”

58. Under the plain meaning of “connection” as an association or relationship, Grotmol discloses claim 23 for the reasons addressed in the Petition. Pet., 52-53. PO notably concedes that Grotmol discloses an “association” between images and instruction sets. POR, 33 (“Grotmol’s disclosure of ... associations ...”).

59. Even if “connection” were narrowly construed to require a link or a specific connection beyond a “mere” association or relationship, Grotmol still anticipates. For example, Grotmol discloses that “[a] given training set may comprise an instance of the video frame[] ... and an instance of a motor control indication configured to cause the robot to execute an action.” EX1005, 3:36-38. Thus, Grotmol organizes the images and motor commands in the training buffer into training sets where each training set specifically links “an instance of the video frame” (e.g., digital picture) to “an instance of a motor control indication” (e.g., instruction set(s)). Grotmol’s images and motor commands are not merely correlated logically, but are specifically connected via the training sets. EX1005, 4:38-42, 46:66-47:3, 33:59-62, 29:37-30:19, 33:24-27, FIGS. 14-15 (depicting connected sets of image features ‘X’ and motor control commands ‘Y’ in training sets of RKNN classifier). By contrast, images and motor commands in the training buffer that belong to different training sets are not specifically connected or linked in the same way as the images and motor commands that are provided in the same

training set. Further, as I addressed above, when two nodes or points of data have a connection, like in FIGS. 16 and 17A of the '583 patent and like the training sets in Grotmol, an “association or relationship” between the nodes or points of data provides a “connection.” The following figure conceptually illustrates this point similar to FIGS. 16 and 17A of the '583 Patent:



J. Claim 29

60. Claim 29 of the '583 patent recites the following language:

29. The system of claim 4, wherein the generating the new one or more digital pictures includes:
detecting one or more objects in the first device's surrounding or the second device's surrounding;
generating one or more representations of the one or more objects; and
generating the new one or more digital pictures that include the one or more representations of the one or more objects.

61. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that claim 29 is anticipated by Grotmol. See EX1003, ¶¶326-330. PO mischaracterizes claim 29 in arguing that it requires a causal chain of “detect[ing] objects in [the] images in order to generate representations of the

detected objects in order to then generate a new digital picture.” POR, 36. Claim 29, however, does not require that object detection be performed specifically “in order to generate representations of the detected objects” or even that the representations of the objects be generated based on or before the object detection step. Nothing in the plain language of claim 29 supports PO’s positions in this regard. For example, claim 29 does not recite “generating one or more representations of the one or more detected objects” as alleged. The generating step only refers to “representations of the one or more objects” regardless of if they have been previously detected.

62. Even if claim 29 were interpreted to require that a new digital picture be generated using a previous detection of an object, Grotmol still anticipates. For example, FIG. 20A of Grotmol shows digital pictures 2000 and 2010 where picture 2000 is acquired at time t1 and picture 2010 is acquired at a later time t2. EX1005, 38:43-39:12, FIG. 2A. Objects such as tree 2002, rectangle 2004, and ball 2008 can be detected from the earlier picture 2000. EX1005, 10:23-34 (“state of the environment (e.g., object size, location)”), 8:57-9:3 (“object recognition”), 9:41-52 (“object recognition and/or classification”), 7:60-63 (“representation of an object edge”), 28:14-40 (“object, at particular locations in the image”). Grotmol explains that objects like those detected in picture 2000 at time t1 are among features that Grotmol uses to determine a next control action or behavior such as to turn the robot

or to perform a target tracking or grasping task. EX1005, 10:23-34, 28:14-40, 23:3-17, 24:31-25:5, 25:30-42, 46:50-13 (“compare individual features”), FIG. 15. The objects detected in picture at time t1 thus cause the robot to perform actions (e.g. turn), thereby positioning the robot to capture picture 2010 (e.g., a new digital picture) at time t2 that includes detected objects from picture 2000, although from a different perspective. EX1005, 38:43-39:12. The new digital picture 2010 includes representations of the earlier detected objects, and generating the new digital picture 2010 includes generating representations (e.g., depictions) of the objects in the digital picture 2010. EX1005, 38:54-55 (“Images 2000, 2010 may comprise object representations ...”), 38:65-66 (“object representations ... may appear”). Grotmol therefore discloses each of the steps recited in claim 29.

