

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

Case Nos. IPR2023-01172, -01173
U.S. Patent No. 11,055,583

Before Honorable Administrative Patent Judges
Barbara A. Parvis, Robert J. Weinschenk, and Russell E. Cass

Table of Contents (pt. 1)

Overview of Anticipation and Obviousness Issues

Overview of Challenged Claims	5
Anticipation Challenges (IPR2023-01172)	10
Obviousness Challenges (IPR2023-00173)	56
Grotmol in view of POSITA's Knowledge	57
Grotmol in view of Hickman	69
Grotmol in view of Hoffman	78

Table of Contents (pt. 2)

Anticipation Issues (IPR2023-01172)

Claim 7	12
Claim 9	16
Claim 14	20
Claim 15	24
Claim 16	30
Claim 17	34
Claim 18	36
Claim 23	41
Claim 29	47
Claim 31	52

Table of Contents (pt. 3)

Obviousness Issues (IPR2023-01173)

Grotmol in view of POSITA's Knowledge	57
Claim 10	58
Claim 11	63
Grotmol in view of Hickman	69
Overview of Hickman and Combination	70
Claim 25	74
Claim 27	76
Grotmol in view of Hoffman	78
Overview of Hoffman and Combination	79
Claim 30	81

Overview of Challenged Claims

Challenged Claims of the '583 Patent

- **Petition 1 (IPR2023-01172)**
 - **§ 102** – Grotmol anticipates claims 1-9, 11-24, 26, 28-29, and 31-34
- **Petition 2 (IPR2023-01173)**
 - **§ 103** – Grotmol in view of the knowledge of a POSITA renders obvious claims 10 and 11
 - **§ 103** – Grotmol in view of Hickman renders obvious claims 25 and 27
 - **§ 103** – Grotmol in view of Hoffman renders obvious claim 30
- Patent Owner did not contest Petitioner's challenges to independent claims 1, 33, or 34, or dependent claims 2-6, 8, 11 (§102), 12-13, 19-22, 24, 26, 28, or 32 during trial.

Overview of the '583 Patent

(12) United States Patent Cosic		(10) Patent No.: US 11,055,583 B1	
		(45) Date of Patent: *Jul. 6, 2021	
(54) MACHINE LEARNING FOR COMPUTING ENABLED SYSTEMS AND/OR DEVICES	(56) References Cited U.S. PATENT DOCUMENTS		
(71) Applicant: Jasmin Cosic , Miami, FL (US)	4,370,707 A 1/1983 Phillips et al. 364/200		
(72) Inventor: Jasmin Cosic , Miami, FL (US)	4,730,315 A 3/1988 Saito et al. 371/19		
(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 67 days. This patent is subject to a terminal dis- claimer.	4,860,203 A 8/1989 Corrigan et al. 364/300		
	5,560,011 A 9/1996 Uyama 395/700		
	5,602,982 A 2/1997 Judd et al. 395/326		
	6,026,234 A 2/2000 Hanson et al. 395/704		
	6,088,731 A 7/2000 Kiraly et al. 709/229		
(21) Appl. No.: 16/584,736	6,106,299 A 8/2000 Ackermann et al. 434/81		
(22) Filed: Sep. 26, 2019	6,126,330 A 10/2000 Knight 395/704		
Related U.S. Application Data		11/2001 Angel et al. 717/4	
(63) Continuation of application No. 15/822,150, filed on Nov. 26, 2017, now Pat. No. 10,474,934.	6,408,262 B1 6/2002 Leerberg et al. 703/2		
(51) Int. Cl. <i>G06K 9/66</i> (2006.01) <i>G06N 5/02</i> (2006.01) <i>G06K 9/62</i> (2006.01) <i>G06T 7/70</i> (2017.01) <i>G06F 9/30</i> (2018.01) <i>G06N 20/00</i> (2019.01) <i>G06F 15/80</i> (2006.01) <i>G06F 15/78</i> (2006.01)	6,643,842 B2 11/2003 Angel et al. 717/130		
	6,728,689 B1 4/2004 Drissi et al. 706/14		
	6,735,632 B1 5/2004 Kiraly et al. 709/229		
	6,754,631 B1 6/2004 Din 704/270		
	(Continued)		
OTHER PUBLICATIONS			
About Event Tracing, retrieved from <URL: http://msdn.microsoft.com/en-us/library/aa363668(d-default&lang=en-us&vs=85).aspx > on Jun. 12, 2014, 2 pages.			
(Continued)			
<i>Primary Examiner</i> — Andrew W Johns			
(52) U.S. Cl. CPC <i>G06K 9/66</i> (2013.01); <i>G06F 9/30076</i> (2013.01); <i>G06K 9/6202</i> (2013.01); <i>G06N</i> <i>20/00</i> (2019.01); <i>G06T 7/70</i> (2017.01); <i>G06F</i> <i>15/78</i> (2013.01); <i>G06F 15/80</i> (2013.01); <i>G06N</i> <i>5/022</i> (2013.01)	(57) ABSTRACT		
(58) Field of Classification Search CPC .. <i>G06K 9/00355</i> ; <i>G06K 9/00389</i> ; <i>G06K 9/66</i> ; <i>G06T 2207/20081</i> ; <i>G06F 3/002</i> ; <i>G06F</i> <i>3/005</i> ; <i>G06F 3/017</i> ; <i>G06N 5/022</i> ; <i>G06N</i> <i>5/025</i> ; <i>G06N 20/00</i>	Aspects of the disclosure generally relate to computing enabled systems and/or devices and may be generally directed to machine learning for computing enabled systems and/or devices. In some aspects, the system captures one or more digital pictures, receives one or more instruction sets, and learns correlations between the captured pictures and the received instruction sets.		
See application file for complete search history.			
34 Claims, 41 Drawing Sheets			




FIG. 31

EX1001, FIG. 36 (cited in 1173-Pap. 2 (Pet.), p. 6)

Independent Claim 1 + Dependent Claim 4

Claim 1

1. A system comprising:
one or more processors; and
one or more non-transitory machine readable media storing machine readable code that, when executed by the one or more processors, causes the one or more processors to perform at least:
receiving or generating a first one or more digital pictures, wherein the first one or more digital pictures depict at least a portion of a first device's surrounding;
receiving or generating a first one or more instruction sets for operating the first device; and
learning the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device.

