

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

TESLA, INC.,
Petitioner,

v.

AUTONOMOUS DEVICES, LLC
Patent Owner

U.S. PATENT NOS. 10,102,449 and 11,055,583

DECLARATION OF JASON JANET

I.	Introduction.....	1
II.	Summary of Opinions.....	1
III.	Qualifications and Compensation.....	2
IV.	Relevant Legal Standards	10
A.	Claim Construction	11
B.	Invalidity	13
1.	Anticipation.....	13
2.	Obviousness	14
C.	Person of Ordinary Skill in the Art	17
V.	Background of the Technology	18
VI.	The Prior Art.....	20
A.	U.S. 9,604,359 (“Grotmol”)	20
B.	U.S. 8,639,644 (“Hickman”).....	23
C.	U.S. 9,443,192 (“Cosic 192”)	28
D.	U.S. 9,283,674 (“Hoffman”)	30
VII.	Claim Construction.....	33
A.	“means for receiving or generating”; “means for learning” (claim 34)	33
VIII.	The ’449 Patent.....	34
A.	The Claims	34
B.	Ground 1: Claims 1-9 and 12-20 Are Anticipated by Grotmol	43
1.	Claim 1	44
a.	[1.1] A system for autonomous device operating, the system comprising:	44

b.	[1.2] one or more processor circuits;.....	45
c.	[1.3] a memory that stores at least a first one or more digital pictures correlated with a first one or more instructions sets for operating a first physical device,.....	46
d.	[1.4] wherein the first physical device includes an actuator for moving at least a portion of the first physical device,.....	50
e.	[1.5] 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 physical device are learned in a learning process that includes operating the first physical device at least partially by a user;.....	52
f.	[1.6] an optical camera that captures digital pictures; and.....	54
g.	[1.7] an artificial intelligence unit that:	57
h.	[1.7.1] receives a new one or more digital pictures from the optical camera;	60
i.	[1.7.2] anticipates the first one or more instruction sets for operating the first physical device based on at least partial match between the new one or more digital pictures and the first one or more digital pictures, wherein the anticipates includes at least one of: determining that a number of at least partially matching portions of the new one or more digital pictures and portions of the first one or more digital pictures exceeds a threshold number, or determining that a percentage of at least partially matching portions of the new one or more digital pictures and portions of the first one or more digital pictures exceeds a threshold percentage; and	62
j.	[1.7.3] causes the one or more processor circuits to execute the first one or more instruction sets for operating the first physical device, wherein the causes is performed in response to the anticipates of the artificial intelligence unit, and wherein the first physical device or a second physical device autonomously performs one or more operations	

	defined by the first one or more instruction sets for operating the first physical device.	65
2.	Claim 6.....	68
a.	[6.1] The system of claim 1, wherein the second physical device autonomously performs one or more operations defined by the first one or more instruction sets for operating the first physical device,.....	68
b.	[6.2] and wherein the first one or more instruction sets for operating the first physical device are modified so that the second physical device can autonomously perform one or more operations defined by the first one or more instruction sets for operating the first physical device.	69
3.	Claim 7: The system of claim 1, wherein the portions of the new one or more digital pictures include portions of the new one or more digital pictures that depict recognized objects and the portions of the first one or more digital pictures include portions of the first one or more digital pictures that depict recognized objects.....	72
4.	Claim 8.....	74
a.	[8.1] The system of claim 1, wherein the first physical device includes a robot, a vehicle, an appliance, an electronic device, or a mechanical machine and	74
b.	[8.2] the second physical device includes a robot, a vehicle, an appliance, an electronic device, or a mechanical machine, and.....	76
c.	[8.3] wherein the one or more processor circuits are included in at least one of: the first physical device, the second physical device, or a third device.	77
5.	Claim 10: The system of claim 1, wherein the first one or more digital pictures includes a one or more still digital pictures or one or more motion digital pictures and the new one or more digital pictures include one or more still digital pictures or one or more motion digital pictures.....	77
6.	Claim 12.....	79
a.	[12.1] The system of claim 1, wherein the memory further stores at least a knowledge structure that	

	includes: a first knowledge cell including the first one or more digital pictures correlated with the first one or more instruction sets for operating the first physical device, and a second knowledge cell including a second one or more digital pictures correlated with a second one or more instruction sets for operating the first physical device, and	80
b.	[12.2] wherein the first knowledge cell is connected with the second knowledge cell by a directed connection that indicates temporal order of the first knowledge cell and the second knowledge cell	83
7.	Claim 13	85
8.	Claim 14: The non-transitory machine readable medium of claim 13, wherein the second physical device autonomously performs one or more operations defined by the first one or more instruction sets for operating the first physical device, and wherein the first one or more instruction sets for operating the first physical device are modified so that the second physical device can autonomously perform one or more operations defined by the first one or more instruction sets for operating the first physical device.	87
9.	Claim 15: The non-transitory machine readable medium of claim 13, wherein the portions of the new one or more digital pictures include portions of the new one or more digital pictures that depict recognized objects and the portions of the first one or more digital pictures include portions of the first one or more digital pictures that depict recognized objects.	87
10.	Claim 16: The non-transitory machine readable medium of claim 13, wherein the memory further stores at least a knowledge structure that includes: a first knowledge cell including the first one or more digital pictures correlated with the first one or more instruction sets for operating the first physical device, and a second knowledge cell including a second one or more digital pictures correlated with a second one or more instruction sets for operating the first physical device, and wherein the first knowledge cell is connected with the second knowledge cell by a	

	directed connection that indicates temporal order of the first knowledge cell and the second knowledge cell.	88
11.	Claim 17	89
12.	Claim 18: The method of claim 17, wherein the autonomously performing of (e) is performed by the second physical device, and wherein the first one or more instruction sets for operating the first physical device are modified so that the second physical device can autonomously perform one or more operations defined by the first one or more instruction sets for operating the first physical device.	91
13.	Claim 19: The method of claim 17, wherein the portions of the new one or more digital pictures include portions of the new one or more digital pictures that depict recognized objects and the portions of the first one or more digital pictures include portions of the first one or more digital pictures that depict recognized objects.....	92
14.	Claim 20: The method of claim 17, wherein the memory further stores at least a knowledge structure that includes: a first knowledge cell including the first one or more digital pictures correlated with the first one or more instruction sets for operating the first physical device, and a second knowledge cell including a second one or more digital pictures correlated with a second one or more instruction sets for operating the first physical device, and wherein the first knowledge cell is connected with the second knowledge cell by a directed connection that indicates temporal order of the first knowledge cell and the second knowledge cell.	92
C.	Ground 2: Claims 3, 5, and 11 Are Obvious Based on Grotmol in View of the Knowledge of a POSITA	93
1.	Claim 3	93
a.	[3.1] The system of claim 1, wherein the first one or more instruction sets for operating the first physical device are obtained, via a hardwired connection, from at least one portion of the one or more processor circuits or from at least one portion of another one or more processor circuits, and	94

b.	[3.2] wherein the one or more processor circuits or the another one or more processor circuits are used at least for operating the first physical device.....	96
2.	Claim 5: The system of claim 1, wherein the causes the one or more processor circuits to execute the first one or more instruction sets for operating the first physical device includes inserting, via a hardwired connection, the first one or more instruction sets for operating the first physical device into at least one portion of the one or more processor circuits.	97
3.	Claim 11: The system of claim 1, wherein the artificial intelligence unit includes at least one of: a hardware element that is included in the one or more processor circuits, a hardware element that is included in a second one or more processor circuits, an application operating on the one or more processor circuits, an application operating on a third one or more processor circuits, or an element coupled to the one or more processor circuits, and wherein the one or more processor circuits include one or more microcontrollers.....	98
D.	Ground 3: Claims 2 and 4 Are Obvious Based on Grotmol in View of Cosic 192	104
1.	Motivation to Combine Grotmol and Cosic 192	104
2.	Claim 2	108
a.	[2.1] The system of claim 1, wherein the first one or more instruction sets for operating the first physical device are obtained from an application for operating the first physical device using at least a tracing of the application for operating the first physical device,.....	108
b.	[2.2] wherein the application for operating the first physical device runs on the one or more processor circuits or on another one or more processor circuits.....	114
3.	Claim 4	114
a.	[4.1] The system of claim 1, wherein the causes the one or more processor circuits to execute the first one or more instruction sets for operating the first physical device includes: instrumenting an	

	application for operating the first physical device with the first one or more instruction sets for operating the first physical device, or instrumenting an application for operating the second physical device with the first one or more instruction sets for operating the first physical device, and	115
b.	[4.2] wherein the application for operating the first physical device or the application for operating the second physical device runs on the one or more processor circuits.	118
E.	Ground 4: Claim 9 Is Obvious Based on Grotmol in View of Hickman	118
1.	Motivation to Combine Grotmol and Hickman.....	119
2.	Claim 9: The system of claim 1, wherein the memory further stores at least a second one or more digital pictures correlated with a second one or more instruction sets for operating the second physical device, 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 second physical device are learned in a second learning process that includes operating the second physical device at least partially by a second user.....	125
IX.	The '583 Patent.....	130
A.	The Claims	130
B.	Ground 1: Claims 1-9, 11-24, 26, 28-29, and 31-34 Are Anticipated by Grotmol.....	142
1.	Claim 1	142
a.	[1.1] A system comprising:	142
b.	[1.2] one or more processors; and	143
c.	[1.3] 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:.....	146

d.	[1.4] 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:	147
e.	[1.5] receiving or generating a first one or more instruction sets for operating the first device; and	150
f.	[1.6] learning the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device.....	153
2.	Claim 2	159
a.	[2.1] The system of claim 1, wherein the receiving the first one or more digital pictures includes receiving the first one or more digital pictures from a picture capturing apparatus,.....	159
b.	[2.2] and wherein the receiving the first one or more instruction sets for operating the first device includes receiving the first one or more instruction sets for operating the device from: an application for operating the first device, one or more microcontrollers, one or more processors, or one or more actuators,.....	161
c.	[2.3] 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 storing the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device into or onto at least one of: at least one non-transitory machine readable medium of the one or more non-transitory machine readable media, one or more volatile memories, one or more non-volatile memories, one or more storage devices, or one or more storage systems.....	163
3.	Claim 3: The system of claim 1, 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 determining that the first one or more instruction sets for operating the first device temporally correspond to the first one or more digital pictures.....	164
4.	Claim 4	167

a.	[4.1] 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:	167
b.	[4.2] receiving or generating a new one or more digital pictures;	167
c.	[4.3] 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.....	169
d.	[4.4] 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.	173
5.	Claim 5: The system of claim 4, 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 device includes executing the first one or more instruction sets for the device.	176
6.	Claim 6: 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 portions of the new one or more digital pictures that represent one or more objects and one or more portions of the first one or more digital pictures that represent one or more objects.	178
7.	Claim 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.	183

8. Claim 8: The system of claim 4, wherein the first one or more instruction sets for operating the first device include one or more information about one or more states of: the first device, or a portion of the first device. 185
9. Claim 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. 186
10. Claim 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. 189
11. Claim 12: The system of claim 4, wherein the new one or more digital pictures depict at least a portion of the first

	device's surrounding, and wherein the first one or more instruction sets for operating the first device are applied to the first device, and wherein the first device is caused to perform the one or more operations defined by the first one or more instruction sets for operating the first device.	191
12.	Claim 13: 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 are applied to the second 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.....	193
13.	Claim 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.	195
14.	Claim 15	199
	a. [15.1] 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 at least:	199
	b. [15.2] modifying: the first one or more instruction sets for operating the first device,.....	199
	c. [15.3] and wherein the determining...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... ..	201
	d. [15.4] and wherein the causing...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	202
15.	Claim 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 ... and wherein the at least partial match between the new one or more digital pictures and the first one or more digital pictures includes ... (iii) at least partial match between the new one or more digital pictures and the modified 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.....	202
16.	Claim 17	205
	a. [17.1] 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:	205
	b. [17.2] modifying: the first one or more digital pictures, or... ..	205
	c. [17.3] and wherein the learning ... 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	207
	d. [17.4] and wherein the determining ... 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	209
17.	Claim 18	211
	a. [18.1] 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:	211
	b. [18.2] modifying: the first one or more instruction sets for operating the first device, or	211

- c. [18.3] and wherein the learning ... 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 215
 - d. [18.4] and wherein the determining ... 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 216
 - e. [18.5] and wherein the causing ... 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 216
- 18. Claim 19: The system of claim 4, wherein the system
further comprising: an artificial intelligence system, 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, at least in
part by the artificial intelligence system, 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, at least in part by the artificial intelligence
system, 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.....217
- 19. Claim 20: The system of claim 19, wherein the artificial
intelligence system includes: one or more inputs for
receiving one or more digital pictures, and one or more
outputs for providing one or more instruction sets, and
wherein the learning ... includes applying the first one or
more digital pictures to the inputs of the artificial
intelligence system, and wherein the determining ...
includes: applying the new one or more digital pictures to

	the inputs of the artificial intelligence system; and receiving the first one or more instruction sets for operating the first device from the outputs of the artificial intelligence system.....	223
20.	Claim 21: The system of claim 20, wherein the one or more inputs for receiving the one or more digital pictures include: one input for one digital picture, multiple inputs for multiple digital pictures, one input for one portion of one digital picture, multiple inputs for multiple portions of one digital picture, multiple inputs for multiple portions of multiple digital pictures, an input for one representation of one digital picture, multiple inputs for multiple representations of multiple digital pictures, one input for one representation of one portion of one digital picture, multiple inputs for multiple representations of multiple portions of one digital picture, or multiple inputs for multiple representations of multiple portions of multiple digital pictures.	227
21.	Claim 22: The system of claim 19, wherein the artificial intelligence system includes: a knowledgebase, or a neural network.....	228
22.	Claim 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.....	228
23.	Claim 24: The system of claim 4, wherein the first one or more digital pictures are: one or more whole digital pictures, ... one or more collections of pixels, or one or more collections of values.	231
24.	Claim 26	232
	a. [26.1] The system of claim 4, wherein at least some elements of the system are included in: a single device, or multiple devices, ...	232
	b. [26.2] and wherein the one or more processors include: one or more microcontrollers, one or more computing devices, or one or more electronic devices, ...	232

- c. [26.3] and wherein the first device includes a robot, a vehicle, an appliance, an electronic device, or a mechanical machine, ...234
 - d. [26.4] and wherein the second device includes a robot, a vehicle, an appliance, an electronic device, or a mechanical machine, ...234
 - e. [26.5] and wherein the first one or more digital pictures include: one or more still digital pictures, or one or more motion digital pictures, and wherein the new one or more digital pictures include: one or more still digital pictures, or one or more motion digital pictures, ...236
 - f. [26.6] and wherein an instruction set of the first one or more instruction sets for operating the first device includes at least one of: only one instruction, a plurality of instructions, one or more inputs, one or more commands, one or more keywords, one or more symbols, one or more operators, one or more variables, one or more parameters, one or more characters, one or more numbers, one or more values, one or more signals, one or more binary bits, one or more functions, one or more function references, one or more objects, one or more object references, one or more data structures, one or more data structure references, one or more states, one or more representations of one or more states, one or more representations of one or more user inputs, one or more codes, one or more data, or one or more information. 236
25. Claim 28: The system of claim 4, wherein the at least partial match between the new one or more digital pictures and the first one or more digital pictures is at least in part based on a product of one or more values that represent one or more portions of the new one or more digital pictures and one or more values that represent one or more portions of the first one or more digital pictures. 237
26. Claim 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.	239
27.	Claim 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.	244
28.	Claim 32: The system of claim 4, wherein the first one or more digital pictures that depict the at least the portion of the first device's surrounding include one or more digital pictures that depict a representation of the at least the portion of the first device's surrounding.	245
29.	Claim 33	247
30.	Claim 34	248
C.	Ground 2: Claims 10 and 11 are Obvious Based on Grotmol in View of the Knowledge of a POSITA	249
1.	Claim 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.	249

2.	Claim 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.	252
D.	Ground 3: Claims 25 and 27 Are Obvious Based on Grotmol in View of Hickman	254
1.	Motivation to Combine Grotmol and Hickman.....	255
2.	Claim 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 % [sic] 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.	255
3.	Claim 27	264
a.	[27.1] 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,.....	264
b.	[27.2] 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;.....	266
c.	[27.3] 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.	267
E.	Ground 4: Claim 30 Is Obvious Based on Grotmol in view of Hoffman.....	269
1.	Motivation to Combine Grotmol and Hoffman	269
2.	Claim 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.	272
X.	Declaration.....	275

I. Introduction

1. I have been retained by Petitioner Tesla, Inc. (“Petitioner” or “Tesla”), as an independent expert witness in the above-captioned *inter partes* reviews (“IPRs”), in which Tesla has requested that the U.S. Patent and Trademark Office cancel as unpatentable all claims of U.S. Patent Nos. 10,102,449 (“the ’449 patent”) and 11,055,583 (“the ’583 patent”) (together, “the Digital Picture Patents”).