K. Claim 31

63. Claim 31 of the '583 patent recites the following language:

31. The system of claim 4, wherein the generating the first one or more instruction sets for operating the first device includes generating the first one or more instruction sets for operating the first device as one or more representations of a second one or more instruction sets for operating the first device.

64. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that claim 31 is anticipated by Grotmol. See EX1003, ¶¶331-333. As I described in my Prior Declaration, Grotmol discloses generating first one or more instruction sets as a representation of second one or more

instruction sets such as by reloading a previous configuration of the robot—including “[b]y loading at time t_1 the configuration saved at time $t_0 < t_1$.” EX1005, 34:25-27, 34:18-21, 34:43-45. The reloaded configuration includes a first one or more instruction sets at time t_1 that represents the prior configuration including the second one or more instruction sets of the robot at time t_0 . *Id.* The reloaded configuration comprises instruction sets that represent the instruction sets of the earlier configuration at least because the reloaded configuration portrays and provides a likeness or reproduction of the instruction sets from the earlier configuration. EX1018, 1481-1482 (“a thing[] ... that depicts a likeness or reproduction of someone or something”); EX1019, 1216-1217 (“a likeness, image, picture, etc.”). Moreover, neither the specification nor file history of the ’583 patent contain any disavowals or redefinitions that would restrict the scope of a “representation” narrower than its plain and ordinary meaning. EX1001, 79:33-36. In a related patent involving similar technology, the inventor instead conveyed a broad understanding of the term “represent[]” to include “any information” about the represented object including object itself. EX1015, 69:5-8.

III. GROUND 2: GROTMOL RENDERS OBVIOUS CLAIMS 10 AND 11

A. Claim 10

65. Claim 10 of the ’583 patent recites the following language:

10. The system of claim 4, wherein the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device are included in a knowledgebase, and wherein the knowledgebase further includes a second one or more digital pictures correlated with a second one or more instruction sets for operating the first device, and wherein at least a portion of the first one or more digital pictures or at least a portion of the first one or more instruction sets for operating the first device is learned in a first learning process that includes operating the first device at least partially by a first user, and wherein at least a portion of the second one or more digital pictures or at least a portion of the second one or more instruction sets for operating the first device is learned in a second learning process that includes operating the first device at least partially by a second user.
--

66. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that Grotmol in view of the knowledge of a POSITA renders claim 10 obvious. *See* EX1003, ¶¶339-341.

67. I disagree with PO's assertion that Grotmol "fails to disclose any second learning process or second user," as recited in claim 10. POR, 11. As I explain further below, Grotmol discloses or renders obvious both (i) multiple learning processes and (ii) multiple users. The claimed "second learning process that includes operating the first device at least partially by a second user" would have been obvious to a POSITA based on Grotmol.

1. Grotmol's Multiple Learning Processes

68. Grotmol discloses multiple learning processes for training different behaviors of the robot, including both "simple tasks" and "complex tasks." EX1005,

24:17-67, 26:1-20. Grotmol describes a “hierarchical” learning framework in which the robot learns complex tasks/behaviors through the individual training of multiple simpler (sub-)tasks/behaviors involved in the complex task. EX1005, 10:35-44; *generally id.*, 24:17-27:17, FIGS. 9-11B. With this approach, different tasks/behaviors can be implemented by different ones of a collection of adaptive controllers (also referred to as “modules”) on the robot, each of which can be trained in a different learning process. FIG. 10D, which is reproduced below for example, shows a pair of “[h]igh-level” adaptive controllers 1 and 2 for more complex tasks that selectively activate different low-level tasks/behaviors controlled by low-level controllers 1-N: controllers 1-N:




FIG. 10D

EX1005, FIG. 10D (annotated); *see also id.*, 24:29-27:17, FIG. 11B.