Claim 4

4. The system of claim 1, 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 a new one or more digital pictures;
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 first one or more digital pictures; and
at least in response to the determining, causing the first device or a second device to perform one or more operations defined by the first one or more instruction sets for operating the first device.

Note: Claims 5-32 each depends from both independent claim 1 and dependent claim 4.

Independent Claims 33 and 34

Claim 33

33. A method implemented using a computing system that includes one or more processors, the method comprising:
receiving or generating a first one or more digital pictures, wherein the first one or more digital pictures depict at least a portion of a first device's surrounding;
receiving or generating a first one or more instruction sets for operating the first device; and
learning the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device.

Claim 34


34. A system comprising:
means for receiving or generating a first one or more digital pictures, wherein the first one or more digital pictures depict at least a portion of a first device's surrounding;
means for receiving or generating a first one or more instruction sets for operating the first device; and
means for learning the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device.

Note: Patent Owner did not contest Petitioner's challenges to claims 33 or 34 during trial.

Anticipation Challenges

IPR2023-01172

Overview of Grotmol



Grotmol, FIG. 1 (cited in 1172-Pap. 2 (Pet.), 56)

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

Grotmol, FIG. 1 (cited in 1172-Pap. 2 (Pet.), 3)

Claim 7

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.**

Grotmol Discloses Detecting Objects in Digital Pictures




FIG. 10A

Grotmol, FIG. 10A (annotated in 1172-Pap. 33 (Reply), p. 2)

Grotmol

As used herein the term feature may be used to describe one or more integer or floating point values characterizing the input, e.g., the presence or absence of an edge, corner, shape, texture, color, object, at particular locations in the image, values of pixels in an image, patches of color texture, brightness in the image, and/or in the image as a whole; properties of the mentioned features, such as size, orientation, intensity, predominance with respect to the surround, of an edge, corner, shape, texture, color, object; the position of one of the features in the image or the relative position of two or more of the above mentioned features; changes in

Grotmol, 28:19-29 (cited in 1172-Pap. 2 (Pet.), 28-29; 1172-Pap. 33 (Reply), 3-4)

Grotmol Discloses Detecting Objects in Digital Pictures

PO's declarant, Dr. Saber, conceded that Grotmol's feature extraction includes object detection

Dr. Saber:

Q So he might not refer to it as a feature, but what he's referring to here is detecting an object in the image?

A Well, he's referring to detecting an object. But, you know, the devil is in the details. There are many ways of detecting objects. So the devil is in the details. So he's calling his term "feature" and he's saying, well, that's a corner, edge. It's very -- kind of a general umbrella.

EX1011, 28:13-22 (cited in 1172-Pap.33 (Reply), p. 4)

Grotmol Discloses Comparing Features of Digital Pictures, Including Comparing Detected Objects

Grotmol:

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

Grotmol, 2:29-37 (cited in 1172-Pap.33 (Pet.), p. 29)

Dr. Janet:

“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.

EX1010, ¶11 (cited in 1172-Pap.33 (Reply), p. 2)

Claim 9

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**.

Grotmol Discloses a “Knowledgebase”

Grotmol’s “Knowledgebase” Includes “Training Sets” that Correlate Digital Pictures with Instruction Sets

Dr. Janet:

244. Although Grotmol does not explicitly use the term “knowledgebase⁸,” a POSITA would have understood Grotmol to disclose the claimed “knowledgebase.” Specifically Grotmol further discloses a memory includes the stored learned correlations:

EX1003, ¶244 (cited in 1172-Pap.2 (Pet.), pp. 30-31)

⁸ I understand that the ’583 patent states that a “Knowledgebase 530 may be or include any data structure or arrangement capable of storing the knowledge of a device’s operation in various visual surroundings.” ’583 patent, 95:20-23.

EX1003, ¶244 FN8 (cited in 1172-Pap.2 (Pet.), pp. 30-31)


Grotmol:

An apparatus and methods for training and/or operating a robotic device to follow a trajectory. A robotic vehicle may utilize a camera and stores the sequence of images of a visual scene seen when following a trajectory during training in an ordered buffer. Motor commands associated with a given image may be stored. During autonomous operation,

Grotmol, Abstract (cited in 1172-Pap.33 (Pet.), p. 30)

Grotmol Discloses a First Device Trained in Multiple “Learning Processes” by a First User

Grotmol:



Grotmol, FIG. 10D (annotated in 1172-Pap. 33 (Reply), p. 6)

FIG. 10D

Dr. Janet:

245. Grotmol further discloses that multiple training processes can be conducted to learn multiple tasks. Ex. 1005, 24:17-67 (describing training of “simple tasks” and “complex tasks,” including behaviors such as “fetch” and “grasping.”), 26:1-20 (describing that “[t]he system may be configured to learn the second behavior from the user.”). Therefore, a POSITA would have understood that Grotmol discloses a memory that stores a knowledgebase including multiple correlations (training sets), where correlations may be learned in separate processes that include operation of a first device by a user.

EX1003, ¶245 (cited in 1172-Pap. 2 (Pet.), p. 31)

PO's Argument that Grotmol Discloses Only a Single "Learning Process" is Incorrect

Patent Owner's Response:

Grotmol fails to disclose the claimed "second learning process." Ex. 2007, ¶¶78-81. Grotmol only discloses *one* user operating *one* device in *one* learning process, and therefore fails to disclose any second learning process, let alone the

1172-Pap. 28 (POR), p. 15

Dr. Saber:

Q So a temporal separation might be one thing that you would look for to determine that the second learning process is separate from the first learning process?