2. This declaration sets forth my analyses and opinions based on my knowledge, experience, and the materials I have considered. As I explain below, it is my opinion that all claims of the ’449 patent and ’583 patent are directed to subject matter that was routine, conventional, and well known in the art before the patents’ priority dates. As would be readily appreciated by one of skill in the art, the claims of the ’449 patent and the ’583 patent are anticipated and/or rendered obvious by the prior art references discussed herein.

3. I have reviewed the documents referenced in this declaration. I understand they have been submitted as exhibits in conjunction with Tesla’s Petitions for IPR.

II. Summary of Opinions

4. Based on my knowledge, experience, and the materials that I have reviewed, it is my opinion that claims 1-20 of the ’449 patent and claims 1-34 of the ’583 patent are anticipated and/or obvious. Specifically:

(i) Claims 1, 6-8, 10, and 12-20 of the '449 patent are anticipated by U.S. Patent No. 9,604,359 ("Grotmol") (Ex. 1005);

(ii) Claims 3, 5, and 11 of the '449 patent are obvious based on Grotmol in view of the Knowledge of a POSITA;

(iii) Claims 2 and 4 of the '449 patent are obvious based on Grotmol in view of U.S. Patent No. 9,443,192 ("Cosic 192") (Ex. 1006);

(iv) Claim 9 of the '449 patent is obvious based on Grotmol in view of U.S. Patent No. 8,639,644 ("Hickman") (Ex. 1007);

(v) Claims 1-9, 11-24, 26, 28-29, and 31-34 of the '583 patent are anticipated by Grotmol;

(vi) Claims 10 and 11 of the '583 patent are obvious based on Grotmol in view of the Knowledge of a POSITA;

(vii) Claims 25 and 27 of the '583 patent are obvious based on Grotmol in view of Hickman; and

(viii) Claim 30 of the '583 patent is obvious based on Grotmol in view of U.S. Patent No. 9,283,674 ("Hoffman") (Ex. 1009).

III. Qualifications and Compensation

5. I have a Ph.D. in Electrical & Computer Engineering from North Carolina State University (or "NCSU"), which I earned in 1998. Prior to that I received a M.S. in Integrated Manufacturing Systems Engineering from North

Carolina State University in 1994 and a B.S. in Mechanical Engineering from the University of Virginia in 1990.

6. Since 1991, I have been active in the robotics, sensors, artificial intelligence, pattern analysis, computer vision, image/video processing, and automation fields. I have authored numerous publications and have co-authored a textbook entitled *Computational Intelligence*. My PhD dissertation is entitled *Pattern Analysis, Tracking and Control for Autonomous Vehicles Using Neural Networks*. It described the use of unsupervised training (i.e., the Hyper-Ellipsoid Clustering Neural Network, or "HECNN"), and supervised training (i.e., the Region-Feature Neural Network, or "RFNN") for dissimilar mobile robots to learn aspects of their common environments through multiple types of sensors (e.g., acoustic rangefinding, infrared proximity, tactile and imaging), build knowledge libraries, and then transfer knowledge between for operational use and/or refinement. Among other outputs, the robots learned to recognize spaces, landmarks within those spaces, and then plan and track motion from the accumulated and shared knowledge. Efficiency and accuracy were shown to increase significantly with the combined unsupervised and supervised learning, and knowledge transfers.

7. I have designed, built, and marketed robots, automated systems, and components thereof, including ground mobile robots, unmanned aerial vehicles (UAV), unattended sensors automated storage and retrieval systems (ASRS),

submersible mobile robots, proof-of-concept extra-terrestrial robots, manually manipulated (“twiddled”) submersible foils, wall climbing mobile robots, and augmented reality systems.

8. I have taught the following courses: Introduction to Robotics and Automation (Duke and NCSU); Introduction to Control Theory (Duke and NCSU); Distributed Real-Time Controls (NCSU); and myriad independent studies courses in the areas of robotics, automation, artificial intelligence, autonomy, pattern analysis, and control systems. I have also served on several MS- and PhD-level graduate student committees, designed qualifying exam problems, served on the NCSU IMSEI Board, and participated in curriculum development for undergraduate and graduate level programs.

9. I have also started, sponsored, and advised multiple champion student teams for international robot competitions including, but not limited to, the NASA/ASCE Extra-Terrestrial Robotics Competition, the DARPA Grand Challenge, the AUVSI/ONR Autonomous Underwater Vehicle Competition, and the European CLAWAR Wall-Climbing Robot competition.

10. I have initiated multiple unmanned aerial vehicle (UAV) projects at both academic and industry levels. Academic UAVs include, but are not limited to, Quadcopters with Hybrid Remote and Autonomous Control, Marsupial UAVs that Deploy and Recover Unmanned Ground Vehicles, and Wall-Climbing UAVs.

Additionally, the AngelFish Cross-Domain Submersible UAV, a DARPA ASW program that I incubated and procured funding for at Teledyne, included a partnership with North Carolina State University.

11. In 1999, while I was employed by Nekton Research (“Nekton”), I captured and managed multiple programs sponsored by the Department of Defense (DoD) and private-sector companies that focused primarily on autonomous underwater vehicles (AUV), remotely operated vehicles (ROV), TwiddleFish toys, and indirect-fire projectiles. During my tenure, Nekton entered into a joint venture agreement with the founders of what eventually became known as Parata Systems – a pharmacy automation solutions company, which I supported launching.

12. After leaving Nekton in March 2002, I joined a then New Jersey-based company called Avionic Instruments, Inc. (“Avionic”). While at Avionic, I continued developing and marketing robots and supported engineering related to various design, manufacturing, quality and assembly issues on the core aerospace product lines. Avionic product lines include, but are not limited to, ducted fans, transformer rectifier units (TRU), regulated TRU (RTRU), auxiliary power unit (APU) control systems, power distribution systems (PDU), frequency converters, corner clamps, and VRAM attractors/thrusters. Customers included DoD, NASA, Boeing, Sikorsky, Augusta-Westland, Dassault, and Lockheed-Martin.

13. In late 2004, soon after Avionic was sold, I, along with others from Avionic, secured funding to develop an alpha-level multi-dispenser robotic system and afforded me time to write the RxMedic business plan and raise multiple rounds of venture capital. In late 2006, RxMedic (called “APDS” until November 2006) was launched as a stand-alone, sole-focus venture. After the operational launch of RxMedic, I served as General Manager and eventually Chief Technical Officer. Through my roles at RxMedic, I oversaw the development of the RxMedic ADS robot, managed the intellectual property portfolio, coordinated sales and marketing, and provided strategic, fiscal, and operational leadership. In May 2010, J.M Smith Corporation acquired RxMedic, and in an effort to assist in the change of ownership, I served as a director of RxMedic until May 2011.

14. Also in late 2004, after Avionic was sold, I, along with others from Avionic, secured funding for Vortex HC through DoD contracts and robot sales, to continue developing the VRAM Mobile Robot Platform (VMRP – a wall-climbing robot), the ARTEMIS AUV (a holonomic submersible robot for counter-mine and counter-obstacle operations), the submersible crawler, the nuclear-grade boiler water reactor (BWR) inspection robot, and other robot products centered around the VRAM. Some DoD programs were/are classified, for which I maintained a SECRET clearance at both the personal and facilities level, and served as the facilities security

officer (FSO). In late 2006, corresponding to the full launch of RxMedic, Vortex HC technologies were largely licensed to Teledyne SeaBotix, SeaRobotics, and HDT.

15. However, I have continued to support Vortex HC licensees and customers to date. After the sale of RxMedic, and after fulfilling my 12-month employment obligation, I joined Teledyne Technologies in summer 2011. I served as the Senior Manager for the RTP division of Teledyne Scientific, and supported multiple DoD sponsored robotic-focused programs. Some of these programs were/are classified, for which I maintained a personal SECRET clearance. Among these programs, were cargo unmanned ground vehicles (CUGV), squad-level autonomous unmanned ground vehicles (UGV), autonomous underwater vehicles (AUV), unmanned underwater vehicles (UUV), and a cross-domain autonomous vehicle capable of transitioning between air-, surface- and underwater-domains. Cross-domain vehicles that were evaluated, designed and prototyped included, but were not limited to, the AngelFish (later called "EagleRay") submersible AUV for anti-submarine warfare, and a ball-shaped robot for countermine operations on the ground, in the beach zone and in the surf zone. Sensor design, refinement and signal processing was a major component of each program. Sensors employed include, but are not limited to: proximity sensors; ranging sensors; electro-optical imaging; long-, short- and mid-wave infrared (IR); inertial measurement units (IMU); optical flow; and radiofrequency (RF). Additionally, control systems were designed, refined and

integrated into the aforementioned systems. Most control systems were closed-loop, in that they utilized sensor-based feedback; others were open-loop, where states were estimated with little or no feedback.

16. In late 2013, Avionic, a Transdigm business unit at that point, requested that I return to turn around a supply-chain issue, and assume management of its Sikorsky S97 Raider helicopter program. The S97 Raider is purported to be the fastest, most maneuverable helicopter, due to its coaxial, counter-rotating variable pitch wings, and an aft-based push-propeller. Avionic also controlled two business units named Acme Aerospace (Acme) and Aerospace Cooling Solutions (ACS). Avionic, Acme and ACS supported the S97 program, which has met milestones and continues to produce multiple successful demonstrations. In 2013, Transdigm expanded my role to include directorship of the Avionic, Acme and ACS sales and marketing team, and to report operational and financial status at six-week intervals. Transdigm required that I move my family to New Jersey in late 2014, which influenced my decision to resign and assume the CEO role of Delta Five Systems in Raleigh, NC.

17. As CEO of Delta Five, I led a private-equity backed, hospitality-focused venture, coordinating the company's launch and strategic path including an early stage pivot. Initially focused on back-end robotics and automation, Delta Five shifted to address the rampant bed bug problem—a top priority for hoteliers—with

a novel unattended sensor and trap, called the Telemetered Pest Monitoring System (TPMS), which is based on computer vision and internet of things (IoT), and has proven capable of scaling to other pests. In addition to the TPMS, Delta Five developed a novel means to mass produce a natural, unscented aggregation pheromone that, in concert with placement and heat, lured invertebrates to the TPMS. After leading Delta Five and serving on its Board for four years, I resigned to join ARA, but continue serving Delta Five as an investor and supportive transition agent.

18. In 2018, I joined ARA as sector CTO. In this role, my primary responsibilities included leading the commercialization and technology transition efforts, which require technical, business and corporate development, as well as the raising of capital and managing matrixed rosters. I assumed my other current role at ARA, Corporate Director of IP and Product Commercialization, in 2020. In this role, I evaluate, quantify, present, and recommend portfolios of IP, product lines, and ventures on the basis of competitive landscapes, addressable markets, sales strategy, resource availability and technology readiness. I also spearhead investments and provide general venture oversight.

19. My curriculum vitae is attached as Ex. 1004 to the IPR petition, and provides further information about my experience, expertise, and presentations.

20. Through my professional experience, I have gained extensive expertise in robotics, sensors, artificial intelligence, automation, computer vision, and image/video processing.

21. I am being compensated at my standard rate of \$500/hour. My compensation is in no way contingent upon my opinions or the outcome of the proceeding.

IV. Relevant Legal Standards

22. I am not an attorney, and therefore my understanding of patent law and the legal standards set forth in this report is based on explanations provided to me by counsel.

23. I understand that for any claim of a patent to claim priority to an earlier application (i.e., to benefit from the earlier application's filing date), the claims of the later patent must be fully supported by the disclosure of the earlier patent application to which priority is claimed. I understand that in order for the claims to be supported, the earlier application's disclosure must be sufficient to allow a person of ordinary skill in the art to reasonably conclude that the inventors were in possession of the claimed invention.

24. I understand that the '449 patent claims priority to U.S. Patent Application No. US 15/818,765, filed on November 21, 2017, and the '583 patent claims priority to U.S. Patent Application No. US 15/822,150, filed on November

26, 2017. For the purposes of my opinions, I have considered the state of the art as of and shortly before these claimed priority dates, and the level of knowledge that a POSITA would have possessed at that time. Unless I state otherwise, whenever I refer to any principle or technical subject matter as having been known or understood, this is meant to denote the knowledge and understanding of a POSITA at or prior to November 21, 2017.

A. Claim Construction

25. It is my further understanding that the numbered paragraphs at the end of the disclosure of a U.S. Patent are the patent “claims” that define the metes and bounds of the alleged invention. I understand that the claims of the Digital Picture Patents are what is being challenged in the present IPR proceedings.

26. I have been informed that, in this proceeding, the Board must determine the scope of the claims by giving the claims their ordinary and customary meaning in light of the specification, as the claims would be interpreted by one of ordinary skill in the art.

27. I understand that patent claims generally include a “transitional” term or phrase, such as “consisting” or “comprising,” which may connect the preamble of the claim to the body of the claim. I have been informed that if a claim uses the term “consisting” as a transition term, that means that the claim is a “closed” claim, which means that the claim is limited to the claim features that follow the transition

term and nothing else. On the other hand, I understand that the transition term “comprising” denotes an “open” claim, which means that the claim is not limited to only the features recited in the claim, and could encompass the listed elements as well as other unrecited elements.

28. I understand that 35 U.S.C. § 112(f) created an exception to the general rule of claim construction called a “means plus function” limitation. These types of terms and limitations should be interpreted to cover only the corresponding structure described in the specification, and equivalents thereof. I also understand that a limitation is presumed to be a means plus function limitation if: (a) the claim limitation uses the phrase “means for”; (b) the “means for” is modified by functional language; and (c) the phrase “means for” is not modified by sufficient structure for achieving the specified function.