69. A POSITA would have understood that the separate training of different tasks/behaviors on different controllers provides separate “learning process[es]” in Grotmol. To this point, Grotmol explains that “[w]hile training a particular behavior, **the user may concentrate on a particular behavior and may not show any examples of other behaviors** thus making a prediction task for the classifier substantially easier.” EX1005, 24:57-61. Then, “**after [a] user is done training a particular behavior, he may create another module and activate it**” such that “[t]he system may be configured to **learn [a] second behavior** from the user.” *Id.*, 26:1-6. In an illustrative example, Gotmol discloses training the robot to fetch a target object (which is a relatively complex task/behavior) using a first controller trained in a first learning process to “control the robot ... to approach the target while avoiding [] obstacles (‘Behavior 1’)” while a second controller is trained in a second learning process to “control the robot ... to turn and grasp the target (‘Behavior 2’).” *Id.*, 25:30-26:34. A separate higher-level controller orchestrates the overall fetch behavior by selectively activating the controllers for the lower-level tasks of the fetch behavior under appropriate conditions. *Id.*

70. Further confirming the presence of multiple learning processes, Grotmol explains that “**after the user is done training** a single complex behavior, **he may train another** switcher to switch between different modules and therefore obtain [a] **second complex behavior**.” EX1005, 26:65-67. Grotmol also teaches

that a user may individually “reset a module to naive state, revert a module to a particular point in history, delete a module, save/load locally, upload, download, and/or other operations.” *Id.*, 25:54-59. These examples show that Grotmol’s controllers and their respective tasks/behaviors can be trained separately, at different times, and by different methodologies suitable to the respective task. *Id.* Grotmol discloses multiple learning processes as claimed. Indeed, Grotmol describes a separate RKNN “learning process” applicable to each of multiple controllers on a robot. EX1005, 38:39-42; *see also id.*, 15:24-26, FIG. 8.

71. For each of these reasons, it is clear that Grotmol’s memory stores a knowledgebase that includes at least a first controller/module containing a first one or more digital pictures correlated with a first one or more instruction sets learned in **a first learning process** and a second controller/module containing a second one or more digital pictures correlated with a second one or more instruction sets learned in **a second learning process**. EX1005, 38:39-42, FIG. 8. And while I have noted that PO characterizes the “knowledgebase” in claim 10 as storing “mixed knowledge,” in apparent reference to the inclusion of knowledge learned through multiple learning processes and/or users, Grotmol’s knowledgebase stored in memory of the robot (e.g., memory 814) also provides mixed knowledge at least because it stores training sets of digital pictures and instructions sets (e.g., motor commands) for multiple tasks/behaviors learned in different learning processes by

one or more users. EX1005, 3:36-39, 38:39-42, FIG. 8.

2. Grotmol's Multiple Users for Training the Robot in Different Learning Processes

72. I disagree with PO's assertions that "Grotmol only discloses one user" and that Grotmol allegedly fails to disclose or render obvious "a second learning process that includes operating the first device at least partially by a second user." POR, 11. Based on my review, Grotmol does not limit the number of users who can train the multiple behaviors for the robot. To the contrary, a POSITA would have understood and found obvious that different users can be involved in training different behaviors of Grotmol's robot, such as different users for different simple or complex tasks EX1003, ¶341.