A Yeah, I mean, generally speaking because you have the same user and you have obviously the same device here. You have the same user, you have the same device. There's a first learning process that's taken place and at seconds later or minutes or whatever, there's a second learning process that's taken place.

EX1011, 50:11-21 (cited in 1172-Pap. 33 (Reply), p. 8)

Grotmol:

After is user is done training a particular behavior, he may create another module and activate it. Such creation is depicted in the figure below, in accordance with one or more implementations.

The system may be configured to learn the second behavior from the user. Figure above shows that the user may select a different set of input features and a different set of outputs for "Behavior 2". The user may select a different type of the adaptive controller ("2"). For example, the user may control the robot using a gamepad (or other input device used to provide "user commands") to turn and grasp the target ("Behavior 2") using input from target tracker ("FE 1") and state of the gripper ("FE 2"). "Adaptive controller"

Grotmol, 1:14 (cited in 1172-Pap. 33 (Reply), pp. 6-7)

Claim 14

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.

Grotmol Discloses Applying Instruction Sets in a “Trained Configuration” from a First Device to a Second Device

Grotmol: (e.g., a smartphone)). Training of the robot’s controller may be based on the user observing robot’s actions and sending one or more target control commands to the robot via the training handset. The trained controller of the robot may comprise a trained configuration configured to enable autonomous operation (e.g., without user input) by the robotic device. The trained configuration may be stored. A saved configuration may be loaded into the robot being trained thereby providing one or more trained behaviors to the robot. In some implementations, the trained configuration may be loaded to one or more other robots in order to provide learned behaviors. Subsequent to loading of the saved configuration, the controller learning process may match process configuration being present during saving of the configuration.

Grotmol, 33:21-35 (cited in 1172-Pap. 2 (Pet.), p. 34)

Dr. Janet: 254. Grotmol further discloses applying instruction sets to a second device. A POSITA would have understood that, during operation, the second device would determine the similarity measure between the new image (new digital picture depicting at least a portion of the second device’s surrounding) and an image from the training set loaded to the second device (stored first digital picture). After a

EX1003, ¶254 (cited in 1172-Pap. 2 (Pet.), p. 35)

Grotmol Discloses “Modifying” Instruction Sets

Exemplary implementations of the methodology described herein may be applicable to controlling trajectory of a robotic device due to (i) position mismatch (e.g., the robot being located physically to the left of a target location); and/or (ii) due to orientation mismatch (e.g., the robot being in the same physical location while oriented towards the left of the target trajectory). 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.

Grotmol, 42:46-58 (cited in 1172-Pap. 2 (Pet.), p. 36)

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.

Grotmol, 43:4-11 (cited in 1172-Pap. 33 (Reply), p. 9)




FIG. 4A

Grotmol, 4A (annotated in 1172-Pap. 33 (Reply), p. 10)

Patent Owner's Witness, Dr. Saber, Admits that Grotmol's Combiner "Modifies" Motor Commands

Dr. Saber:

Q Do any of the transfer functions that Grotmol lists here modify the motor commands in signal 418?

MR. McCOMMAS: Objection. Form.

A Well, there isn't any of the transfer functions. There is -- that transfer function may comprise one or more of, so obviously the addition, multiplication, union, logical "and," and a logical "or" would modify -- would take X as described by Grotmol and which X comprise, you know, 408 and 418 and would produce H of X, which would, you know, be a modification because otherwise it would just say it will pass it along and you don't need a combiner.

EX1011, 60:18-61:9 (cited in 1172-Pap. 33 (Reply), p. 11)

Claim 15

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, ... and ...

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, ... and ...

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


Grotmol Discloses “Modifying” Instruction Sets

Exemplary implementations of the methodology described herein may be applicable to controlling trajectory of a robotic device due to (i) position mismatch (e.g., the robot being located physically to the left of a target location); and/or (ii) due to orientation mismatch (e.g., the robot being in the same physical location while oriented towards the left of the target trajectory). 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.

Grotmol, 42:46-58 (cited in 1172-Pap. 2 (Pet.), pp. 36-37)

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.

Grotmol, 43:4-11 (cited in 1172-Pap. 33 (Reply), pp. 9, 12, 14)



Grotmol, 4A (annotated in 1172-Pap. 33 (Reply), p. 10)

Grotmol Determines Modified Instruction Sets to Perform Based on Matching A New Digital Picture to a Stored Digital Picture

Grotmol: 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.

Grotmol, 43:4-11 (cited in 1172-Pap. 2 (Pet.), pp. 37-38)

Dr. Janet: Grotmol teaches that each stored digital image is correlated with a motor command. When the robotic vehicle determines that the best matching image is shifted or mismatched and subsequently selects an instruction set, it issues a motor command that is different from the first instruction set. Said differently, the robotic vehicle is selecting a modified instruction set.

EX1003, ¶264 (cited in 1172-Pap. 2 (Pet.), p. 38)

PO Errs in Arguing that Claim 15 Requires the First Instruction Sets to be Modified Before the Determination Step Begins

PO's Claim Interpretation is Inconsistent with Preferred Embodiments of the '583 Patent

Dr. Janet:

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:

EX1010, ¶134 (cited in 1172-Pap. 33 (Reply), pp. 14-15)

'583 Patent:

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. In further embodiments, Con-

EX1001, 89:57-62 (cited in 1172-Pap. 33 (Reply), pp. 14-15)

PO Errs in Arguing that Claim 15 Requires the First Instruction Sets to be Modified Before the Determination Step Begins

PO's Claim Interpretation that the Modified Instruction Set Must be the Correlated Instruction Set Improperly Imports a Non-Recited "Learning" Step

Claim 15:

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:

No recited "learning" of modified instructions

Claim 18:

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

"Learning" Step Recited

Grotmol Discloses Claim 15 Even Under PO's Erroneous Interpretation

Grotmol determines modified instruction sets after the modification step, in some examples

Grotmol:

learning process, and/or a combination thereof. The corrector 412, the predictor 422 and the combiner 414 may cooperate to produce a control signal 420 for the robotic platform 410. In one or more implementations, the control signal 420 may comprise one or more motor commands (e.g., pan camera to the right, turn right wheel forward), sensor acquisition commands (e.g., use high resolution camera mode), and/or other commands.