29. I understand that a structure is considered structurally equivalent to the corresponding structure identified in the specification only if the differences between them are insubstantial. For instance, if the structure performs the same function in substantially the same way to achieve substantially the same result. I further understand that a structural equivalent must have been available at the time of the issuance of the claim.

B. Invalidity

30. I understand that Tesla bears the burden of proving that the challenged claims of the Digital Picture Patents are invalid, and must prove this by a preponderance of the evidence, which means that invalidity must be shown to be more likely than not.

1. Anticipation

31. I have been asked to consider the question of whether the claims of the Digital Picture Patents would have been anticipated by certain prior art references.

32. I understand that determining anticipation of a patent claim requires a comparison of the claim language to the prior art on a limitation-by-limitation basis.

33. I understand that a prior art reference “anticipates” an asserted claim, and thus renders the claim invalid, if all elements of the claim are disclosed in that prior art reference, either explicitly or inherently (i.e., necessarily present or implied), as those elements are arranged in the claim so that the disclosure effectively puts the public in possession of the invention. I understand that a prior art reference anticipates an asserted claim when each and every element of the claim is disclosed in a single embodiment as arranged in the claim. The mere fact that a certain outcome may result from a given set of circumstances is insufficient for inherency; rather, inherency requires that the claimed product or process at issue must be a necessary result from the operation as taught by the prior art reference. I

understand that, although references cannot be combined for anticipation, additional references may be used to interpret an anticipating reference and shed light on what it would have meant to those skilled in the art at the time of the invention.

34. In order to anticipate a claimed invention, a prior art reference must enable one of ordinary skill in the art to make the invention without undue experimentation, i.e., provide an enabling disclosure. A prior art reference need not enable its full disclosure; it only needs to enable the portions of its disclosure alleged to anticipate the claimed invention.

2. Obviousness

35. I have been asked to consider the question of whether the claims of the Digital Picture Patents would have been obvious. I understand that this analysis must be conducted from the perspective of the person of ordinary skill in the art, and whether the skilled artisan would consider any differences between the prior art and what is claimed to have been obvious. To make this assessment, I have been informed that the concept of patent obviousness involves four factual inquiries: (1) the scope and content of the prior art; (2) the differences between the claimed invention and the prior art; (3) the level of ordinary skill in the art; and (4) secondary considerations of non-obviousness. I have been instructed that one must not engage in hindsight. Rather, I understand that one should instead consider what the person

of ordinary skill in the art would have reason to pursue further, and steps that were routinely done, such as in response to known problems, or obstacles.

36. It is my understanding that the following is a non-exhaustive list of rationales that may support the obviousness of an invention: combining prior art elements according to known methods to yield predictable results; simple substitution of one known element for another to obtain predictable results; use of a known technique to improve a similar device (method, or product) in the same way; applying a known technique to a known device (method, or product) ready for improvement to yield predictable results; choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; and some teaching, suggestion, or motivation in the prior art that would have led a POSITA to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention.

37. It is my understanding that the motivation to combine prior art references may be implicit and may be found in the knowledge of one of ordinary skill in the art, or in the nature of the problem to be solved. Specifically, it is my understanding that an implicit motivation to combine exists not only when a suggestion may be gleaned from the prior art as a whole, but when the “improvement” is technology-independent and the combination of references results in a product or process that is more desirable, for example because it is stronger,

cheaper, cleaner, faster, lighter, smaller, more durable or more efficient. It is my further understanding that the motivation to combine references may be found in the nature of the problem to be solved where prior art references are directed to precisely the same problem.

38. I also understand that prior art may be relied on for its express disclosure and teachings. I also understand that the prior art may be relied upon for a teaching of features that are necessarily present in the prior art reference even if that specific feature is not expressly or explicitly disclosed.

39. I understand that before reaching any final conclusion on obviousness, the obviousness analysis requires consideration of objective indicia of non-obviousness, if any such indicia are offered. These must be considered to ensure that, for example, there were not some unanticipated problems, obstacles or hurdles that may seem easy to overcome in hindsight, but which were not readily overcome prior to the relevant invention date of the patents/claims at issue here. I understand that these objective indicia are also known as “secondary considerations of non-obviousness,” and may include long-felt but unmet need and unexpected results, among others. I also understand, however, that any offered evidence of secondary considerations of non-obviousness must be commensurate with the scope of the challenged claims. This means that for any offered evidence of secondary considerations of non-obviousness to be given substantial weight, I understand the

proponent of that evidence must establish a “nexus,” or a sufficient connection or tie between that evidence and the merits of the claimed invention, which I understand specifically incorporates any novel element(s) of the claimed invention. If the secondary consideration evidence offered actually results from something other than the merits of the claim, then I understand that there is no nexus or tie to the claimed invention. I also understand it is the Patent Owner who has the burden of proving that a nexus exists, and I understand that secondary considerations will not overcome a strong showing of obviousness.

C. Person of Ordinary Skill in the Art

40. I have been asked to review the Digital Picture Patents from the perspective of a person of ordinary skill in the art (“POSITA”) as of the earliest claimed priority date for the patents—November 21, 2017.

41. It is my opinion that a POSITA relevant to the Digital Picture Patents would have had at least: (1) a bachelor’s degree in electrical engineering, computer engineering, computer science, or equivalent course work with two to three years of work experience in computer vision, image/video processing, and/or machine learning; or (2) a master’s degree in electrical engineering, computer engineering, computer science, or equivalent course work with a focus in computer vision, image/video processing, and/or machine learning.

42. Based on my education, training and experience, it is my opinion that I can accurately represent the views of a POSITA as of the earliest claimed priority date of November 21, 2017 for the '449 patent and the '583 patent. The opinions I provide in this declaration are provided using the viewpoint of the POSITA as of that date.

V. Background of the Technology

43. The asserted claims of the Digital Picture Patents are directed to the well-known practice of employing “devices, systems, methods, and/or applications for automating devices and/or systems.” '449 patent, Abstract. Specifically, the asserted claims are directed to a system that “captures one or more digital pictures, receives one or more instruction sets, and learns correlations between the captured pictures and the received instruction sets.” '583 patent, Abstract. Devices' operations have been learned by intelligent systems for decades, at least as far back as several decades. Teach pendants, for example, have been integral to robots dating back to at least the 1990's, where human operators would guide robots through various tasks based on sensory feedback, and the robots would learn to operate autonomously (i.e., without constant operator control) based on perceived circumstances.

44. Other types of supervised, unsupervised, and/or human-in-the-loop (HITL) training have been applied to mobile robots for decades, as well. For

example, AAAI and IJCAI competitions from the mid 1990's demonstrated mobile robots that would learn to recognize objects and features within perceived environments and, based on supervised training, initiate one or more actions including avoidance, pursuit, and manipulation. These mobile robots utilized computer vision, whiskers, ultrasonic, and infrared sensors for perception. Many would also evaluate the scene(s) utilizing a variety of well-known sensor processing techniques such as color histograms, transforms, moments, codes, masks, and neural networks. Feedback from operator(s) in the loop might adjust, refine and optimize parameters (e.g., neural network synaptic weights and biases), and then be stored as "knowledge." Multi-mobile robot systems, sometimes referred to as "multi-agent systems," from as far back as the 1990's were also shown to exchange knowledge and/or compile their respective observations to create a universal knowledge base that can be broadly used by any of a multitude of mobile robots. During subsequent operation, "knowledge" can be locked, deleted, added, and/or modified, depending on the demand for adaptability.

45. At its core, both supervised and unsupervised learning utilizes pattern analysis from well-known machine learning models, including statistics, Euclidean distances, moments, discriminators, masks, and the like. Unsupervised learning employs algorithms that learn patterns that were not previously "labeled" or "identified"; and learning features or other aspects about input data (e.g., images)

can be used for clustering and/or certain aspects of supervised learning. In contrast, supervised learning maps input data, including meta-level features and topographies gleaned from unsupervised learning, to target outputs. For example, the target output of a mobile robot vehicle might be to avoid obstacles and/or react to traffic guidance mechanisms, and a human in the loop would provide the mapping from input data (e.g., images) to safe mobile robot vehicle operation. Supervised and unsupervised learning models that integrate HITL allow humans to provide direct feedback to the model(s) for predictions, identification and decision making. As with most control models, training and operation can be based on individual images, image sequences, accumulated image aspects and even the rate at which aspects of images change. Through single and multiple events, the model refines its understanding of the world around it, and significantly improves its abilities to, for example, safely control a driver-optional mobile robot. Examples of HITL AI models abound, including supervised and unsupervised models, and were well understood at least as far back as the 1990s.

VI. The Prior Art

A. U.S. 9,604,359 (“Grotmol”)

46. Grotmol is a U.S. patent filed on January 27, 2015, which issued on March 28, 2017. Grotmol discloses robotic devices that can be trained to operate

autonomously to follow a trajectory. I understand that Grotmol was not of record during prosecution of the applications that became the '449 patent or the '583 patent.

47. Grotmol discloses an apparatus and methods for training and operating a robotic device to follow a trajectory. Ex. 1005, Abstract. As Grotmol notes, a “typical approach” to autonomous navigation training entailed the recording of user controlled movements along a desired path and then a “simpl[e] replay” of the same navigation autonomously. *Id.*, 1:62-2:5. The training approach described in Grotmol is more dynamic. Grotmol utilizes associations between user commands and camera images captured during training, such that during subsequent autonomous operation, a newly acquired image “may be compared with one or more images from the training buffer” to determine how to control the robot. *Id.*, Abstract.

48. Grotmol’s training process first involves a “listen mode,” in which a robot observes a user’s commands in order to “learn an association.” *Id.*, 23:3-11. Grotmol discloses that the robot may then switch—either manually or automatically—into an “autonomous mode,” in which the robot performs tasks autonomously based on learned associations. *Id.*, 23:46-48, 24:1-16. Additionally, Grotmol describes “override” and “control” modes that use varying amounts of user operation to facilitate the learning process. In “override mode,” a robot will operate autonomously until it receives a user’s command, then apply the command, and “try to remember it for the next time the same situation occurs.” *Id.*, 23:12-22. In “correct

mode,” a user may provide “small corrections” to otherwise autonomous operation, where the robot “may sum the teacher’s command with what the robot already knows to get a final motor command.” *Id.*, 23:23-29, 64-67.

49. Grotmol discloses that the associations are learned by storing images captured during training in a “memory buffer (training buffer),” along with “[m]otor commands issued by a trainer corresponding to time instances when the images are acquired.” *Id.*, 39:13-22. Then, “[d]uring 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.” *Id.*, 39:23-28. “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.” *Id.*, 43:5-9.

50. Grotmol further teaches that a robot’s “trained configuration may be stored” and “loaded to one or more other robots in order to provide learned behaviors.” *Id.*, 33:12-35. Thus, a POSITA would have understood Grotmol to disclose that the memory component of one robot can store behaviors learned by another robot.

51. Grotmol discloses that multiple learning processes can be performed to train multiple behaviors, including “simple tasks” and “complex task[s].” *Id.*, 24:17-67, 26:1-20. Grotmol does not explicitly limit the number of users that can train the

multiple behaviors. Nevertheless, a POSITA would have understood the advantages of utilizing different users, of varying skills and knowledge, to train different behaviors appropriately.

B. U.S. 8,639,644 (“Hickman”)

52. Hickman is a U.S. patent filed on May 4, 2012, which issued on January 28, 2014. Hickman discloses a shared robot knowledge base that can instruct robots to interact with objects and learn based on feedback from the robots. I understand that Hickman was not of record during the prosecution of the applications that became the ’449 patent or the ’583 patent. Hickman is cited on page 2 of Grotmol in the list of “References Cited.”

53. Hickman discloses “shared robot knowledge bases for use with cloud computing systems.” Ex. 1007, Abstract. Hickman teaches that “information in the shared robot knowledge bases is updated based on robot experiences so that any particular robot may benefit from prior experiences of other robots.” *Id.* One embodiment of Hickman’s shared robot knowledge base is illustrated in Figure 4, copied below, in which the results of one robot’s learning process for grasping a cup with the appropriate force can be shared with other robots:




FIGURE 4

Ex. 1007, Fig. 4

Id., Fig. 4, 25:3-27:26.

54. As Hickman shows in Figure 4 above, a first robot obtains image data associated with a first object, shown as a cup, communicates with the shared knowledge base to figure out how to grasp the cup, and then attempts to grasp the cup but crushes it. *Id.*, 26:38-56.

In operation, the robot 404 obtains identification information associated with an object 405 via one or more sensors associated with the robot 404. For example, in FIG. 4, **the robot 404 obtains image data (e.g., an image, a 3D image, a point cloud, etc.) associated with the object 405 via optical sensors.**

After obtaining image data associated with the object 405, the robot 404 sends an identification query 406 to the cloud processing engine 402. **The identification query 406 may include at least some of the image data associated with the object 405.** After receiving image data associated with the object 405 from the robot 404, the cloud processing engine 402 may analyze the image data.

After determining that the object is a cup 405 and obtaining information about the cup 405, the cloud computing system 401 can send object data 409 associated with the cup 405 to the robot 404 in response to the identification query 406 received from the robot 404. The object data 409 may include both the identity of the cup 405 and instructions for interacting with the cup 405 based on the application or task that the robot 404 has been instructed to execute or perform.

After receiving the instructions in the object data 409 from the cloud computing system 401, the robot 404 can execute the instructions to grasp the cup 405. However, in this example, when the robot 404 grasps the cup 405 according to the instructions it received in the object data 409 from the cloud computing system 401, **the robot 404 crushes the cup 405** and

causes the coffee to spill 410. In this example interaction, the grasping force specified in the instructions was too high, and as a result, the robot crushed the cup and caused a mess.

Id., 25:14-19, 25:24-30, 26:21-28, 26:38-47¹.

55. Hickman further explains that a human can then take control over the first robot to show it the appropriate force to use when grasping the cup, which will in turn update the shared knowledge base with the appropriate force for grasping the cup:

When an accident happens, a human can “coach” the robot on how much force to use by, for example, manually controlling the robot's hand to grasp a new cup. The robot can capture the grasping force that it used to successfully grasp the cup under the manual control of the human, and then send feedback 411 to the cloud processing system 402 with the modified grasping force. **The cloud processing system 402 can then update 412 task and object data in the shared robot knowledge base 403** to improve how the “grasp” task is applied this particular “cup” object.

Id., 26:47-56.

56. If a second robot later encounters a second cup similar to the first, the second robot can grasp the cup with the appropriate amount of force learned based

¹ Emphasis added unless otherwise stated.

on the updated training data shared by the first robot, and confirm the accuracy of the learned force to the shared knowledge base:

Later, a second robot 413 may encounter an object 414 while performing the same (or a similar) application that requires grasping the same type of cup. The second robot 413 may send image data associated with the object 414 in an identification query 415 to the cloud computing system 401.