73. As I explained above, Grotmol describes a first learning process that involves a user teaching the robot "to approach [a] target while avoiding [] obstacles ('Behavior 1')" and a second learning process that involves a user teaching the robot "to turn and grasp the target ('Behavior 2')." EX1005, 25:30-26:34. Both learning processes are user-led and rely on a user to instruct the robot how to perform the task at hand. *Id.* There is no impediment in Grotmol to different users leading different learning processes. EX1005, 24:19-30; *generally id.* 24:19-27:40. In fact, switching users would be as simple as handing off a gamepad or other remote device used for controlling the robot from one user to another. EX1005, 33:18-20. A POSITA

would have been familiar with prior art devices like Grotmol's that were commonly trained by multiple users in different learning processes. It would have been obvious for a POSITA to involve multiple users in training Grotmol's robot at least to benefit from the varying skills and knowledge of the different users in training the robot to perform different tasks/behaviors.

74. Moreover, Grotmol explains that learning is user dependent. In particular, Grotmol discloses that:

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 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.*

EX1004, 33:21-35.

75. Considering that Grotmol's learning is user dependent, it would have been obvious to a POSITA that multiple users could operate a single robot and train the robot differently according to the users' instructions with Grotmol's system.

B. Claim 11

76. Claim 11 of the '583 patent recites the following language:

11. The system of claim 4, wherein the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device are included in a knowledgebase, and wherein the knowledgebase further includes a second one or more digital pictures correlated with a second one or more instruction sets for operating a third device, and wherein at least a portion of the first one or more digital pictures or at least a portion of the first one or more instruction sets for operating the first device is learned in a first learning process that includes operating the first device at least partially by a first user, and wherein at least a portion of the second one or more digital pictures or at least a portion of the second one or more instruction sets for operating the third device is learned in a second learning process that includes operating the third device at least partially by: the first user, or a second user.

77. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that Grotmol in view of the knowledge of a POSITA renders claim 11 obvious. See EX1003, ¶¶342-344.

78. PO errs in arguing that “Grotmol only discloses one user operating one device in one learning process.” POR, 17. As I explained in detail above with respect to claim 10, it is not true that Grotmol provides just one user in one learning process. Indeed, the reasons indicating obviousness of claim 10 with respect to these features would have been accentuated when different users are used to train different devices in different learning processes as recited in claim 11. A POSITA would have understood that different devices are commonly trained by different users, especially when the devices belong to different users as they can be in Grotmol.

79. Grotmol also does not limit the knowledge provided by its knowledgebase in memory 814 to training sets learned on just a single device as

alleged. Grotmol instead discloses that “[t]he trained controller of the robot may comprise a trained configuration configured to enable autonomous operation (e.g., without user input) by the robotic device,” where “[t]he trained configuration may be stored” and “[a] saved configuration may be loaded into the robot being trained thereby providing one or more trained behaviors to the robot.” EX1005, 33:24-30. Further, a “trained configuration may be loaded to **one or more other robots** in order to provide learned behaviors.” EX1005, 33:30-32. Thus, the training sets and trained configuration stored on a first robot can include training sets learned on another robot (e.g., learned by operating a third robot and shared with the first robot). EX1005, 3:36-39 (describing training sets including digital pictures correlated with motor commands learned during training).

80. Contrary to PO’s arguments, a POSITA would have understood and found obvious from Grotmol’s disclosure at col. 33:30-32 that Grotmol’s knowledgebase can include so-called **mixed knowledge** including training sets learned in a first learning process on a first device and training sets learned in a second learning process on a third device. For PO’s position to be true (i.e., that Grotmol’s knowledgebase only stores training sets learned on a single device), a POSITA would need to assume that the knowledge stored in Grotmol’s memory would include only a single trained configuration and that any trained configurations learned on the first robot would be flushed from the memory when another trained

configuration is shared by another robot (e.g., a third device) and saved or loaded on the first robot. But this is not how a POSITA would have understood Grotmol and is inconsistent with Grotmol's repeated descriptions of its robots storing **multiple** trained configurations from which a desired configuration can be loaded from memory.