In some implementations, the predictor 422 and the combiner 414 components may be configured to operate a plurality of robotic platforms. The control signal 420 may be adapted by a decoder component 424 in accordance with a specific implementation of a given platform 430. In one or more implementations of robotic vehicle control, the adaptation by the decoder 424 may comprise translating binary signal representation 420 into one or more formats (e.g., pulse code modulation) that may be utilized by given robotic vehicle. U.S. patent application Ser. No. 14/244,890 entitled "LEARNING APPARATUS AND METHODS FOR CONTROL OF ROBOTIC DEVICES", filed Apr. 3, 2014 describes some implementations of control signal conversion.

In some implementations of the decoder 424 corresponding to the analog control and/or analog corrector 412 implementations, the decoder may be further configured to rescale the drive and/or steering signals to a range appropriate for the motors and/or actuators of the platform 430.

"modify"

Dr. Janet:

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.

EX1010, ¶34 (cited in 1172-Pap. 33 (Reply), pp. 14-15)

"determine"

Grotmol, 15:26-52 (cited in 1172-Pap. 33 (Reply), pp. 14-15)

Claim 16

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[or] 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, ... (iii) at least partial match between the new one or more digital pictures and the modified the first one or more digital pictures, ... [or] (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, ...**

Grotmol Discloses Pre-Processing Operations that Modify Stored Digital Pictures

Grotmol:

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. Motor commands issued by a trainer corresponding to time instances when the images are acquired may be stored. Additional sensory information (e.g., vehicle motion information, ambient environment information, vehicle operational parameters) corresponding to time instances when the images are acquired may be stored.

Grotmol, 39:13-22 (cited in 1172-Pap. 2 (Pet.), p. 40)

Image shift parameters determined from the image correlation operation may be used when determining which image(s) from the match search set may be considered as a match to the given image. In some implementations, the given image may be shifted by amount determined from the image matching operation. By way of an illustration, image 2010 may be shifted to the right by amount depicted by arrow 2015 in FIG. 20A. The shifted image may be trimmed from the edge towards which the shift is performed (e.g., right edge of the image 2010 in FIG. 20A). The matching image (e.g., 2000 in FIG. 20A) may be trimmed on the left by the amount of the shift.

Grotmol, 40:59-61 (cited in 1172-Pap. 33 (Reply), p. 17)

Dr. Janet:

POSITA would have understood that the pre-processing, normalizing, and equalizing of digital images comprises modifying images. Accordingly, stored digital pictures are stored with said modifications. Thus, when Grotmol approaches its best matching step, in my opinion, the stored image it is using for said comparison is a modified digital picture.

EX1003, ¶269 (cited in 1172-Pap. 2 (Pet.), p. 40)

Grotmol Discloses Using the Modified and Stored Digital Pictures in a Matching Process

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

Grotmol, Abstract (cited in 1172-Pap. 2 (Pet.), p. 20)

Dr. Janet: 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.

EX1010, ¶38 (cited in 1172-Pap. 33 (Reply), p. 16)

Grotmol Also Discloses Pre-Processing Operations That Modify “New” Digital Pictures Used in a Matching Process

Grotmol:

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

Grotmol, 4:18-22 (cited in 1172-Pap. 33 (Reply), pp. 17-18)

Dr. Janet:

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.

EX1010, ¶40 (cited in 1172-Pap. 33 (Reply), p. 18)

Claim 17

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, ... 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, ... 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

....

Grotmol Discloses Modifying Digital Pictures, and Learning and Matching the Modified Digital Pictures

Grotmol:

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. Motor commands issued by a trainer corresponding to time instances when the images are acquired may be stored. Additional sensory information (e.g., vehicle motion information, ambient environment information, vehicle operational parameters) corresponding to time instances when the images are acquired may be stored.

← **Modify**
← **Learn**

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

← **Match**

Grotmol, 39:13-28 (cited in 1172-Pap. 2 (Pet.), p. 8, 41)

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

Grotmol, 33:59-62 (cited in 1172-Pap. 2 (Pet.), p. 41)

Claim 18

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, ... 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, ... 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, ... 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, ...

Grotmol Discloses “Modifying” Instruction Sets and “Learning” the Modified Instruction Sets

Grotmol:

TABLE 3

Mode	Action of the robot in presence of training input	Action of the robot in absence of training input
Control	Executes the action in accordance with the training input; Does not learn the association	Idle
Listen	Executes the action in accordance with the training input; Learns the association	Idle
Override	Executes the action in accordance with the training input; <u>Learns the association</u>	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. <u>Learn the resulting association</u>	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

Grotmol, 22:27-46 (cited in 1172-Pap. 33 (Reply), p. 20)

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

In the correct mode, the robot may integrate what the teacher commands with what the robot already knows. In some implementations, the robot may sum the teacher's command with what the robot already knows to get a final motor command. The teacher's correction may operate in this case as a deviation from the predicted command determined by the robot.

Grotmol, 22:27-46 (cited in 1172-Pap. 2 (Pet.), p. 44 and 1172-Pap. 33 (Reply), pp. 20-22)

Grotmol Discloses “Determining” and “Causing” Performance of Operations Defined by the Modified Instruction Sets

Grotmol:

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.