The robot 413 may be performing a task that requires the robot 413 to grasp the cup 414. The instructions for interacting with the cup 414 may include instructions for positioning the robot hand around the cup 414, the amount of force to use when grasping the cup 414, and how tightly to close the fingers of the robot hand around the cup 414. Importantly, **the cloud computing system 401 updated the amount of force to use when grasping the particular type of cup that the robot 413 must now grasp (i.e., cup 414) based on the feedback 411 previously received from robot 404 regarding the same type of cup.** As a result, when the robot 413 applies the “grasp” task to the type of “cup” object here (i.e., when robot 413 grasps cup 414), the robot 413 successfully grasps the cup without crushing the cup 419. Thus, **the second robot 413 has in effect learned from the experience of the earlier robot 404.**

After successfully grasping the cup 419, the robot 413 may send feedback 420 to the cloud processing engine 402 that confirms

the accuracy of the instructions that the robot 413 received in the object data 418 from the cloud computing system 401.

Id., 26:57-61, 27:7-26.

C. U.S. 9,443,192 (“Cosic 192”)

57. Cosic 192 is a U.S. patent filed on August 30, 2015, which issued on September 13, 2016. I understand that Cosic 192 has the same inventor as the Digital Picture Patents, and is cited on both. ’449 patent, p.2; ’583 patent, p. 2.

58. Cosic 192 discloses “devices, systems, methods, and/or applications for learning the operation of a computing device or software application, storing this knowledge in a knowledgebase, neural network, or other repository, and enabling autonomous operation of the computing device or software application with partial, minimal, or no user input.” Ex. 1006, Abstract. Cosic 192 describes its system as a “universal artificial intelligence engine” or “UAIE.” *Id.*, 10:14-16 (“In some embodiments, the system further comprises: an universal artificial intelligence engine (UAIE) for autonomous application operating”).

59. Cosic 192 discloses an autonomous robot learning process that is substantially similar to the process claimed in the ’449 patent, in which a robot may record instructions and/or contextual information regarding, for example, an object detected with a camera, and recall and execute learned instructions based on previous circumstances:

In further exemplary embodiments, Robot 70 may be an **autonomous robot**. As such, Robot 70 can be controlled by Control Application 120 that provides autonomous functionalities to Robot 70 through various non-UAIE systems or processes. In one example, Robot 70 may detect an object through its physical contact sensor and Control Application 120 may execute one or more instructions (i.e. Instruction Sets 600 or Operations 610, etc.) to move Robot 70 around the object. In another example, Robot 70 may detect an object through its vision sensor (i.e. video or other camera, etc.) and Control Application 120 may execute one or more instructions (i.e. Instruction Sets 600 or Operations 610, etc.) to move Robot 70 closer to the object... Any other sensory inputs can be received and corresponding operations can be performed in other examples. In some embodiments involving autonomous robots, **UAIE can serve as a learning system to learn or record Control Application's 120 instructions (i.e. Instruction Sets 600 or Operations 610, etc.) and/or contextual information (i.e. Extra Info 630, etc.) used to implement autonomous Robot's 70 operations in various circumstances.** This knowledge (i.e. Knowledgebase 530, Knowledge Cells 800, etc.) can be loaded into another robot, device, machine, and/or application to provide pre-learned knowledge as previously described. The robot, device, machine, and/or application into which the knowledge is loaded may then implement anticipatory instructions (i.e. anticipatory Instruction Sets 600 or Operations 610, etc.) included in the loaded Knowledgebase 530 or

Knowledge Cells 800... UAIE can then implement the subsequent operations autonomously based on similar operations performed in prior similar circumstances, thereby implementing anticipation of Control Application's 120 operations as previously described.

Ex. 1006, 201:45-202:52. Cosic 192 describes various autonomous robot embodiments. *See generally Id.*, 196:60-202:51.




FIG. 63

Ex. 1006, Fig. 63




FIG. 7

'449 patent, Fig. 7

D. U.S. 9,283,674 ("Hoffman")

60. Hoffman is a U.S. patent filed on October 6, 2014, published on July 9, 2015 as U.S. Publication No. 2015/0190925, and issued on March 15, 2016. I understand that Hoffman was not of record during prosecution of the applications that became the '449 patent or the '583 patent.

61. The disclosure of Hoffman relates to a "method of operating a robot includ[ing] electronically receiving images and augmenting the images by

overlaying a representation of the robot on the images.” Ex. 1009, Abstract. Hoffman discloses methods for “[a]utonomous and semi-autonomous robots” that can be operated “without human intervention or be teleoperated by a remote user.” *Id.*, 7:22-26. Hoffman describes embodiments wherein the robot outputs data such as images to a cloud service that processes the data and returns a processed data set such as a layout map of the environment to the robot. *Id.*, 8:1-8 (“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).”




FIG. 6C




FIG. 6D

Ex. 1009, Figs. 6C and 6D

62. As shown above, the user, using the described telepresence software application, can observe the remote video window showing the camera view of the robot and rendered flythrough view along with a map showing a top down view of the surroundings and rendered robot depicting its location on the map. *See, e.g., id.*, 17:19-18:45.

VII. Claim Construction

63. I understand that the words of the claims of a patent are generally given the ordinary and customary meaning as interpreted by a person of ordinary skill in the art at the time of invention. I have been informed that patent claims are construed in light of the claim language, the specification (including the drawings), and the prosecution history (including the references cited during the prosecution of the patent).

A. “means for receiving or generating”; “means for learning” (claim 34)

64. Claim 34 of the '583 patent requires a “means for receiving or generating,” and a “means for learning.” In determining the corresponding structure for these claimed functions, a POSITA would have looked to the '583 specification which describes one or more processor circuits for performing both functions.

In some aspects, the disclosure relates to a non-transitory computer storage medium having a computer program stored thereon, the program including instructions that when executed by one or more processor circuits cause the one or more **processor circuits to perform operations comprising: receiving a first digital picture from a picture capturing apparatus.** The operations may further include **receiving one or more instruction sets** for operating a device. **The operations may further include learning the first digital picture correlated with the one or more instruction sets for operating**

the device. The operations may further **include receiving a new digital picture** from the picture capturing apparatus.

'583 patent, 16:8-35.

VIII. The '449 Patent

A. The Claims

65. I understand that the petition challenges claims 1-20 of the '449 patent. Independent claim 1 of the '449 patent is directed to a system for autonomous device operation in which a device anticipates and autonomously executes instruction(s) corresponding to a newly received digital picture based on comparison to a previously stored digital picture. As I explain in this Declaration, a POSITA, in and prior to 2017, would have understood that there was nothing inventive about this system, which was well known in the art at the time. Claim 1 is reproduced below:

1. A system for autonomous device operating, the system comprising:
 - (a) one or more processor circuits;
 - (b) a memory that stores at least a first one or more digital pictures correlated with a first one or more instructions sets for operating a first physical device,
 - (c) wherein the first physical device includes an actuator for moving at least a portion of the first physical device,
 - (d) 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 physical device are learned in a learning process that includes operating the first physical device at least partially by a user;

(e) an optical camera that captures digital pictures; and

(f) an artificial intelligence unit that:

(g) receives a new one or more digital pictures from the optical camera;

(h) anticipates the first one or more instruction sets for operating the first physical device based on at least partial match between the new one or more digital pictures and the first one or more digital pictures, wherein the anticipates includes at least one of: determining that a number of at least partially matching portions of the new one or more digital pictures and portions of the first one or more digital pictures exceeds a threshold number, or determining that a percentage of at least partially matching portions of the new one or more digital pictures and portions of the first one or more digital pictures exceeds a threshold percentage; and

(i) causes the one or more processor circuits to execute the first one or more instruction sets for operating the first physical device, wherein the causes is performed in response to the anticipates of the artificial intelligence unit, and

(j) wherein the first physical device or a second physical device autonomously performs one or more operations defined by the first one or more instruction sets for operating the first physical device.

66. Independent claim 13 is nearly identical to claim 1, but recites a different preamble that replaces the functionalities tied to claim 1's "artificial intelligence unit" with "[a] non-transitory machine readable medium having a stored thereon, instructions that when executed by one or more processor circuits cause the one or more processor circuits to perform operations comprising"

67. Independent claim 17 is also nearly identical to claim 1, but claims a method rather than an apparatus, and replaces claim 1's "artificial intelligence unit" with "one or more processor circuits."

68. Dependent claims 2-12, which depend from claim 1, claims 14-16, which depend from claim 13, and claims 18-20, which depend from claim 17, can be summarized as follows in the table below based on the additional steps or features that they require. As I further explain in this Declaration, the steps and features added by these dependent claims were also well-known at or prior to 2017, and a POSITA would have understood that there was nothing inventive about the combination of steps and/or features recited in these claims.

Dependent Claim	Additional Features Claimed
Claim 2	In addition to the limitations of claim 1, claim 2 further requires the first instruction sets to be obtained from an application for operating the first device by tracing of the application, where the application runs on one or more processor circuits. + 1st instructions obtained from application for operating 1st device using tracing of application + application for operating 1st device runs on processor circuit(s)
Claim 3	In addition to the limitations of claim 1, claim 3 further requires the first instruction sets to be obtained from the one or more processor circuits via a hardwired connection, where the processor circuits are used to operate the first device. + 1st instructions obtained from processor circuit(s) via a hardwired connection + Processor circuit(s) used to operate 1st device

Dependent Claim	Additional Features Claimed
Claim 4	In addition to the limitations of claim 1, claim 4 further requires causing the one or more processor circuits to execute the first instruction sets to include instrumenting an application for operating the first or second device, where the application runs on one or more processor circuits. + Causing processor circuit(s) to execute 1st instructions includes instrumenting an application for 1st or 2nd device with instructions for 1st device + application for operating 1st or 2nd device runs on processor circuit(s)
Claim 5	In addition to the limitations of claim 1, claim 5 further requires causing the one or more processor circuits to execute the first instruction sets includes inserting the first instruction sets into the one or more processor circuits via a hardwired connection. + 1st instructions inserted into processor circuit(s) via a hardwired connection

Dependent Claim	Additional Features Claimed
Claim 6	In addition to the limitations of claim 1, claim 6 further requires the first instructions sets to be modified for and then performed by the second physical device. + Autonomous performance of 1st instructions by 2nd device + Modification of 1st instructions for 2nd device
Claim 7	In addition to the limitations of claim 1, claim 7 further requires the first, stored digital picture(s) and the second, new digital picture(s) to depict recognized objects. + 1st and 2nd DPs depict recognized objects
Claim 8	In addition to the limitations of claim 1, claim 8 further requires the first and second devices to be robots or vehicles, and the processor circuits to be included in the first, second, or third devices. + 1st device and 2nd device are robots or vehicles + Processors circuits are contained in the 1st, 2nd, or 3rd device

Dependent Claim	Additional Features Claimed
Claim 9	In addition to the limitations of claim 1, claim 9 further requires the memory to also store second digital picture(s) correlated with second instruction set(s) for operating the second device, where the second digital picture(s) or instruction set(s) are learned in a second learning process with a second user operating the second device. + 2nd DP correlated with 2nd instructions for 2nd device + Portion of 2nd DP or 2nd instructions learned with 2nd user operating 2nd device
Claim 10	In addition to the limitations of claim 1, claim 10 further requires the first and new digital pictures to include still digital pictures or motion digital pictures. + 1st and 2nd DPs include still DPs or motion DPs
Claim 11	In addition to the limitations of claim 1, claim 11 further requires the artificial intelligence unit to include a hardware element included in the one or more processor circuits, which include microcontrollers. + AI unit includes a hardware element included in processor circuit(s) + Processor circuit(s) include microcontroller(s)

Dependent Claim	Additional Features Claimed
Claim 12	In addition to the limitations of claim 1, claim 12 further requires the memory to store first and second knowledge cells containing correlated digital pictures and instruction sets, where a directed connection between the first and second knowledge cells indicates a temporal order between the two. + 1st and 2nd knowledge cells • 1st knowledge cell includes 1st DP correlated with 1st instructions • 2nd knowledge cell includes 2nd DP correlated with 2nd instructions + directed connection between 1st and 2nd knowledge cells indicating temporal order
Claim 14	In addition to the limitations of claim 13, claim 14 further requires the first instructions sets to be modified for and then performed by the second physical device. + Autonomous performance of 1st instructions by 2nd device + Modification of 1st instructions for 2nd device

Dependent Claim	Additional Features Claimed
Claim 15	In addition to the limitations of claim 13, claim 15 further requires the first, stored digital picture(s) and the second, new digital picture(s) to depict recognized objects. + 1st and 2nd DPs depict recognized objects
Claim 16	In addition to the limitations of claim 13, claim 16 further requires the memory to store first and second knowledge cells containing correlated digital pictures and instruction sets, where a directed connection between the first and second knowledge cells indicates a temporal order between the two. + 1st and 2nd knowledge cells • 1st knowledge cell includes 1st DP correlated with 1st instructions • 2nd knowledge cell includes 2nd DP correlated with 2nd instructions + directed connection between 1st and 2nd knowledge cells indicating temporal order
Claim 18	In addition to the limitations of claim 17, claim 18 further requires... + Autonomous performance of 1st instructions by 2nd device + Modification of 1st instructions for 2nd device

Dependent Claim	Additional Features Claimed
Claim 19	In addition to the limitations of claim 17, claim 19 further requires the first, stored digital picture(s) and the second, new digital picture(s) to depict recognized objects. + 1st and 2nd DPs depict recognized objects
Claim 20	In addition to the limitations of claim 17, claim 19 further requires the memory to store first and second knowledge cells containing correlated digital pictures and instruction sets, where a directed connection between the first and second knowledge cells indicates a temporal order between the two. + 1st and 2nd knowledge cells • 1st knowledge cell includes 1st DP correlated with 1st instructions • 2nd knowledge cell includes 2nd DP correlated with 2nd instructions + directed connection between 1st and 2nd knowledge cells indicating temporal order

B. Ground 1: Claims 1-9 and 12-20 Are Anticipated by Grotmol

69. As set forth in detail below, claims 1, 6-8, 10, and 12-20 of the '449 patent are anticipated by Grotmol. The discussion below specifies where each

element of the claims is found in Grotmol, and includes a detailed explanation of the significance of the quotations and citations.

1. Claim 1

- a. [1.1] A system for autonomous device operating, the system comprising:**

Limitation [1.1]	Claimed Features
	system

70. Grotmol discloses a system implemented as a “self-contained robot apparatus”:

Another aspect of the disclosure relates to a self-contained robotic apparatus. The apparatus may comprise a platform, a first sensor, a second sensor, a memory component, and one or more physical processors.

An apparatus and methods for training and/or operating a robotic device to follow a trajectory.

In some implementations, a user (e.g., human operator may train the system to Switch between tasks based on the sensory context. FIG. 2 shows a diagram of the **system**, according to some implementations.

Ex. 1005, 3:27-30, Abstract, 10:45-48.

b. [1.2] one or more processor circuits;

71. Grotmol discloses that its robot may contain “one or more physical processors.” *Id.*, 3:27-30. Although Grotmol does not specifically use the term “processor circuits,” a POSITA would have understood the “one or more processors” of Grotmol to include “one or more processor circuits.” The ’449 patent only defines “processor circuits” as “includ[ing] a logic circuit,” which itself “may include a microcontroller.” ’449 patent, 1:67-2:2. Grotmol defines “processor” as follows:

As used herein, the terms “processor”, “microprocessor” and “digital processor” are meant generally to include all types of digital processing devices including, without limitation, digital signal processors (DSPs), reduced instruction set computers (RISC), general-purpose (CISC) processors, microprocessors, gate arrays (e.g., field programmable gate arrays (FPGAs)), PLDs, reconfigurable computer fabrics (RCFs), array processors, secure microprocessors, and application-specific integrated circuits (ASICs). Such digital processors may be contained on a single unitary IC die, or distributed across multiple components.