81. For example, Grotmol explains that a version of a trained configuration from an earlier time can be loaded from memory in place of the current configuration using an “undo” and/or “time machine” operation. EX1005, 34:18-48. “The undo operation may be effectuated using a timeline comprising, e.g., a plurality of bookmarks (e.g., shown in FIG. 5) indicative of one or more date/time, context, action, and/or other attributes of association learning.” *Id.*, 34:39-45. The availability of earlier trained configurations indicates that Grotmol's knowledgebase does not include only the current configuration, but further includes other configurations and their associated training sets that may be unused at a given time. Requiring a trained configuration learned on a first robot to be removed from memory when another configuration is loaded from another robot would be inconsistent with Grotmol's approach to maintaining multiple training configurations in memory. Grotmol explains that “[s]ubsequent to loading of [a] saved configuration, the controller learning process may match process configuration being present during saving of the configuration,” which indicates that

multiple configurations are available to be matched and that the multiple configurations are saved. EX1005, 33:32-35.

82. Maintaining multiple configurations and training sets from multiple devices in memory also would have been obvious to a POSITA, who would have appreciated that such an approach was conventional in the art by the relatively late timeframe of the '583 patent. In fact, a POSITA would have been familiar with other prior art publications by the same assignee of Grotmol (Brain Corp.) that similarly disclose *multiple* trained configurations (e.g., brain images) stored in memory. EX1020, [0153], [0157] (“The user then selects one or more computational **brain images**. The computational **brain images** may be transmitted for local storage on the AM of a robotic device”; “Full **images** may be loaded on to an autonomy module (AM) and run as a self-sufficient control application. Partial **images** may lack the full functions necessary to run certain features of the robotic device. Thus, partial images may be used to augment or upgrade (downgrade) a pre-loaded computational brain image or a stored computational brain image. ... In some implementations, **two or more partial computational brain images may be combined to form full computational brain images.**”); EX1021, [0436] (“Network node and connection information for a current task may be **saved for future use** and flushed. **Previously stored network configuration may be loaded** in place of the network node and connection information for the current task.”).

83. Grotmol's teaching that "[a] trained configuration may be loaded to one or more other robots in order to provide learned behaviors" does not mean that the receiving robot would discard other trained configurations saved on the robot, particularly since multiple trained configurations can themselves be used to perform complex tasks. EX1005, 33:27-32. As discussed above (claim 10), Grotmol's robots instead use a collection of controllers, each of which provides a trained configuration for one or more different tasks or behaviors learned in a respective learning process. *Id.*, 33:27-32 ("The **trained controller** of the robot may comprise a **trained configuration** ..."), FIGS. 10B-10D (depicting **multiple adaptive controllers**—each having a respective trained configuration). The trained configuration shared from one robot to another robot in Grotmol is therefore not always a trained configuration that defines all the tasks/behaviors of the robot, but rather corresponds to the task(s)/behavior(s) of a particular controller of which there may be one or many in Grotmol). Additionally, the fact that Grotmol's robot can include multiple controllers each having a respective trained configuration indicates that the robot would have sufficient hardware capabilities to implement the behaviors defined by multiple controllers in series, parallel, or both.

84. In this context, a POSITA would have understood Grotmol's disclosure at col. 33:30-32 to mean that a trained configuration containing behavior(s) learned on a third robot and loaded on a first robot can be used together with the trained

configurations for other controllers—i.e., controllers that Grotmol expressly discloses are learned locally through learning processes on the first robot. The result is a knowledgebase on the first robot that includes knowledge learned in learning process on both the first robot and other robots. *Id.*




FIG. 10B

EX1004, FIG. 10B (annotated).

85. To the extent PO argues that Grotmol must provide mixed knowledge within a single controller to render claim 11 obvious, Grotmol would still render these features obvious. A single controller on a first robot in Grotmol can contain a training set that includes correlated images and motor commands, a portion of which were learned in a learning process on the first robot and a portion of which were learned in another learning process on a second robot. This would occur, for

example, when the first robot loads a trained configuration from a second robot and a user then modifies the trained configuration through additional training on the first robot. EX1005, 23:12-47 (describing “override” and “correct” modes), 33:30-32.