Grotmol, 43:5-11 (cited in 1172-Pap. 2 (pet.), p. 46)

Dr. Janet:

39, 33:59-62, Abstract, 43:5-11, 42:46-58, 42:37-40. A POSITA would have understood that, following the robot modifying its instruction set, the next time a robot encounters the same situation and uses the modified instruction set the robot is causing the robotic vehicle’s processors to perform the modified instruction set.

EX1003, ¶290 (cited in 1172-Pap. 2 (Pet.), p. 47)

PO's Argument that Grotmol's "Override Mode" Does Not Involve "Modifying" Instruction Sets is Wrong

Patent Owner:

command. Ex. 2007, ¶¶106-107. *Grotmol is replacing an original, whole, unmodified command with another (i.e., overridden), whole, unmodified command.* Replacing an element with another element is fundamentally different than modifying the original element. *Id.* Petitioner has failed to show that Grotmol's

1172-Pap. 28 (POR), pp. 30-31

'583 Patent:

Executing may include modifying an element of a logic circuit with one or more alternate instruction sets, redirecting the logic circuit to one or more alternate instruction sets, replacing the inputs into the logic circuit with one or more alternate inputs or instruction sets, and/or replacing the outputs from the logic circuit with one or more alternate outputs or instruction sets. In further embodiments, a pro-

EX1001, 147:36-42 (cited in 1172-Pap. 33 (Reply), p. 21)

produce two output values. Modifying execution and/or functionality of Logic Circuit 250 can be implemented by replacing its input values with anticipatory input values (i.e. anticipatory Instruction Sets 526, etc.). Artificial Intelligence

EX1001, 140:15-18 (cited in 1172-Pap. 33 (Reply), p. 21)

state of its target. Ptrace can be used to modify a running application such as modifying an application with alternate instruction sets (i.e. anticipatory Instruction Sets 526, etc.).

EX1001, 136:9-11 (cited in 1172-Pap. 33 (Reply), p. 21)

Grotmol's "Correct Mode" Modifies Instruction Sets

Grotmol:

In the correct mode, the robot may integrate what the teacher commands with what the robot already knows. In some implementations, the robot may sum the teacher's command with what the robot already knows to get a final motor command. The teacher's correction may operate in this case as a deviation from the predicted command determined by the robot.

Grotmol, 23:23-29 (cited in 1172-Pap. 33 (Reply), p. 22)

Claim 23

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.

PO's Construction is Unclear and Unsupported

Patent Owner:

In the context of claim 23 of the '583 patent, including in light of the claims, specification, and drawings, a “connection” should be construed as a “relationship that includes a link.” Ex. 2007, ¶¶53-59. This construction is supported by the

1172-Pap. 28 (POR), p. 7

Dr. Janet:

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

EX1010, ¶¶54-55 (cited in 1172-Pap. 33 (Reply), pp. 23-24)

Plain and Ordinary Meaning of “Connection” is Not Limited to a “Link”

EX. 2013

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

EX2013, 481 (cited in 1172-Pap. 33 (Reply), pp. 23-24)

EX. 2012

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

EX2012, 399-400 (cited in 1172-Pap. 33 (Reply), pp. 23-24)

Examples and Embodiments of “Connection” in the ’583 Patent are Non-Limiting

Dr. Janet:

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.

EX1010, ¶¶53 (cited in 1172-Pap. 33 (Reply), pp. 23-24)

Grotmol Discloses the Claimed “Connection[s]”

Grotmol:

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

Grotmol, 3:35-39 (cited in 1172-Pap. 2 (Pet.), p. 21)

Dr. Janet:

307. In my opinion, Grotmol’s storing an association between a digital picture and an action to be a “connection”⁹ between the two. At least the data structures described in Grotmol linking digital images and instruction sets together in memory, for example, would comprise “one or more connections” as recited in the claim:


EX1003, ¶307 (cited in 1172-Pap. 2 (Pet.), p. 53)

Even under PO's Construction, Grotmol Discloses the Claimed "Connection[s]"

Dr. Janet:

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,

EX1010, ¶59 (cited in 1172-Pap. 33 (Reply), pp. 24-25)



Grotmol: "Training Set" Example

EX1010, ¶59 (cited in 1172-Pap. 33 (Reply), p. 25)

Claim 29

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.

PO Errs in Arguing that Claim 29 Requires a Causal Chain of Detecting Objects to Generate Representations of the Detected Objects for New Digital Pictures

Dr. Janet:

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.

EX1010, ¶61 (cited in 1172-Pap. 33 (Reply), p. 25)

Pet. Reply:

Dr. Saber could not even identify support for such a method in the specification of the ’583 patent.³ EX1011, 65:9-67:1.

1172-Pap. 33 (Reply), pp. 25-26

Grotmol Discloses Claim 29 Even Under PO's Interpretation

Objects detected at time t_1 cause robot to turn and generate image including representations of the objects at time t_2

Grotmol:




FIG. 20A

Grotmol, FIG. 20A (cited in 1172-Pap. 2 (Pet.), p. 60)

respectively. The image 2000 may comprise representations one or more objects, e.g., a tree 2002, a rectangle, a ball 2008, and/or other objects. During motion of the vehicle 160 relative position and/or orientation of the camera and the visual scene may change. By way of an illustration, the vehicle 160 may execute a right turn. The image 2010 may be obtained subsequent to the turn execution. Representation of objects in the image 2010 may be shifted relative representations of the same objects in the frame obtained prior to the turn execution. Additional object representations (e.g.,

Grotmol, 38:56-65 (cited in 1172-Pap. 2 (Pet.), p. 60)

As used herein the term feature may be used to describe one or more integer or floating point values characterizing the input, e.g., the presence or absence of an edge, corner, shape, texture, color, object, at particular locations in the image, values of pixels in an image, patches of color texture, brightness in the image, and/or in the image as a whole; properties of the mentioned features, such as size, orientation, intensity, predominance with respect to the surround, of an edge, corner, shape, texture, color, object; the position of one of the features in the image or the relative position of two or more of the above mentioned features; changes in

Grotmol, 28:19-29 (cited in 1172-Pap. 33 (Reply), p. 26)

Grotmol Discloses Claim 29 Even Under PO's Interpretation

Objects detected at time t_1 cause robot to turn and generate image including representations of the objects at time t_2

Grotmol:

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. Motor commands issued by a trainer corresponding to time instances when the images are acquired may be stored. Additional sensory information (e.g., vehicle motion information, ambient environment information, vehicle operational parameters) corresponding to time instances when the images are acquired may be stored.