Id., 8:5-15. Given this definition, a POSITA would have understood Grotmol’s use of the term “processor” to include processor circuits that include logic circuits that may include a microcontroller.

- c. [1.3] a memory that stores at least a first one or more digital pictures correlated with a first one or more instructions sets for operating a first physical device,

Limitation [1.3]	Claimed Features
	access memory including 1st DP correlated with 1st instructions for 1st Device

72. Grotmol discloses accessing a memory that stores multiple training sets, each including an image of a scene from a robot vehicle's operation (digital picture) correlated with a motor control indication configured to cause the robot to execute an action (instruction set):

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.

A robotic vehicle may utilize a camera and stores the sequence of images of a visual scene when following a trajectory during training in an ordered buffer. **Motor commands associated with a given image may be stored.**

In one or more implementations, such as object recognition, and/or obstacle avoidance, the output provided by the sensor component 166 may comprise a stream of pixel values associated with one or more digital images.

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).

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.

Id., 3:35-39, Abstract, 8:66-9:3, 10:23-26, 39:13-22. Grotmol also teaches that one embodiment for this method is autonomous vehicle navigation. *Id.*, 38:36-58.

73. Grotmol discloses that the stored images are “associated with” or “correspond to” contemporaneously issued motor control indications. *Id.*, 33:59-62

(“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”), Abstract (“Motor commands *associated with* a given image may be stored.”). A POSITA would have understood the associations between the motor commands and images in Grotmol to correspond to the claimed “correlat[ion].” Grotmol further explains that the associated motor commands may correspond to the time when the images are captured. *Id.*, 39:17-19 (“Motor commands issued by a trainer *corresponding to* time instances when the images are acquired may be stored.”), 43:6-9 (“[M]otor 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.”).

74. A POSITA would understand the “instruction sets” of the ’449 patent to correspond to the motor commands of Grotmol, due to the way the term is broadly described in the specification. The ’449 patent discloses instruction sets of any programming language or construct that can be executed on a processor or logic circuit, used by an actuator to operate a device, and can be received by any element of the system.

In some embodiments, an instruction set (i.e. Instruction Set 526, etc.) may be used or executed by a processor (i.e. Processor 11, etc.) for operating a device... may be used, executed, or

produced by a logic circuit (i.e. Logic Circuit 250, etc.) for operating a device... [or] may be used by an actuator for operating a device. For example, such instruction set may be or include one or more inputs into an actuator.

In some designs, an instruction set can be received from a processor, application program, logic circuit, and/or other processing element... In some aspects, an instruction set can be received from memory (i.e. Memory 12, etc.), hard drive, or any other storage element or repository. In other aspects, an instruction set can be received over a network such as Internet, local area network, wireless network, and/or other network... In general, **an instruction set can be received by any element of the system.**

Instruction Set 526 may include one or more instructions or commands of a high-level programming language such as Java or SQL, a low-level language such as assembly or machine language, an intermediate language or construct such as bytecode, and/or **any other language or construct.**”

’449 patent, 141:62-142:9, 142:36-143:13, 82:67-83:19.

75. Grotmol shows in Figure 8, copied below, that the “memory” component is connected to the “sensory” functionality through the “processing”

functionality, consistent with claim 1's requirements that the images stored in the memory are detected by the robot's sensors:




FIG. 8

Ex. 1005, Fig. 8 (annotated).

76. Grotmol further discloses accessing the training sets stored in the memory component to perform autonomous operation. *Id.*, Abstract (“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.”); *see also id.*, 39:23-26 (“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.”).

- d. **[1.4] wherein the first physical device includes an actuator for moving at least a portion of the first physical device,**

	Claimed Features
--	-------------------------

Limitation [1.4]	Actuator for moving 1st Device
-----------------------------------	---------------------------------------

77. The '449 patent provides various examples of an “actuator,” including, inter alia, a “motor.” '449 patent, 69:34-39 (“Examples of an actuator include *a motor*, a linear motor, a servomotor, a hydraulic element, a pneumatic element, an electro-magnetic element, a spring element, and/or other actuators. Examples of types of actuators include a rotary actuator, a linear actuator, and/or other types of actuators.”)

78. Grotmol discloses a robotic device capable of autonomous operation that includes motors and other “actuators” for moving the device and controlling its behaviors:

Another aspect of the disclosure relates to a self-contained robotic apparatus. The apparatus may comprise a platform, a first sensor, a second sensor, a memory component, and one or more physical processors. **The platform may comprise a motor.**

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.

The user may select which **actuators** are controlled by a particular behavior. For example, navigation tasks may not need to learn to control the gripper

The higher level inputs may be provided by a sensorimotor control process, which may control the desired movement of the robot (e.g., output of the **motor control system**) and/or the desired focus point of the camera (e.g., output of the vision control system).

E.g., Ex. 1005, 3:27-31, Abstract, 24:66-25:2, 44:31-35. A POSITA would understand these disclosures to equate

- e. **[1.5] 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 physical device are learned in a learning process that includes operating the first physical device at least partially by a user;**

Limitation	Claimed Features
[1.5]	portion of 1st DP or 1st Instructions learned with user operation

79. Grotmol discloses that a human² can operate the robot during a training process and “[t]rained 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.” *Id.*, 33:59-62. In both the “listen mode” and “override mode,” Grotmol teaches that the robot will listen to commands sent by a user “and try to remember it for the next time the same situation occurs.” *Id.*, 23:20-21; *see also id.*, 24:57-61 (“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).”), 34:50-55 (“In some implementations of supervised training of robots, control instructions (also referred to as corrections) produced by the ***trainer*** (e.g., human) may be combined with control instructions produced by the robot controller instructions (predictions).”). Grotmol describes this training as a learning process. *E.g., id.*, 24:17-20 (“Apparatus and methods for hierarchical ***learning*** are disclosed herein, in accordance with one or more implementations. In ***supervised learning***, a user may

² The challenged claims of the ’344 patent recite “user(s),” which in my analysis I have considered to include humans. However, the ’449 patent does not limit the term “user” to only humans. ’449 patent, 68:64-66 (“User 50 (also referred to simply as user or other suitable name or reference) comprises a human user or non-human user.”).

train simple tasks by demonstrating example tasks to the system.”), claim 12 (reciting “one or more physical processors configured to operate a random k-nearest neighbors *learning process*”).

- f. [1.6] an optical camera that captures digital pictures;
and

Limitation [1.6]	Claimed Features
	optical camera

80. Grotmol discloses that the robot includes a sensor comprising an optical camera that captures digital pictures:

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

FIG. 1 depicts a mobile robotic apparatus that may be configured with an adaptive controller in accordance with one or more implementations of e.g., the learning apparatuses illustrated in FIGS. 4A-4B, *infra*. The robotic apparatus 160 may comprise a sensor component 166. The sensor component 166 may be characterized by an aperture/field of view 168 (e.g., an extent of the observable world that may be captured by the sensor at a given moment). The sensor component 166 may provide

information associated with objects within the field-of-view 168. In one or more implementations, such as object recognition, and/or obstacle avoidance, the output provided by the sensor component 166 may comprise a stream of pixel values associated with one or more digital images.

In some implementations of robotic navigation in an arbitrary environment, **the sensor component 166 may comprise a camera** configured to provide an output comprising a plurality of digital image frames refreshed at, e.g., 25 Hz frame rate. The sensor output may be processed by a learning controller, e.g., as illustrated and described with respect to FIG. 4A.

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).

One or more implementations described herein may provide a mechanism for enabling a robot to learn navigating a target trajectory while reducing deviation from a target path. In some implementations, the robot may comprise a robotic vehicle (e.g., 160 in FIG. 1) **comprising a camera** (e.g., 166). The robot may comprise a controller configured to store a sequence of images obtained during training phase. In some implementations, the

images may be stored in an ordered buffer in a non-volatile memory (e.g., memory 814 in FIG. 8).

FIG. 20A illustrates exemplary images for use with training of path navigation. Images 2000, 2010 may be obtained with **a variety of sensors, e.g., a video camera, IR camera, RADAR, SONAR, LiDAR, ultrasonic imager, and/or other device.** In some implementations, the camera may be spanning the entire horizontal view field of 360 degrees. A 360 degree horizontal view field may provide additional information about robot's environment and enable cyclical shift of images.

In one or more implementations, images 2000, 2010 may be obtained with **a camera 166 mounted on a robotic vehicle 160** of FIG. 1. Images 2000, 2010 may comprise representations of visual field that may be obtained at times $t_1 < t_2$, respectively. The image 2000 may comprise representations one or more objects, e.g., a tree 2002, a rectangle, a ball 2008, and/or other objects.

During 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.

Ex. 1005, Abstract, 8:57-9:3, 10:15-26, 38:33-58, 39:13-22.

g. [1.7] an artificial intelligence unit that:

Limitation [1.6]	Claimed Features
	artificial intelligence unit

81. The remaining limitations of claim 1 require the use of an “artificial intelligence unit” that must perform the following functions:

- (1) receive a new picture (claim [1.7.1] “receives a new one or more digital pictures from the optical camera”);
- (2) anticipate an instruction set based on matching the new picture with a picture stored in memory (*see infra* claim [1.7.2]); and
- (3) cause the processor to execute the instruction set in response to determining a match between the new picture and picture stored in memory (*see* ’449 patent, claim [1.7.3]).

82. I understand that during prosecution, the applicant for the ’449 patent stated that the artificial intelligence unit “is specifically described in the specification to include hardware, programs, or a combination thereof.” Ex. 1002, 99. The applicant further stated that the artificial intelligence unit may “be implemented entirely or in part as a computer program.” *Id.* at 100. The ’449 patent itself gives a

broad description of the term “artificial intelligence unit” as a generic piece of hardware or software:

In certain embodiments, the artificial intelligence unit is coupled to the picture capturing apparatus. In further embodiments, the artificial intelligence unit is coupled to the memory unit. In further embodiments, **the artificial intelligence unit is part of, operating on, or coupled to the processor circuit.** In further embodiments, the system further comprises: a second processor circuit, wherein the artificial intelligence unit is part of, operating on, or coupled to the second processor circuit. In further embodiments, the artificial intelligence unit is part of, operating on, or coupled to a remote computing device, the remote computing device coupled to the processor circuit via a network. In further embodiments, the artificial intelligence unit includes a circuit, a computing apparatus, or a computing system attachable to the processor circuit. In further embodiments, the artificial intelligence unit includes a circuit, a computing apparatus, or a computing system attachable to the device. In further embodiments, the artificial intelligence unit is attachable to an application for operating the device, the application running on the processor circuit. In further embodiments, the artificial intelligence unit includes a circuit, a computing apparatus, or a computing system built into the processor circuit. In further embodiments, the artificial intelligence unit includes a circuit, a computing apparatus, or a computing system built into the device. In further embodiments, the artificial intelligence unit is

built into an application for operating the device, the application running on the processor circuit. In further embodiments, the artificial intelligence unit is provided as a feature of the processor circuit. In further embodiments, **the artificial intelligence unit is provided as a feature of an application running on the processor circuit.** In further embodiments, the artificial intelligence unit is provided as a feature of the device. In further embodiments, the artificial intelligence unit is further configured to: take control from, share control with, or release control to the processor circuit. In further embodiments, the artificial intelligence unit is further configured to: take control from, share control with, or release control to an application or an object of the application, the application running on the processor circuit. In further embodiments, the artificial intelligence unit is further configured to: take control from, share control with, or release control to a user or a system.

'449 patent, 2:31-3:6; *see also id.*, claim 11 (reciting “wherein the artificial intelligence unit includes . . . an application operating on the one or more processor circuits”).

83. In my opinion, a POSITA would understand Grotmol to disclose an “artificial intelligence unit” as a part of or operating on its processor circuit. *See, e.g.*, Ex. 1005, claim 12 (reciting “one or more physical processors configured to operate a random k-nearest neighbors learning process to determine a motor control indication”).

- h. [1.7.1] receives a new one or more digital pictures from the optical camera;

Limitation	Claimed Features
[1.7.1]	receipt of 2nd DP from optical camera

84. A POSITA would understand Grotmol to disclose an artificial intelligence unit—i.e., its processor or a feature of an application running thereon—that receives new digital pictures from the robot’s camera before comparing them with stored digital pictures:

In some implementations, the robot may comprise a robotic vehicle (e.g., 160 in FIG.1) comprising a camera (e.g., 166). The robot may comprise a controller configured to store a **sequence of images** obtained during training phase.

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.

Ex. 1005, 38:36-39, 39:23-28.

85. As depicted in Figure 8 of Grotmol below, the “processing” functionality of Grotmol is connected to the “sensory” functionality, which may include camera input:

[T]he output provided by the sensor component 166 may comprise a stream of pixel values associated with one or more digital images.

In one or more implementations, images 2000, 2010 may be obtained with a camera 166 mounted on a robotic vehicle 160 of FIG. 1. Images 2000, 2010 may comprise representations of visual field that may be obtained at times $t_1 < t_2$, respectively. The image 2000 may comprise representations one or more objects, e.g., a tree 2002, a rectangle, a ball 2008, and/or other objects.




FIG. 8

Id., 9:1-3, 38:52-58, Fig. 8 (annotated).

86. A POSITA would also appreciate that both Grotmol and the '449 patent disclose that the digital picture data may come from, for example, a CCD image sensor in the form of RGB-encoded pixels. *Compare* '449 patent, 69:58-64, 112:26-32 with Ex. 1005, 9:8-22.

- i. **[1.7.2] anticipates the first one or more instruction sets for operating the first physical device based on at least partial match between the new one or more digital pictures and the first one or more digital pictures, wherein the anticipates includes at least one of:**

determining that a number of at least partially matching portions of the new one or more digital pictures and portions of the first one or more digital pictures exceeds a threshold number, or

determining that a percentage of at least partially matching portions of the new one or more digital pictures and portions of the first one or more digital pictures exceeds a threshold percentage; and

Limitation	Claimed Features
[1.7.2]	Anticipation of 1st Instructions by matching 1st DP and 2nd DP

87. Grotmol discloses anticipating motor commands (instruction sets for operating the first device) based on a partial match between the new image (new one or more digital pictures) and a stored image (first one or more digital pictures). *Id.*, 43:6-9 (“[M]otor 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”); *see also* Abstract (“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.”).

88. A POSITA would have understood Grotmol's disclosures for determining the best matching stored image by calculating a similarity measure between a current image and a stored image, and then determining whether that similarity measure breaches a threshold to correspond to the claimed "at least partial match between the new one or more digital pictures and the first one or more digital pictures."

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

[R]esponsive to the similarity measure breaching a threshold, selecting one or more training sets from the plurality of training feature sets..

Ex. 1005, 4:30-39, claim 1; *see also id.*, 39:13-41, 40:64-41:18. Grotmol clarifies that the term "feature" as used in the paragraph excerpted above can include values of image pixels or other image qualities. *Id.*, 28:19-24 ("[T]he 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 . . .”). Grotmol also clarifies that the term “training set” used in the paragraph excerpted above includes the video frame instances that a POSITA would have understood correspond to the claimed digital pictures. *Id.*, 3:36-39 (“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”).