IV. GROUND 3: THE GROTMOL-HICKMAN COMBINATION RENDERS OBVIOUS CLAIMS 25 AND 27


A. Combining Grotmol and Hickman Would Have Been Obvious to a POSITA

86. I disagree with PO’s argument that Hickman would not have been obvious to combine with Grotmol because Hickman allegedly “stor[es] instruction sets and images separately in separate knowledgebases (images in Unique Object Knowledge Base 307 and instructions in Task Knowledge Base 309[] ...), 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, 21.

87. To start, PO’s argument is besides the point because the Grotmol-Hickman combination that I described in my Prior Declaration did not rely upon the specific data structures or organization of Hickman’s knowledgebases. EX1003, ¶¶159-166, 346. The combination instead applied Hickman’s suggested use of a “shared robot knowledgebase” to Grotmol’s system to facilitate knowledge sharing between robots through the cloud. *Id.* In the combination of Grotmol and Hickman, the training configuration learned by operation of a robot described in Grotmol

would be uploaded to the server containing the ‘shared robot knowledge base’ as described in Hickman. EX1003, ¶356. Thus, the knowledgebase in the combination stores the robot’s “training configuration” as described in Grotmol, and the training configuration itself includes training sets comprised of digital pictures correlated with instruction sets. EX1005, 3:35-39.

88. I also do not agree that Hickman teaches away from correlated images/instruction sets like those described in Grotmol. Based on my review, Hickman does not criticize or disparage such correlations in any way. Hickman instead affirmatively describes correlated images and instruction sets. For example, as shown below in FIG. 3, Hickman discloses a data structure 314 for a general object knowledge base 308 (both highlighted in yellow). EX1007, 15:28-37. The data structure 314 includes, among other data, object identifier (GID), image data (IMG), and task data (TASKS). EX1007, 15:28-37.



EX1007, FIG. 3 (annotated).

89. Hickman explains that the task 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.” EX1007, 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 [sic] **may not require a separate task knowledge base 309,**” although, in

other embodiments, the task data “can be looked up in a separate task knowledge base 309.” EX1007, 19:30-38. Thus, Hickman discloses embodiments in which task data that includes instructions sets are provided in the **same knowledgebase** (e.g., general object knowledge base 308) as the images, and the images and tasks are correlated in the data structure 314. PO’s argument that Hickman is somehow limited to “storing instruction sets and images separately in separate knowledgebases” contradicts Hickman’s disclosure.

90. Even if PO argues that only the unique object knowledge base 307 should be considered in the Grotmol-Hickman combination, 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.” EX1007, 12: 5-18. Therefore, 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 since Hickman itself contemplates the addition of other fields to the unique object knowledge base 307 and does not exclude the addition of task data.

91. PO also failed to acknowledge that the Grotmol-Hickman combination applies aspects of an entire “shared robot knowledge base” in the combination,

which in Hickman encompasses each of the lower-order knowledge bases 307, 308, and 309. EX1007, FIG. 3; Pet. 27-28. Regardless of how data is stored across knowledge bases 307-309, the data is nonetheless all contained within knowledge base 306, and images for both individual unique objects and general objects are correlated with tasks at least through the UIDs, GIDs, and TIDs. EX1007, FIG. 3; *generally id.*, 10:37-22:51. This further undercuts PO's and Dr. Saber's position that Hickman teaches away from the correlated digital pictures and instruction sets.

B. Claim 25

92. Claim 25 of the '583 patent recites the following language:

25. The system of claim 4, wherein the system further comprising: a server that receives from the first device at least one of: the first one or more digital pictures, or the first one or more instruction sets for operating the first device, and wherein the second device receives from the server at least one of: the first one or more digital picture % or the first one or more instruction sets for operating the first device, and wherein the second device is caused to perform the one or more operations defined by the first one or more instruction sets for operating the first device.
--

93. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that the Grotmol-Hickman combination renders claim 25 obvious. See EX1003, ¶¶347-357.