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

Grotmol, 39:13-28 (cited in 1172-Pap. 33 (Reply), p. 3)

Grotmol Discloses Claim 29 Even Under PO's Interpretation

Objects detected at time t1 cause robot to turn and generate image including representations of the objects at time t2

Dr. Janet:

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.

EX1010, ¶62 (cited in 1172-Pap. 33 (Reply), pp. 26-27)

Claim 31

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.

Grotmol's Reloaded Configuration Comprises Instruction Sets Generated as Representations of Instruction Sets in an Earlier Configuration

Petition:

Grotmol discloses that the system may be configured to remove a learned behavior. Ex. 1005, 34:18-21 (“The BrainOS software platform may be configured to enable users to selectively remove a learned behavior (and/or a portion thereof) via an undo and/or time machine operation.”). In implementation, the system reloads the previous configuration to unlearn a behavior, specifically, “[b]y loading at time t_1 the configuration saved at time $t_0 < t_1$, the robot controller effectively ‘forgets’ what it learned in time interval $t_0 < t < t_1$.” *Id.*, 34:25-27. *See also id.*, 34:43-45 (“A user may select a given bookmark in order to restore (undo) the learning configuration to the state corresponding to time of the bookmark.”). Decl. ¶ 332. When the system then uses the reloaded configuration, the system is generating the first instruction set as a representation of a second instruction set. Decl. ¶ 332.

1172-Pap. 2 (Pet.), pp. 71-72

Plain and Ordinary Meaning of “Representation” Supports Petitioner’s Mapping

Dr. Janet:

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

EX1010, ¶64 (cited by 1172-Pap. 33 (Reply), p. 28

New Oxford Am. Dictionary:

rep·re·sen·ta·tion /repriːzenˈtāʃən, -zən/ ▶ *n.*
1 the action of speaking or acting on behalf of someone or the state of being so represented: *asylum-seekers should be guaranteed good legal advice and representation.*
2 the description or portrayal of someone or something in a particular way or as being of a certain nature: *the representation of women in newspapers.* ■ the depiction of someone or something in a picture or other work of art: *Picasso is striving for some absolute representation of reality.* ■ a thing, esp. a picture or model, that depicts a likeness or reproduction of someone or something: *a striking representation of a vase of flowers.* ■ (in some theories of perception) a mental state or concept regarded as corresponding to a thing perceived.

EX1018, 1481 (cited by 1172-Pap. 33 (Reply), p. 28

Webster’s New World College Dictionary:

rep·re·sen·ta·tion (rep’ri zen tā’shan) *n.* [ME < MFr < L. *representatio*] 1 a representing or being represented (in various senses); specif., the fact of representing or being represented in a legislative assembly 2 legislative representatives, collectively 3 a likeness, image, picture, etc. 4 [often *pl.*] a description, account, or statement of facts, allegations, or arguments, esp. one intended to influence action, persuade hearers, make protest, etc. 5 the production or performance of a play, etc. 6 *Law* a statement or implication of fact, oral or written, as made by one party to induce another to enter into a contract.

EX1019, 1216-1217 (cited by 1172-Pap. 33 (Reply), p. 28

The '583 Patent Does Not Limit the Breadth of “Representation” Beyond the Plain and Ordinary Meaning

Dr. Janet:

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.

EX1010, ¶64 (cited by 1172-Pap. 33 (Reply), p. 28

Petition:

⁸ The '583 patent specification does not define or even contain the term “representations of a [first or second or any] instruction sets.” The most similar term used in the specification is “representations of that particular instruction set” which is described in the context of obtaining instruction sets implemented through assembly language. *See* Ex. 1001, 79:33-36 (“Syntax, addressing modes, operands, and/or other elements of an assembly language instruction set may translate directly into numeric (i.e. binary, etc.) representations of that particular instruction set.”).

1172-Pap. 2 (Pet.), p. 61, FN8

Obviousness Challenges

IPR2023-01173

DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Obviousness based on Grotmol in view of the Knowledge of a POSITA

IPR2023-01173


DEMONSTRATIVE EXHIBIT – NOT EVIDENCE

Claim 10

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**.

Grotmol Discloses a Knowledgebase that Includes Correlations Learned in a “Second Learning Process”

Grotmol:



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

Grotmol, 3:35-39 (cited in 1173-Pap. 2 (Pet.), p. 24)

FIG. 10D Grotmol, FIG. 10D (cited in 1173-Pap. 32 (Reply), p. 3)

After is user is done training a particular behavior, he may create another module and activate it. Such creation is depicted in the figure below, in accordance with one or more implementations.

The system may be configured to learn the second behavior from the user. Figure above shows that the user may

Grotmol, 26:1-6 (cited in 1173-Pap. 32 (Reply), p. 3)

(“base approach”). While 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 (e.g., less computationally intensive). The user may have an

Grotmol, 24:57-61 (cited in 1173-Pap. 32 (Reply), p. 3)

Claim 10's Recitation of a "Second User" is not Entitled to Patentable Weight

Reply:

² Petitioner also notes that the language in claim 10 directed to a "second user" different from the first does not further define any structure or function of the claimed "system" and is not entitled to patentable weight. Even if it were given weight, it would have been obvious based on the teachings of Grotmol in view of the knowledge of a POSITA.