89. A POSITA would have appreciated that both Grotmol and the ’449 patent discuss pixel-wise difference analysis between two digital pictures to determine a “partial match.” ’449 patent, 21:12-17 (discussing “portion”), 67:27-31 (discussing “matching”), 112:23-54 (discussing RGB pixel matching); Ex. 1005, 3:14-16 (“The first input type may comprise a digital image comprising a plurality of pixel values.”), 41:8-18 (“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.”).

- j. [1.7.3] causes the one or more processor circuits to execute the first one or more instruction sets for operating the first physical device, wherein the causes is performed in response to the anticipates of the artificial intelligence unit, and wherein the first physical device or a second physical device autonomously performs one or more operations defined by the first one or more instruction sets for operating the first physical device.

Limitation [1.6]	Claimed Features
	executing 1st instructions on 1st or 2nd device autonomously

90. As I discussed above for limitation [1.7.2], Grotmol discloses an artificial intelligence unit that anticipates the learned instruction sets for operating the device by comparing a new picture with a previously learned picture at least by matching portions of the pictures and determining whether those matching portions exceed a threshold number.

91. In my opinion, Grotmol further discloses that, in response to the “anticipates,” the system “causes the one or more processor circuits to execute the first one or more instruction sets for operating the first physical device.” Grotmol describes that after it is determined that the similarity measure between the new image and image from the training set breaches a threshold, one or more potential control signals corresponding to the image from the training set that were obtained

during a training operation of the robot (the instruction sets correlated with the first digital picture) are determined and executed:

The method may comprise, responsive to the similarity measure breaching a threshold, selecting one or more training sets from the plurality of training sets. The method may comprise determining one or more potential control signals for the robot. Individual ones of the one or more potential control signals may be associated with a corresponding training set of the plurality of training sets.

In some implementations, the control signal may be configured **to cause the robot to execute the action.**

Ex. 1005, 2:35-42, 3:12-13; *see also id.*, 43:5-9 (“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.”).

92. The ’449 patent similarly describes that this step may comprise simply causing an application to execute an instruction. ‘449 patent, 10:52-58 (“In certain embodiments, the causing the processor circuit to execute the one or more instruction sets for operating the device correlated with the first digital picture includes causing an application for operating the device to execute the one or more

instruction sets for operating the device correlated with the first digital picture, the application running on the processor circuit.”).

93. Grotmol discloses autonomous performance of the operations defined by the first instruction set for operating the device. Ex. 1005, 34:56-59 (“In some implementations, the trainer may be provided with the control of the robot during training. Upon completion of the training, the robot may be configured to operate autonomously.”); Abstract (“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.”), 43:5-9 (“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.”).

94. Grotmol also discloses that training sets and training configurations, which include pictures and associated instruction sets for operating the first device, can be stored or saved, and then loaded into the robot to provide the trained behavior *Id.*, 33:27-30 (“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.”). Grotmol explains that the saved training configuration can also be loaded into robots other than the one that was operated during the learning. *Id.*, 33:30-32 (“[T]he trained configuration may be loaded to ***one or more other robots*** in order to provide learned behaviors”). A POSITA would have understood

this disclosure to mean that the robot autonomously performing the operations defined by the first instruction set can be the robot that was used during the training mode (*i.e.*, first physical device), or it can be a different robot (*i.e.*, second physical device) that is loaded with the training set of the first.

95. Accordingly, in my opinion, Grotmol discloses an artificial intelligence unit that causes a processor circuit to execute the anticipated instruction, and a first physical device or a second physical device autonomously performs operations defined by the first instruction.

2. Claim 6

Claim 6	Additional Features to Claim 1
	+ Autonomous performance of 1st instructions by 2nd device
	+ Modification of 1st instructions for 2nd device

- a. **[6.1] The system of claim 1, wherein the second physical device autonomously performs one or more operations defined by the first one or more instruction sets for operating the first physical device,**

Limitation [6.1]	Claimed Features
	Autonomous performance of 1st instructions by 2nd device

96. As I discussed above for limitation [1.7.3], Grotmol discloses that training sets and training configurations, which include pictures and associated instruction sets for operating a first device, can be stored or saved, and then loaded

into the robot to provide the trained behavior Ex. 1005, 33:27-32 (“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.”). Grotmol explains that the saved training configuration can also be loaded into robots other than the one that was operated during the learning. *Id.*, 33:30-32 (“[T]he trained configuration may be loaded to ***one or more other robots*** in order to provide learned behaviors”). Thus, a POSITA would have understood Grotmol to explicitly disclose a second physical device (“one or more other robots”) that can autonomously perform one or more operations defined by the first one or more instruction sets for operating the first physical device (training configuration from first robot (e.g., training sets with images correlated to instruction sets) that is loaded into a second one or more other robots to “provide learned behaviors.”

- b. **[6.2] and wherein the first one or more instruction sets for operating the first physical device are modified so that the second physical device can autonomously perform one or more operations defined by the first one or more instruction sets for operating the first physical device.**

Limitation	Claimed Features
[6.2]	Modification of 1st instructions for 2nd device

97. Grotmol discloses that the robot may adjust its trajectory to compensate for, “mismatch[es]” in position or orientation, and further discloses modifying the motor command where a current image is shifted relative to a stored image:

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**.

During operation when the robot may follow a target trajectory the shift amount determined using the image matching process may be close to 0 (this configuration may be referred to as “the robot stays on track”). 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. In some implementations, the target error level may comprise zero displacement.

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.

Ex. 1005, 42:46-43:3, Abstract.

98. Grotmol further explains that motor commands corresponding to a given image and time may be combined with the output from a PID controller to “stabilize” the robot’s operation:

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.

Id., 43:4-11.

99. Accordingly, a POSITA would understand that when a training configuration is loaded into a second physical device (“one or more other robots” –

see supra limitation [6.1]), and the second robot encounters “position mismatch” or “orientation mismatch” from the first robot, the instruction set is modified to account for this mismatch between the first and second robot’s position or orientation. Thus, in my opinion, Grotmol discloses that the anticipated first instruction for operating the first physical device is modified so the second physical device can perform the operation defined by the instruction.

3. **Claim 7: The system of claim 1, wherein the portions of the new one or more digital pictures include portions of the new one or more digital pictures that depict recognized objects and the portions of the first one or more digital pictures include portions of the first one or more digital pictures that depict recognized objects.**

Claim 7	Additional Features to Claim 1
	+ 1 st and 2 nd DPs depict recognized objects

100. Grotmol discloses that the images taken by or stored on the robot include recognized objects. For example, Grotmol discloses that images 2000 and 2010, shown below in Figure 20A, may “comprise representations of one or more objects, e.g., a tree 2002, a rectangle, a ball 2008, and/or other objects.” Ex. 1005, 38:56-58.




FIG. 20A
Ex. 1005, Fig. 20A

101. A POSITA would have appreciated the similarity between these disclosures and the examples of objects in pictures from the '449 patent, such as “car, pedestrian, door, building, animal, person, etc.” '449 patent, 88:9-20. The '449 patent further describes various approaches to object recognition, such as “identifying and/or analyzing object features such as lines, edges, ridges, corners, blobs, regions, and/or their relative positions, sizes, shapes, etc.” *Id.*, 88:10-12.

102. Grotmol further discloses that the sensor used to capture these images “may provide information associated with objects within the field-of-view In one or more implementations, such as *object recognition*, and/or obstacle avoidance, the output provided by the sensor component 166 may comprise a stream of pixel values associated with one or more digital images.” Ex. 1005, 8:65-9:3; *see, e.g., id.*, 9:41-52 (“In one or more implementations, object recognition and/or classification

may be implemented”). Grotmol, like the ’449 Patent, further includes edge detection. *Id.*, 8:60-63 (“[T]he term feature may refer to a representation of an object edge”).

103. Grotmol describes that the autonomous operation of a robot may be based in part on recognized objects and their characteristics. *E.g.*, *id.*, 10:26-34 (noting that “tasks of the robot may be configured based on a context” such as “state of the environment (e.g., object size, location)”), 24:31-25:5 (describing embodiment in which a robot is trained to “find an object, grasp it, and bring it back to base”).

4. Claim 8

Claim 8	Additional Features to Claim 1
	+ 1st device and 2nd device are robots or vehicles + Processor circuits are contained in the 1st, 2nd, or 3rd device

- a. **[8.1] The system of claim 1, wherein the first physical device includes a robot, a vehicle, an appliance, an electronic device, or a mechanical machine and**

Limitation [8.1]	Claimed Features
	1st device is a robot or vehicle

104. Grotmol discloses embodiments of its devices as both robots and autonomous vehicles:

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.

One aspect of the disclosure relates to a method of determining a control signal for a **robot**.

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 robot** may comprise a robotic vehicle (e.g., 160 in FIG.1)

Implementations of the principles of the disclosure may be applicable to a wide assortment of applications including computer-human interaction (e.g., recognition of gestures, voice, posture, face, and/or other interactions), controlling processes (e.g., processes associated with **an industrial robot, autonomous and other vehicles**, and/or other processes), augmented reality applications, access control (e.g., opening a

door based on a gesture, opening an access way based on detection of an authorized person), detecting events (e.g., for visual surveillance or people or animal counting, tracking)

Ex. 1005, Abstract, 2:13-14, 33:27-30, 38:36-37, 46:16-26; *see supra* claim 1.

- b. [8.2] the second physical device includes a robot, a vehicle, an appliance, an electronic device, or a mechanical machine, and

Limitation [8.2]	Claimed Features
	2nd device is a robot or vehicle

105. Grotmol discloses embodiments in which a “second physical device” is also a robot or autonomous vehicle. Grotmol discloses that training sets and training configurations can be stored or saved, and then loaded into the robot to provide the trained behavior. *Id.*, 33:27-32 (“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.”). Grotmol explains that the saved training configuration can also be loaded into robots other than the one that did the learning. *Id.*, 33:30-32 (“the trained configuration may be loaded to **one or more other robots** in order to provide learned behaviors”). Thus, a POSITA would have understood Grotmol to disclose a second device (“one or more other robots”) that can autonomously perform one or more

operations defined by one or more instruction sets for operating the first device that is loaded into a second one or more other robots to “provide learned behaviors.”

- c. **[8.3] wherein the one or more processor circuits are included in at least one of: the first physical device, the second physical device, or a third device.**

Limitation	Claimed Features
[8.3]	Processors circuits are contained in the 1st, 2nd, or 3rd device

106. Grotmol discloses that the “one or more processor circuits” (*see supra* claim [1.2]) are included in a first, second, or third device. In my opinion, a POSITA would understand that the “one or more processor circuits” may be included in any of the robots described in Grotmol (i.e., the “first,” the “second,” or a “third” device). *E.g.*, Ex. 1005, 3:27-39 (“Another aspect of the disclosure relates to a self-contained robotic apparatus. The apparatus may comprise a platform, a first sensor, a second sensor, a memory component, and *one or more physical processors.*”); *see also id.*, claim 12 (reciting “[a] self-contained robotic apparatus, the apparatus comprising . . . one or more physical processors configured to operate a random k-nearest neighbors learning process to determine a motor control indication”).

5. **Claim 10: The system of claim 1, wherein the first one or more digital pictures includes a one or more still digital pictures or one or more motion digital pictures and the new one or more digital pictures include one or**

more still digital pictures or one or more motion digital pictures.

Claim 10	Additional Features to Claim 1
	+ 1st and 2nd DPs include still DPs or motion DPs

107. Grotmol discloses that digital pictures may comprise single images and/or image frames from a video:

The input signal may comprise a **sequence of images and/or image frames**. The sequence of images and/or image frame may be received from a CCD camera via a receiver apparatus and/or downloaded from a file. The image may comprise a two-dimensional matrix of red/green/blue (RGB) values **refreshed at a 25 Hz frame rate**. It will be appreciated by those skilled in the arts that the above image parameters are merely exemplary, and many other image representations (e.g., bitmap, CMYK, HSV, HSL, grayscale, and/or other representations) and/or frame rates are equally useful with the present disclosure.

Images 2000, 2010 may be obtained with a variety of sensors, e.g., **a video camera, IR camera, RADAR, SONAR, LiDAR, ultrasonic imager, and/or other device**.

A self-contained robotic apparatus, the apparatus comprising...
a first sensor component configured to provide a signal

configured to convey a **video frame comprising a plurality of pixels...**

Ex. 1005, 9:12-22, 38:44-47, claim 12. In my opinion, a POSITA would understand these disclosures of standalone or sequentially captured images to correspond to the claimed “still digital pictures” and “motion digital pictures.”

6. Claim 12

Claim 12	Additional Features to Claim 1
	+ 1st and 2nd knowledge cells • 1st knowledge cell includes 1st DP correlated with 1st instructions • 2nd knowledge cell includes 2nd DP correlated with 2nd instructions + directed connection between 1st and 2nd knowledge cells indicating temporal order

- a. [12.1] The system of claim 1, wherein the memory further stores at least a knowledge structure that includes: a first knowledge cell including the first one or more digital pictures correlated with the first one or more instruction sets for operating the first physical device, and
- a second knowledge cell including a second one or more digital pictures correlated with a second one or more instruction sets for operating the first physical device, and

Limitation [12.1]	Claimed Features
	1st and 2nd knowledge cells • 1st knowledge cell includes 1st DP correlated with 1st instructions • 2nd knowledge cell includes 2nd DP correlated with 2nd instructions

108. Although Grotmol does not explicitly use the term “knowledge cell,” a POSITA would have understood Grotmol to disclose the claimed “knowledge structure” that comprises “a first knowledge cell” and a “second knowledge cell” for the reasons I set forth below.

109. The ’449 patent describes that a “knowledge cell” may “compris[e] the first digital picture correlated with the one or more instruction sets for operating the device.” ’449 patent, 6:64-7:1.


Knowledge Cell 800 comprises the functionality for storing one or more Digital Pictures 525 correlated with any Instruction Sets

526 and/or Extra Info 527. Knowledge Cell 800 includes a unit of knowledge of how Device 98 operated in a visual surrounding.

Once created or generated, Knowledge Cells 800 can be used in/as neurons, nodes, vertices, or other elements in any of the data structures or arrangements (i.e. neural networks, graphs, sequences, etc.) used for storing the knowledge of Device's 98 operation in various visual surroundings, thereby facilitating learning functionalities herein.

Knowledge Cell 800 may include any data structure or arrangement that can facilitate such storing.

Id., 90:36-40, 90:40-46, 90:66-67.



'449 patent, Fig. 13 (excerpt)

110. The '449 patent further describes that “[k]nowledge cells may be connected, interrelated, or interlinked into knowledge structures using statistical, artificial intelligence, machine learning, and/or other models or techniques.” ’449 patent, 143:65-144:1. “In some embodiments, Knowledge Structuring Unit 520 may correlate a Digital Picture 525 with one or more temporally corresponding Instruction Sets 526 and/or Extra Info 527.” *Id.*, 91:22-25.