94. PO's arguments challenging the obviousness of claim 25 are deficient for many of the same reasons as those addressed above. *Above*, Section IV.A. For

instance, PO again focuses on Hickman’s allegedly “simple knowledgebase that includes uncorrelated images and instruction sets.” POR. 26. But this characterization of Grotmol is both inaccurate and immaterial to the proposed combination. *Above*, Section IV.A. PO also argues that claim 25 “requires a sophisticated approach where knowledge is stored and processed using instruction sets correlated with a digital picture in a specific manner and ultimately transferred from one device to be used on another device.” POR, 26-27. While it is unclear what “sophisticat[ion]” is required, the Grotmol-Hickman combination communicates “trained configurations”—including digital pictures correlated with instruction sets as described in Grotmol—from a first device to a second device via a server as suggested by Hickman and consistent with the requirements of claim 25. EX1003, ¶¶347-357. Moreover, Hickman demonstrates that the sharing of learned behaviors and instructions for autonomous devices was well-known in the art. EX1007, 11:30-41. Storing and sharing digital pictures along with instruction sets is disclosed by both Grotmol and Hickman. *See e.g.*, EX1005, FIGS. 2, 4A, 4B, 3:27-39, Abstract, 8:66-9:3, 10:23-26, 38:36-58, 39:13-22; EX1005, FIGS. 3, 4, 13:35-24:64.

C. Claim 27

95. Claim 27 of the ’583 patent recites the following language:

27. The system of claim 4, wherein the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device are included in a knowledgebase, and wherein the knowledgebase further includes a second one or more digital pictures correlated with a second one or more instruction sets for operating a third device, and wherein the machine readable code, when executed by the one or more processors, causes the one or more processors to further perform at least:

- receiving or generating an additional one or more digital pictures;
- determining the second one or more instruction sets for operating the third device based on at least partial match between the additional one or more digital pictures and the second one or more digital pictures;
- and
- at least in response to the determining the second one or more instruction sets for operating the third device, causing a fourth device to perform one or more operations defined by the second one or more instruction sets for operating the third device.

96. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that the Grotmol-Hickman combination renders claim 27 obvious. See EX1003, ¶¶358-364.

97. PO's arguments challenging the obviousness of claim 27 suffer from several of the same deficiencies as those addressed above with respect to claims 10, 11, and 25.

98. First, PO's arguments based on Hickman's alleged "failure to disclose and teaching away from ... the claimed correlation of the claimed digital picture with the claimed instruction set" are both incorrect and irrelevant for the reasons discussed above. *Above*, Section IV.A.

99. Second, PO's arguments that it would not be obvious to implement a

“knowledgebase that includes mixed knowledge” from multiple devices lacks merit for at least the reasons discussed above with respect to claim 11. Additionally, Hickman describes a cloud-based, shared robot knowledgebase that would improve the ability of Grotmol’s robots to share knowledge and to operate based on mixed knowledge learned across multiple devices. EX1003, ¶¶358-364.

100. Third, PO’s arguments that Grotmol-Hickman combination does not provide the claimed determination of instruction sets for operating a third device based on “at least partial match” between digital pictures and, in response, causing a fourth device to perform operations defined by those instructions sets also fail based on PO’s erroneous assumptions that the combination lacks a mixed knowledgebase. As discussed above (claim 11) and in my Prior Declaration, the Grotmol-Hickman combination discloses a shared knowledgebase in the cloud that includes trained configurations including correlated digital pictures and instructions sets learned on multiple devices (e.g., including a first device and a third device). EX1003, ¶¶358-364. The trained configurations can be provided to any other robot, including a fourth device. *Id.* When a match is determined between pictures newly acquired by the fourth device and pictures correlated with the instruction sets learned on the third device as described in Grotmol, the fourth device would perform operations defined by the correlated instruction sets. EX1007, 26:14-26, 26:57-27:6, 26:38-56, 26:66-27:27, Abstract, FIG. 4; EX1005, 3:35-39, 39:23-26, 43:5-11.