1173-Pap. 32 (Reply), p. 7 FN2

Grotmol Renders Obvious a “Second User” Operating the First Device in a Second Learning Process

Dr. Janet:

341. Grotmol does not limit the number of users that can train the multiple behaviors for the robot, thus in my opinion it would have been obvious that different users can train different behaviors. Given that Grotmol describes the training of both simple and complex tasks, and in order to take advantage of the varying skills and knowledge of trainers, it would have been obvious that different trainers could be used to train behaviors involving different levels of complexity in order to take advantage of the varying skills and knowledge of trainers. In my opinion, there is no

EX1010, ¶341 (cited by Pap. 2 (Pet.), pp. 24-25)

Grotmol Renders Obvious a “Second User” Operating the First Device in a Second Learning Process

Dr. Saber:

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

EX1016, 166:3-18 (cited by 1173-Pap. 32 (Reply), pp. 7-8).

Claim 11

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**.

PO Did Not Contest Petitioner's § 102 Challenge to Claim 11 in the -1172 Proceeding, Only Contested § 103 Challenge in -1173 Proceeding

Petitioner's Reply:

II. THE POR FAILS TO REFUTE THE PETITION'S SHOWING THAT GROTMOL ANTICIPATES CLAIMS 1-9, 11-24, 26, 28-29, AND 31-34

As an initial matter, the Petition challenged each of claims 1-9, 11-24, 26, 28-29, and 31-34 of the '583 patent. The POR noticeably addresses only dependent claims 7, 9, 14-18, 23, 29, and 31, leaving the Petition's challenges to claims 1-6, 8, 11-13, 19-22, 24, 26, 28, and 32-34 uncontested. This Reply responds to PO's arguments on claims 7, 9, 14-18, 23, 29, and 31. The remaining claims are unpatentable for the reasons set forth in the Petition. EX1010, ¶8.

1172-Pap. 33 (Reply), p. 1

Grotmol Discloses Sharing of Trained Configurations to Share Learned Behaviors with Other Robots

Grotmol:

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

Grotmol, 33:21-32 (cited in 1173-Pap. 2 (Pet.), p. 26)




FIG. 10B

Grotmol, FIG. 10B (cited in 1173-Pap. 32 (Reply), p. 14)

Dr. Janet:

Moreover, a POSITA would have recognized the advantages of training across multiple robots—e.g., to enable continued training while a robot was undergoing maintenance and increase the amount of training that could be done which would maximize performance.

EX1003, ¶1343 (cited in 1173-Pap. 2 (Pet.), p. 26)

Obvious for the Receiving Robot to Maintain its Locally Learned Knowledge When Receiving a Trained Configuration Shared by Another Robot

PO Wrongly Contends that Grotmol's Receiving Robot Would Include Knowledge Learned only on One Device

POR:

¶¶132-136. As discussed above with respect to claim 10, Grotmol only discloses *one* user operating *one* device in *one* learning process. *Supra* Section VI.A.1; Ex. 2007,

¶¶132-133. The knowledge gained from this process in Grotmol, when loaded into another device, is still knowledge of one user operating one device in one learning process, which is about loading the knowledge previously learned on one robot into

1172-Pap. 27 (POR), p. 17

Dr. Janet:

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.

EX1003, ¶80 (cited in 1173-Pap. 32 (Reply), p. 11)

Obvious for the Receiving Robot to Maintain its Locally Learned Knowledge When Receiving a Trained Configuration Shared by Another Robot

PO Wrongly Contends that Grotmol's Receiving Robot Would Include Knowledge Learned only on One Device

POR: ¶¶132-136. As discussed above with respect to claim 10, Grotmol only discloses *one* user operating *one* device in *one* learning process. *Supra* Section VI.A.1; Ex. 2007, ¶¶132-133. The knowledge gained from this process in Grotmol, when loaded into another device, is still knowledge of one user operating one device in one learning process, which is about loading the knowledge previously learned on one robot into

1172-Pap. 27 (POR), p. 17

Dr. Janet: 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.

EX1010, ¶80 (cited in 1173-Pap. 32 (Reply), p. 11)

Obvious for the Receiving Robot to Maintain its Locally Learned Knowledge When Receiving a Trained Configuration Shared by Another Robot

PO Wrongly Contends that Grotmol's Receiving Robot Would Include Knowledge Learned only on One Device

Dr. Janet:

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.

EX1010, ¶81 (cited in 1173-Pap. 32 (Reply), p. 11)

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

EX1010, ¶83 (cited in 1173-Pap. 32 (Reply), pp. 13-14)

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.”).

EX1010, ¶82 (cited in 1173-Pap. 32 (Reply), p. 12)

Obviousness based on the Grotmol- Hickman Combination

IPR2023-01173

Overview of Hickman

Hickman discloses a cloud-based shared knowledge base.

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

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

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




FIGURE 3

Hickman, FIG. 3 (cited by Pap. 32 (Reply), p. 18)

Overview of Hickman

Hickman discloses a cloud-based shared knowledge base.

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

Hickman, 24:8-25

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

Hickman, 11:31-41.

A POSITA Would Have Combined Grotmol with Hickman

Petition:

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

Pap. 2 (Pet.), p. 28

Hickman:

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

Hickman, 11:31-41 (cited by Pap. 2 (Pet.), p. 28)

Dr. Janet:

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

EX1003, ¶159 (cited by Pap. 2 (Pet.), p. 30)

PO's Non-Combinability Arguments are Baseless

Neither Hickman nor the Grotmol-Hickman combination requires separate knowledge bases for digital pictures and instructions sets

Dr. Janet:

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,

EX1010, ¶87 (cited by Pap. 32 (Reply), p. 16)

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

EX1010, ¶89 (cited by Pap. 32 (Reply), p. 16)

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.