111. Grotmol describes storing sequences of images and the corresponding instruction sets (i.e., a sequence of knowledge cells). Ex. 1005, Abstract (“Motor commands associated with a given image may be stored.”); *see also id.*, 38:21-42. Grotmol further discloses that its training data is stored as a “sequence” and may be stored in “an ordered buffer in a non-volatile memory.” *Id.*, 38:39-42. I understand that an “ordered buffer” is a region of memory that is allocated for storing data in a specific order or sequence. Grotmol describes that the training buffer may be ordered according to when the images were acquired. *See id.*, 39:65-40:4; *see also id.*, 39:17-19 (“Motor commands issued by a trainer corresponding to time instances when the images are acquired may be stored.”). A POSITA would have understood the training data stored in memory in a sequence corresponds to the claimed “knowledge cells” in a “knowledge structure.”

112. Furthermore, as I explained with respect to limitations [1.3] and [1.5] of claim 1, Grotmol discloses that multiple images may be acquired during a training

process, with images respectively stored in association with corresponding motor commands issued at the time the image was acquired. Thus, a POSITA would have understood Grotmol to disclose first, second, third, etc. knowledge cells, where the knowledge cells include digital pictures correlated with instruction sets.

- b. **[12.2] wherein the first knowledge cell is connected with the second knowledge cell by a directed connection that indicates temporal order of the first knowledge cell and the second knowledge cell**

Limitation	Claimed Features
[12.1]	directed connection between 1st and 2nd knowledge cells indicating temporal order

113. The '449 patent does not define or discuss “directed connection” in its specification. The '449 patent does describe, however, that knowledge cells can be “connected” to other knowledge cells to form a knowledgebase:

Knowledge cells may be connected, interrelated, or interlinked into knowledge structures using statistical, artificial intelligence, machine learning, and/or other models or techniques. Such interconnected or interrelated knowledge cells can be used for enabling autonomous device operation. The interconnected or interrelated knowledge cells may be stored or organized into a **knowledgebase** (i.e. Knowledgebase 530, etc.). . . . In further embodiments, knowledgebase may be or include a sequence (i.e. Sequence 533, etc.). In further embodiments, knowledgebase may be or include a collection of knowledge cells (i.e. Collection of Knowledge Cells 530 d, etc.). In general,

knowledgebase may be or include any data structure or arrangement, and/or repository capable of storing the knowledge of a device's operation in various visual surroundings.

'449 patent, 143:65-144:18. As I discussed above in relation to limitation [12.1], the '449 patent discloses that these knowledge cells contain the claimed correlated digital pictures and instruction sets.

114. In my opinion, Grotmol discloses these directed “connections” between knowledge cells in the knowledgebase. Grotmol describes storing sequences of images and the corresponding instruction sets, which a POSITA would have understood to correspond to a sequence of knowledge cells. Ex. 1005, Abstract (“Motor commands associated with a given image may be stored.”); 38:21-41 (“The robot may comprise a controller configured to store a sequence of images obtained during training phase. In some implementations, the images may be stored in an ordered buffer in a non-volatile memory...”). As I noted above, I understand that an ordered buffer is a region of memory that is allocated for storing data in a specific order or sequence. A POSITA would have understood Grotmol to describe that the training buffer may be ordered temporally, according to when the images were acquired, due to its disclosures of images in the training buffer that precede and follow each other in time:

The given image may be compared to images of the primary particle set. In some implementations, the primary particle set

may comprise a previously used image IO (e.g., 2000 in FIG. 20A), one or more (e.g., 2) images **preceding the IO image in time** in the training buffer, and one or more (e.g., 5) images **following the IO image in time** in the training buffer.

Motor commands issued by a trainer corresponding to time instances when the images are acquired may be stored.

Ex. 1005, 39:65-40:4, 39:17-19. Thus, a POSITA would understand that Grotmol describes connected knowledge cells because it describes knowledge cells stored in temporal order – i.e., order one after another based on when the image-instruction correlation was established.

7. Claim 13

115. Independent claim 13 is a system claim that corresponds to system claim 1. My analysis above explaining how Grotmol discloses the limitations of claim 1 applies to the limitations of claim 13 as set forth in the chart below.

Claim 13	Disclosure of Grotmol
[13.1] A non-transitory machine readable medium having a stored thereon, instructions that when executed by one or more processor circuits cause the one or more processor circuits to perform operations comprising:	To the extent the preamble is limiting, Grotmol discloses “[a] non-transitory machine readable medium having a stored thereon, instructions that when executed by one or more processor circuits cause the one or more processor circuits to perform operations.” <i>E.g.</i> , Ex. 1005, claim 17 (reciting a substantially similar preamble); <i>see supra</i> limitation [1.2].

[13.2] accessing a memory that stores at least a first one or more digital pictures correlated with a first one or more instructions sets for operating a first physical device, wherein	<i>See analysis for limitation [1.3] above.</i>
[13.3] the first physical device includes an actuator for moving at least a portion of the first physical device, and wherein	<i>See analysis for limitation [1.4] above.</i>
[13.4] 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 physical device are learned in a learning process that includes operating the first physical device at least partially by a user;	<i>See analysis for limitation [1.5] above.</i>
[13.5] receiving a new one or more digital pictures from an optical camera;	<i>See analysis for limitations [1.6] and [1.7.1] above.</i>
[13.6] anticipating the first one or more instruction sets for operating the first physical device based on at least partial match between the new one or more digital pictures and the first one or more digital pictures, wherein	<i>See analysis for limitation [1.7.2] above.</i>
[13.7] the anticipating includes at least one of: determining that a number of at least partially matching portions of the new one or more digital pictures and portions of the first one or more digital pictures exceeds a threshold number, or determining that a percentage of at least partially matching portions of the new one or more digital pictures and portions of the first one or more digital pictures exceeds a threshold percentage; and	<i>See analysis for limitation [1.7.2] above.</i>
[13.8] causing the one or more processor circuits or another one or more processor circuits to execute the first one or more instruction sets for operating the first physical device, the	<i>See analysis for limitation [1.7.3] above.</i>

causing performed in response to the anticipating, wherein	
[13.9] the first physical device or a second physical device autonomously performs one or more operations defined by the first one or more instruction sets for operating the first physical device.	See analysis for limitation [1.8] above.

8. **Claim 14: The non-transitory machine readable medium of claim 13, wherein the second physical device autonomously performs one or more operations defined by the first one or more instruction sets for operating the first physical device, and wherein the first one or more instruction sets for operating the first physical device are modified so that the second physical device can autonomously perform one or more operations defined by the first one or more instruction sets for operating the first physical device.**

Claim 14	Additional Features to Claim 13
	+ Autonomous performance of 1st instructions by 2nd device + Modification of 1st instructions for 2nd device

116. Claim 14 requires that the second physical device autonomously performs the operation defined by the anticipated instruction, and the instruction is modified so that the second physical device can perform the operation. Grotmol discloses claim 14 for the same reasons I discussed above in Section VIII.B.2, *supra*, for claim 6.

9. **Claim 15: The non-transitory machine readable medium of claim 13, wherein the portions of the new**

one or more digital pictures include portions of the new one or more digital pictures that depict recognized objects and the portions of the first one or more digital pictures include portions of the first one or more digital pictures that depict recognized objects.

Claim 15	Additional Features to Claim 13
	+ 1st and 2nd DPs depict recognized objects

117. Claim 15 requires that the new pictures and stored pictures include portions that depict recognized objects. Grotmol discloses claim 15 for the same reasons I discussed above in Section VIII.B.3, *supra*, for claim 7.

- 10. Claim 16: The non-transitory machine readable medium of claim 13, wherein the memory further stores at least a knowledge structure that includes: a first knowledge cell including the first one or more digital pictures correlated with the first one or more instruction sets for operating the first physical device, and a second knowledge cell including a second one or more digital pictures correlated with a second one or more instruction sets for operating the first physical device, and wherein the first knowledge cell is connected with the second knowledge cell by a directed connection that indicates temporal order of the first knowledge cell and the second knowledge cell.**

Claim 16	Additional Features to Claim 13
-----------------	--

	+ 1st and 2nd knowledge cells • 1st knowledge cell includes 1st DP correlated with 1st instructions • 2nd knowledge cell includes 2nd DP correlated with 2nd instructions + directed connection between 1st and 2nd knowledge cells indicating temporal order
--	---

118. Claim 16 requires the memory stores a knowledge structure that includes a first knowledge cell including a digital picture correlated with an instruction, a second knowledge cell including a digital picture correlated with an instruction, and the first knowledge cell and second knowledge cell are connected by a directed connection that includes temporal order of the cells. Grotmol discloses claim 16 for the same reasons I discussed above in Section VIII.B.6, *supra*, for claim 12.

11. Claim 17

119. Independent claim 13 is a system claim that corresponds to system claim 1. My analysis above explaining how Grotmol discloses the limitations of claim 1 applies to the limitations of claim 13 as set forth in the chart below. In my opinion, Grotmol anticipates claim 17.

Claim 17	Disclosure of Grotmol
[17.1] A method comprising:	To the extent the preamble is limiting, Grotmol discloses a method

	implemented using a “self-contained robotic apparatus” that includes “one or more physical processors.” Ex. 1005, 3:27-30; <i>see also id.</i> , 2:14-18, 7:30-42.
[17.2] (a) accessing a memory that stores at least a first one or more digital pictures correlated with a first one or more instructions sets for operating a first physical device, wherein	<i>See</i> analysis for limitation [1.3] above.
[17.3] the first physical device includes an actuator for moving at least a portion of the first physical device, and wherein	<i>See</i> analysis for limitation [1.4] above.
[17.4] 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 physical device are learned in a learning process that includes operating the first physical device at least partially by a user, the accessing of (a) performed by one or more processor circuits;	<i>See</i> analysis for limitation [1.5] above.
[17.5] (b) receiving a new one or more digital pictures from an optical camera, the receiving of (b) performed by the one or more processor circuits;	<i>See</i> analysis for limitations [1.6] and [1.7.1] above.
[17.6] (c) anticipating the first one or more instruction sets for operating the first physical device based on at least partial match between the new one or more digital pictures and the first one or more digital pictures, wherein the anticipating [of (c)] includes at least one of: determining that a number of at least partially matching portions of the new one or more digital pictures and portions of the first one or more digital pictures exceeds a threshold number, or determining that a percentage of at least partially matching portions of the new one or more digital pictures and portions	<i>See</i> analysis for limitation [1.7.2] above.

of the first one or more digital pictures exceeds a threshold percentage, the anticipating of (c) performed by the one or more processor circuits;	
[17.7] (d) executing the first one or more instruction sets for operating the first physical device, the executing of (d) performed by the one or more processor circuits or by another one or more processor circuits in response to the anticipating of (c); and	<i>See analysis for limitation [1.7.3] above.</i>
[17.8] (e) autonomously performing, by the first physical device or by a second physical device, one or more operations defined by the first one or more instructions ets [sic] for operating the first physical device.	<i>See analysis for limitation [1.8] above.</i>

- 12. Claim 18: The method of claim 17, wherein the autonomously performing of (e) is performed by the second physical device, and wherein the first one or more instruction sets for operating the first physical device are modified so that the second physical device can autonomously perform one or more operations defined by the first one or more instruction sets for operating the first physical device.**

Claim 18	Additional Features to Claim 17
	+ Autonomous performance of 1st instructions by 2nd device + Modification of 1st instructions for 2nd device

120. Claim 18 requires that the second physical device autonomously performs the operation defined by the anticipated instruction, and the instruction is modified so that the second physical device can perform the operation. Grotmol

discloses claim 18 for the same reasons I discussed above in Section VIII.B.2, *supra*, for claim 6.

- 13. Claim 19: The method of claim 17, wherein the portions of the new one or more digital pictures include portions of the new one or more digital pictures that depict recognized objects and the portions of the first one or more digital pictures include portions of the first one or more digital pictures that depict recognized objects.**

Claim 19	Additional Features to Claim 17
	+ 1st and 2nd DPs depict recognized objects

121. Claim 19 requires that the new pictures and stored pictures include portions that depict recognized objects. Grotmol discloses claim 19 for the same reasons I discussed above in Section VIII.B.3, *supra*, for claim 7.

- 14. Claim 20: The method of claim 17, wherein the memory further stores at least a knowledge structure that includes: a first knowledge cell including the first one or more digital pictures correlated with the first one or more instruction sets for operating the first physical device, and a second knowledge cell including a second one or more digital pictures correlated with a second one or more instruction sets for operating the first physical device, and wherein the first knowledge cell is connected with the second knowledge cell by a directed connection that indicates temporal order of the first knowledge cell and the second knowledge cell.**

Claim 20	Additional Features to Claim 17
-----------------	--

	+ 1st and 2nd knowledge cells • 1st knowledge cell includes 1st DP correlated with 1st instructions • 2nd knowledge cell includes 2nd DP correlated with 2nd instructions + directed connection between 1st and 2nd knowledge cells indicating temporal order
--	---

122. Claim 20 requires the memory stores a knowledge structure that includes a first knowledge cell including a digital picture correlated with an instruction, a second knowledge cell including a digital picture correlated with an instruction, and the first knowledge cell and second knowledge cell are connected by a directed connection that includes temporal order of the cells. Grotmol discloses claim 20 for the same reasons I discussed above in Section VIII.B.6, *supra*, for claim 12.

C. Ground 2: Claims 3, 5, and 11 Are Obvious Based on Grotmol in View of the Knowledge of a POSITA

123. As set forth below, Grotmol renders claims 3, 5, and 11 obvious in view of the knowledge of a POSITA.

1. Claim 3

Claim 3	Additional Features to Claim 1
----------------	---------------------------------------

	+ 1st instructions obtained from processor circuit(s) via a hardwired connection + Processor circuit(s) used to operate 1st device
--	---

- a. **[3.1] The system of claim 1, wherein the first one or more instruction sets for operating the first physical device are obtained, via a hardwired connection, from at least one portion of the one or more processor circuits or from at least one portion of another one or more processor circuits, and**

Limitation	Claimed Features
[3.1]	1st instructions obtained from processor circuit(s) via a hardwired connection

124. Limitation [3.1] recites that the “first one or more instruction sets” are “obtained,” but the claim does not identify what is doing the obtaining. Claim 1 recites that the instructions are stored in a memory. *See supra* limitation [1.3]. For the purposes of this Declaration, I interpret limitation [3.1] to encompass a memory receiving the “first one or more instruction sets.”

125. As shown in Figure 8, Grotmol discloses a system in which “[p]rocessing 816” is connected to “[m]emory 814.”




FIG. 8

Ex. 1005, Fig. 8 (annotated)

126. As I discussed above with respect to limitation [1.3], Grotmol discloses “a memory that stores at least a first one or more digital pictures correlated with a first one or more instructions sets for operating a first physical device.” In my opinion, a POSITA would have understood from Figure 8 that the memory in Grotmol “obtains” the instruction sets for storage from the processor.

127. In my opinion, a POSITA would further have understood that the connection between Grotmol’s processor and memory is wired, and thus corresponds to the claimed “hardwired connection.” Grotmol discloses a “self-contained robotic apparatus” that comprises a “memory component” and “one or more physical processors.” *Id.*, 3:26-30. In my opinion, this suggests that a hardwired connection would be used, rather than a wireless one, because the two

components are not remote. Moreover, hardwired connections between a processor and its memory are typical in robotic devices.