**V. GROUND 4: THE GROTMOL-HOFFMAN COMBINATION
RENDERS OBVIOUS CLAIM 30**

**A. Combining Grotmol and Hoffman Would Have Been Obvious to a
POSITA**

101. PO and Dr. Saber have not refuted the analysis in my Prior Declaration that shows a POSITA would have found it obvious to combine Grotmol with Hoffman in a manner that would have provided claim 30 of the '583 patent. EX1003, ¶¶365-371. For example, PO argues that the combination would not have been obvious because “Grotmol discloses a system for autonomous robot operation, whereas Hoffman is a system for user-controlled (i.e., manual) robot operation.” POR, 36 (emphasis in original). But this is an over-simplification of both Grotmol’s and Hoffman’s systems and demonstrably false. Hoffman explains that “the user 10 can select different levels of human control over the robot 200, ranging from a teleoperation mode, in which the user 10 directly controls the ... robot 200, to autonomous operation[]” EX1009, 11:31-36. Grotmol similarly provides different operating modes that involve varying degrees of user input for both training and autonomous operation. EX1005, 33:15-24, 21:64-23:29, Table 3.

102. I also do not agree with PO’s assertion that the record lacks explanation as to “how the map in Hoffman would be used with the digital pictures claimed in Grotmol.” As I explained in my Prior Declaration, the combination uses Hoffman’s augmented map generation with Grotmol’s autonomous robotic system in order to

provide more detailed information about the robot's surroundings to the user.
EX1003, ¶¶365-371.

B. Claim 30

103. Claim 30 of the '583 patent recites the following language:

30. The system of claim 29, wherein the detecting the one or more objects in the first device's surrounding or the second device's surrounding includes detecting at least one or more locations of the one or more objects, and wherein the new one or more digital pictures include a top-down view of the one or more representations of the one or more objects.

104. Based on my review, PO and Dr. Saber have not refuted the analysis from my Prior Declaration that the Grotmol-Hoffman combination renders claim 30 obvious. *See* EX1003, ¶¶372-375.

105. PO begins its argument for claim 30 by addressing base claim 29. But these arguments fail for the reasons addressed above. *Above*, Section II.J (Claim 29). I also disagree with PO's arguments directed to the additional limitations in claim 30. PO asserts that "Grotmol only discloses detecting objects but does not disclose detecting positions or locations of those objects. But Grotmol teaches that "object ... location" is detected as part of the context for a task and that objects are located as part of a fetch task/behavior that includes "find[ing] an object, grasp[ing] it, and bring[ing] it back to the base." EX1005, 10:23-34, 24:31-25:5; Pet., 42-43. Further, in the combination, and consistent with my explanation in my Prior


Declaration, the Grotmol-Hoffman system would predictably generate a map like that described in Hoffman in the context of Grotmol's process as described above with respect to claim 29. The rendering of the augmented map (e.g., the new digital picture) is dependent on the robot's location as taught in Hoffman, and in the combination with Grotmol, the robot's location would in turn depend on actions or trajectories taken by the robot in response to previously detected objects as taught in Grotmol (and as I described further above with respect to claim 29). EX1009, 18:38-40 ("The map 184 provides the current location of the robot 200"), 18:46-19:6, FIGS. 6A-6F; *supra*, Claim 29. Further, as I described above in analyzing claim 29, generating the new digital picture (e.g., the map) includes generating depictions (e.g., representations) of walls, obstacles, and other objects detected in a vicinity of the robot. *Id.*

VI. CONCLUSION

106. After reviewing the PORs and Dr. Saber's declaration, I maintain my opinion from my Prior Declaration that the respective claims of the '583 patent challenged in each of the IPRs are obvious.

107. 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 15th day of July, 2024.

A handwritten signature in black ink, appearing to read 'Jason Janet', is written over a horizontal line.

Jason Janet