EX1010, ¶88 (cited by Pap. 32 (Reply), p. 16)

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.

EX1010, ¶91 (cited by Pap. 32 (Reply), p. 16)

Claim 25

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 pictures 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.

The Grotmol-Hickman Combination Renders Obvious Claim 25

PO's arguments against the claim 25 challenge rest on similar errors as PO's non-combinability arguments

Dr. Janet:

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.

EX1010, ¶94 (cited by Pap. 32 (Reply), p. 21)

Claim 27

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.**

The Grotmol-Hickman Combination Renders Obvious Claim 27

PO's arguments against the claim 27 challenge rest on similar errors as PO's non-combinability arguments

Dr. Janet:

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.

EX1010, ¶98 (cited by Pap. 32 (Reply), p. 23)

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,

EX1010, ¶100 (cited by Pap. 32 (Reply), p. 23)

Obviousness based on the Grotmol-Hoffman Combination

IPR2023-01173

Overview of Hoffman

(57)

ABSTRACT

A method of operating a robot includes electronically receiving images and augmenting the images by overlaying a representation of the robot on the images. The robot representation includes user-selectable portions. The method includes electronically displaying the augmented images and receiving an indication of a selection of at least one user-selectable portion of the robot representation. The method also includes electronically displaying an intent to command the selected at least one user-selectable portion of the robot representation, receiving an input representative of a user interaction with at least one user-selectable portion, and issuing a command to the robot based on the user interaction.

EX1009, Abstract (cited in 1173-Pap. 2 (Pet.), p. 10)

and/or analysis. Moreover, the remote user 10 may access the received sensor data 502 for review, analysis, and control of the robot 200. In some examples, the cloud service 22 executes an application outputting data 23 that may be communicated to the robot 200. For example, the cloud service 22 may process received image data 502 (captured by the robot 200) and return a processed data set 23 (e.g., a layout map of the environment) to the robot 200 (e.g., to a controller 400).

EX1009, 8:1-8 (cited in 1173-Pap. 2 (Pet.), p. 43)




FIG. 6C




FIG. 6D

EX1009, FIGS. 6C, 6D (cited in 1173-Pap. 2 (Pet.), p. 11)

A POSITA Would Have Combined Grotmol with Hoffman

Dr. Janet:

366. In my opinion, a person of ordinary skill in the art would have been motivated to combine the autonomous robotic system of Grotmol with the handheld device and augmented display of Hoffman. As detailed in VI.A above, Grotmol teaches a training methodology for autonomous navigation by one or more robots. See Section VI.A. A POSITA would have looked to expand upon the functionality of the robots and training methods of Grotmol by using the teachings of Hoffman, which, as discussed above in Section VI.D, discloses controlling a robot remotely that can receive and display images from the robot, including augmented images.

EX1003, ¶366 (cited by Pap. 2 (Pet.), p. 37)

Claim 30

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**.

Grotmol Discloses Detection of Objects and their Locations

Grotmol:




FIG. 20A

Grotmol, FIG. 20A (cited by 1173-Pap. 2 (Pet.), p. 42)

In one or more implementations, object recognition and/or classification may be implemented using a spiking neuron classifier comprising conditionally independent subsets as described in co-owned U.S. patent application Ser. No. 13/756,372 filed Jan. 31, 2013, and entitled “SPIKING NEURON CLASSIFIER APPARATUS AND METHODS USING CONDITIONALLY INDEPENDENT SUBSETS” and/or co-owned U.S. patent application Ser. No. 13/756,382 filed Jan. 31, 2013, and entitled “REDUCED LATENCY SPIKING NEURON CLASSIFIER APPARATUS AND METHODS”, each of the foregoing being incorporated herein by reference in its entirety.

Grotmol, 9:41-52 (cited by 1173-Pap. 2 (Pet.), p. 42)

In some implementations of robotic vehicle navigation, output of the sensor **166** in FIG. 1 may comprise representations of one or more objects (e.g., targets, and/or obstacles). The tasks of the robot may be configured based on a context. In one or more implementations, the context may comprise one or more of robot state (e.g., location or motion information, (position, orientation, speed), platform state or configuration (e.g., manipulator size and/or position), available power and/or other), state of the environment (e.g., object size, location), environmental state (wind, rain), previous state information (e.g., based on historic states of robot motions), and/or other characteristic state information.

Grotmol, 10:23-34 (cited by 1173-Pap. 2 (Pet.), p. 42)

Hoffman Discloses Generating New Digital Pictures Including a Top-Down View of Detected Objects

Dr. Janet:

373. Hoffman describes generating an augmented map of a robot's surroundings (top-down view) that includes the locations of representations of objects, such as walls and obstacles, in the robot's surroundings:

EX1003, ¶373 (cited by 1173-Pap. 2 (Pet.), p. 42)

Ex. 1009, 8:5-8, Fig. 6F. In my opinion, a POSITA reviewing Grotmol would be inclined to look to Hickman to augment Grotmol's disclosed object detection features. Specifically, it would be advantageous to have a full view of the object in relation to the robot in space in order for the robot to optimize its next task.

374. In my opinion, it would have been obvious to combine 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 and to the robot's decision-making system. Such combination would have simply been the use of a known technique to improve a similar device in the same way. A POSITA would have had a reasonable expectation of success in combining the teachings of Grotmol and Hoffman, which describe complementary systems.

EX1003, ¶¶373-74 (cited by 1173-Pap. 2 (Pet.), p. 42)

Hoffman:




FIG. 6F

Hoffman, FIG. 6F (cited by 1173-Pap. 2 (Pet.), p. 43)

Thank You