128. Even assuming that Figure 8 of Grotmol does not expressly disclose a hardwired connection, A POSITA would find the use of a “hardwired connection” to obtain the first one or more instruction sets from the processor to the memory obvious because “there are a finite number of identified, predictable solutions, [and] a person of ordinary skill in the art has good reason to pursue the known options within his or her technical grasp.” *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. at 402. In particular, there are two ways to transmit an electronic signal: with wires (hardwired) and without wires (wireless)—there are no other options. A POSITA would have been motivated to try hardwired connections for the reasons discussed above. Furthermore, a POSITA would have had a reasonable expectation of success in implementing a hardwired connection in Grotmol because it is a well-known, routine way of connecting memory and a processor.

- b. **[3.2] wherein the one or more processor circuits or the another one or more processor circuits are used at least for operating the first physical device.**

Limitation	Claimed Features
[3.2]	Processor circuit(s) used to operate 1st device

129. As I discussed above for limitation [1.2], Grotmol discloses one or more processor circuits for operating the first physical device (robotic device). *See supra* limitation [1.2].

2. **Claim 5: The system of claim 1, wherein the causes the one or more processor circuits to execute the first one or more instruction sets for operating the first physical device includes inserting, via a hardwired connection, the first one or more instruction sets for operating the first physical device into at least one portion of the one or more processor circuits.**

Claim 5	Additional Features to Claim 1
	+ 1st instructions inserted into processor circuit(s) via a hardwired connection

130. Claim 5 is similar to claim 3, but recites that the instruction sets are “inserted . . . into” the processor rather than being “obtained . . . from” the processor. Claim 5 also refers back to the “causes the one or more processor circuits” step, from limitation [1.7.3], which the ’449 patent requires the artificial intelligence unit to perform. As I discussed above in limitation [1.7], Grotmol’s artificial intelligence unit is part of or operating on its processor circuit, for example as software.

131. As shown in Figure 8, Grotmol discloses a system in which “[p]rocessing 816” is connected to “[m]emory 814.”




FIG. 8

Ex. 1005, Fig. 8 (annotated)

132. As I discussed above with respect to limitation [1.3] Grotmol discloses “a memory that stores at least a first one or more digital pictures correlated with a first one or more instructions sets for operating a first physical device.” Grotmol also discloses that the “one or more processor circuits” ultimately “execute the first one or more instruction sets.” *See supra* limitation [1.7.3]. In my opinion, a POSITA would have understood from Figure 8 that the instruction sets to be executed are “inserted” into the processor via the same memory connection. Further, it also my opinion that it would have been obvious to a POSITA that this connection is hardwired for the same reasons I discussed above for claim 3. *See supra* claim 3.

3. **Claim 11: The system of claim 1, wherein the artificial intelligence unit includes at least one of: a hardware element that is included in the one or more processor circuits, a hardware element that is included in a second one or more processor circuits, an application**

operating on the one or more processor circuits, an application operating on a third one or more processor circuits, or an element coupled to the one or more processor circuits, and wherein the one or more processor circuits include one or more microcontrollers.

Claim 11	Additional Features to Claim 1
	+ AI unit includes a hardware element included in processor circuit(s) + Processor circuit(s) include microcontroller(s)

133. Claim 11 requires that the artificial intelligence unit include at least one of a hardware element that is included in the one or more processor circuits, and wherein the one or more processor circuits include one or more microcontroller.

134. As I discussed above for claim 1, Grotmol discloses an artificial intelligence unit operating on one or more processor circuits. *See supra* limitation [1.2], [1.7]; Ex. 1005, 2:14-20, 33:15-18, 38:36-39, 39:23-28. In my opinion, a POSITA would understand that the processor(s) described in Grotmol would include a hardware element. *Id.*, 8:5-15 (defining “processor”).

135. Grotmol discloses the use of a “microcontroller” as a type of “processor” or “computing device.” Grotmol states that its use of the term “processor” is meant generally to include all types of digital processing devices, explicitly including “microprocessors”:

As used herein, the terms “processor”, “microprocessor” and “digital processor” are meant generally to include **all types of**

digital processing devices including, without limitation, digital signal processors (DSPs), reduced instruction set computers (RISC), general-purpose (CISC) processors, **microprocessors**, gate arrays (e.g., field programmable gate arrays (FPGAs)), PLDs, reconfigurable computer fabrics (RCFs), array processors, secure microprocessors, and application-specific integrated circuits (ASICs). Such digital processors may be contained on a single unitary IC die, or distributed across multiple components.

Ex. 1005, 8:5-9. Grotmol further states that its use of “computer” encompasses literally any device capable of executing a set of instructions and processing an incoming data signal:

As used herein, the terms “computer”, “computing device”, and “computerized device”, include, but are not limited to, personal computers (PCs) and minicomputers, whether desktop, laptop, or otherwise, mainframe computers, workstations, servers, personal digital assistants (PDAs), handheld computers, embedded computers, programmable logic device, personal communicators, tablet or “phablet” computers, portable navigation aids, J2ME equipped devices, smart TVs, cellular telephones, smart phones, personal integrated communication or entertainment devices, or **literally any other device capable of executing a set of instructions and processing an incoming data signal.**

Ex. 1005, 7:30-42. In my opinion, a microcontroller is both a “digital processing device” and a “device capable of executing a set of instructions and processing an incoming data signal.” Thus, a POSITA would understand that Grotmol discloses a system in which the “one or more processor circuits” includes “one or more microcontrollers.”

136. In addition, Grotmol discloses several processor circuits that a POSITA would have understood include one or more microcontrollers. For example, Grotmol discloses adaptive controllers that are designed to process and execute specific tasks. Grotmol discloses an “Adaptive controller” of type “1” that may learn to autonomously approach a target by generating correction torques on the robot’s wheels, but that may not control the robot’s gripper. *Id.*, 25:38-42 (“After a while ‘Adaptive controller’ of type ‘1’ (e.g., kNN classifier) may learn to approach the target autonomously generating correction torques on robot's wheels (‘output1’ and ‘output2’) and may not control the gripper (‘output N’). ‘Sensory input’ may include input from the cameras and encoders.”). Grotmol also discloses an “Adaptive controller of type 2” that is specifically designed to learn to turn the robot with only certain wheels *and* grasp targets with the gripper. *Id.*, 26:14-19 (“‘Adaptive controller’ of type ‘2’ (e.g., perceptron classifier) may learn to turn the robot appropriately and grasp the target by autonomously effectuating correct torques on

the robot's wheel ('output 2') and gripper ('output N'), while not controlling another wheel ('output 1').”).

137. Grotmol further discloses a “learning controller” that is used to process images that are outputted from the camera:

In some implementations of robotic navigation in an arbitrary environment, the sensor component 166 may comprise a camera configured to provide an output comprising a plurality of digital image frames refreshed at, e.g., 25 Hz frame rate. The sensor output may be processed by a **learning controller**, e.g., as illustrated and described with respect to FIG. 4A.

Ex. 1005, 10:16-22. Grotmol further discloses a “predictor” that is specifically programmed and used to assign priorities to available tasks.

In some implementations, a user (e.g., human operator) may train the system to switch between tasks based on the sensory context. FIG. 2 shows a diagram of the system, according to some implementations. **A predictor may assign priorities to the available tasks based on the sensory context**; the priorities may be trained by the user via a corrector plus combiner scheme. The predicted priorities may be filtered by a “persistent winner-take-all” module that only switches to a different task if the priority breaches a certain threshold.

Ex. 1005, 10:45-53. Thus, in my opinion, a POSITA would have understood that Grotmol's "adaptive controllers," "learning controller," and "predictor" include one or more microcontrollers.

138. Even if Grotmol did not expressly disclose the use of a microcontroller, it would have been obvious in view of the knowledge of a POSITA. As I discussed above, Grotmol discloses the use of FPGAs or ASICs in place of a microprocessor. A POSITA would consider FPGAs and ASICs to be programmable logic devices similar to a microcontroller, which the '449 patent confirms. The term "microcontroller" is primarily used in the '449 patent when describing "Logic Circuit 250." The '449 patent describes that its "Logic Circuit 250" may include "a microcontroller, field-programmable gate array (*FPGA*), application-specific integrated circuit (*ASIC*), and/or other computing circuit or device." '449 patent, 81:56-82:5. The use of a microcontroller in place of an FPGA or ASIC would have been "the simple substitution of one known element for another." *KSR*, 550 U.S. at 401. In my opinion, a POSITA would have been motivated to use a microcontroller instead of an FPGA, for example, because microcontrollers can be less complex and expensive. It is also my opinion that a POSITA would have had a reasonable expectation of success in implementing a microcontroller in Grotmol in place of an FPGA because the two devices are similar in operation. Implementing software on a microcontroller would have been well within the skill of a POSITA in 2017.

D. Ground 3: Claims 2 and 4 Are Obvious Based on Grotmol in View of Cosic 192

139. As set forth below, Grotmol in view of Cosic 192 renders obvious claims 2 and 4.

1. Motivation to Combine Grotmol and Cosic 192

140. A POSITA would have been motivated to combine the autonomous robotic system of Grotmol with Cosic 192's teachings regarding tracing and instrumentation in a similar system.

141. Cosic 192 describes a substantially similar system to Grotmol in which an "autonomous robot" records instructions and/or contextual information, uses this information to implement the robot's operations in various circumstances, and can subsequently recall and execute instructions based on similar operations performed in previous circumstances:

In further exemplary embodiments, Robot 70 may be an **autonomous robot**. As such, Robot 70 can be controlled by Control Application 120 that provides autonomous functionalities to Robot 70 through various non-UAIE systems or processes. In one example, Robot 70 may detect an object through its physical contact sensor and Control Application 120 may execute one or more instructions (i.e. Instruction Sets 600 or Operations 610, etc.) to move Robot 70 around the object. In another example, Robot 70 may detect an object through its

vision sensor (i.e. video or other camera, etc.) and Control Application 120 may execute one or more instructions (i.e. Instruction Sets 600 or Operations 610, etc.) to move Robot 70 closer to the object... Any other sensory inputs can be received and corresponding operations can be performed in other examples. In some embodiments involving autonomous robots, **UAIE can serve as a learning system to learn or record Control Application's 120 instructions (i.e. Instruction Sets 600 or Operations 610, etc.) and/or contextual information (i.e. Extra Info 630, etc.) used to implement autonomous Robot's 70 operations in various circumstances.** This knowledge (i.e. Knowledgebase 530, Knowledge Cells 800, etc.) can be loaded into another robot, device, machine, and/or application to provide pre-learned knowledge as previously described. The robot, device, machine, and/or application into which the knowledge is loaded may then implement anticipatory instructions (i.e. anticipatory Instruction Sets 600 or Operations 610, etc.) included in the loaded Knowledgebase 530 or Knowledge Cells 800... **UAIE can then implement the subsequent operations autonomously based on similar operations performed in prior similar circumstances,** thereby implementing anticipation of Control Application's 120 operations as previously described.

Ex. 1006, 201:45-202:51. Cosic 192 teaches that the system may obtain instructions via “tracing or profiling,” which it describes as a software design technique for

“outputting an application’s instructions, data and/or other information during the application’s execution (runtime).”

Instruction Acquisition Unit 111 can employ various techniques for attaching to and/or obtaining Software Application's 120 instructions, data, and/or other information. In one example, Instruction Acquisition Unit 111 can attach to and/or obtain Software Application's 120 instructions, data, and/or other information through **tracing or profiling**, or other techniques. **Tracing or profiling is used in software design as a technique of outputting an application's instructions, data, and/or other information during the application's execution (runtime).** Tracing or profiling may include adding trace code (i.e. instrumentation, etc.) to an application and/or outputting trace information to a specific target

Ex. 1006, 72:63-73:8. Cosic 192 further teaches that the autonomous system can execute “anticipatory instructions” via “instrumentation” from “Artificial Intelligence Unit 130.”

For example, when an event occurs and an appropriate event handler is called, **instrumentation can be implemented** in the application's source code to insert (including modifying, rewriting, overwriting, etc. of the source code) alternate instructions such as anticipatory instructions (i.e. anticipatory Operations 610 or Instruction Sets 600) from Artificial Intelligence Unit 130 at which point the modified source code can be recompiled and/or executed.

Ex. 1006, 114:45-53.

142. In my opinion, a POSITA would have been familiar with the teachings of both Grotmol and Cosic 192 and would have readily appreciated the advantages of their combination. Both references are directed to the same field of the training robots for autonomous operation, and both references discuss substantially similar approaches. Grotmol describes an autonomous robotic system in which, as part of a learning process, “[m]otor commands associated with a given image may be stored” (Ex. 1005, Abstract), but Grotmol does not expressly discuss how these commands are recorded in the robot’s programming. A POSITA looking for further disclosures regarding recording commands would appreciate Cosic 192’s teaching that *tracing* may be used to record motor commands during operation. Grotmol further describes that during autonomous operation, the stored motor commands may be executed when a new, similar image is detected (*Id.*, Abstract, 43:4-11), but Grotmol does not expressly discuss how the stored motor commands are executed in the robot’s programming. A POSITA looking for further disclosures regarding execution of stored motor commands would appreciate Cosic 192’s teach that *instrumentation* may be used to insert and execute the anticipatory instructions before returning the application to its prior state.

143. In my opinion, the resulting combination of Grotmol’s autonomous robotic system and Cosic 192’s tracing and instrumentation teachings would have

been expected and predictable, amounting to no more than the “[u]se of known technique to improve similar devices (methods, or products) in the same way.” M.P.E.P. § 2143. In my opinion, tracing and profiling are both common, well-understood programming techniques. Cosic 192 simply teaches how these techniques can be advantageously applied to a system similar to Grotmol’s.

144. For at least these reasons, it would have been obvious to a POSITA to combine Grotmol and Cosic 192.

2. Claim 2

Claim 2	Additional Features to Claim 1
	+ 1st instructions obtained from application for operating 1st device using tracing of application + application for operating 1st device runs on processor circuit(s)

- a. [2.1] The system of claim 1, wherein the first one or more instruction sets for operating the first physical device are obtained from an application for operating the first physical device using at least a tracing of the application for operating the first physical device,

Limitation [2.1]	Claimed Features
	1st instructions obtained from application for operating 1st device using tracing of application

145. In the field of computer science, “tracing” generally refers to logging information about a program’s execution. The ’449 patent describes that “tracing” may be used “to obtain instruction sets, data, and/or other information and provide

them as input into Artificial Intelligence Unit 110.” ’449 patent, 73:30-32. The ’449 patent describes several embodiments for obtaining instruction sets using existing tracing software.

Also, independent vendors may provide portable tools with similar functionalities that can be utilized across different computing systems and/or platforms. These native and portable tools may provide a wide range of **functionalities to obtain runtime and other information such as instrumentation, tracing or profiling...**

In one example, **obtaining instruction sets, data, and/or other information can be implemented through the .NET platform's native tools for application tracing or profiling** such as System.Diagnostics.Trace, System.Diagnostics.Debug, and System.Diagnostics.TraceSource classes for tracing execution flow, and System.Diagnostics.Process, System.Diagnostics.EventLog, and System

In a further example, **obtaining instruction sets, data, and/or other information can be implemented through Java platform's APIs for application tracing or profiling** such as Java Virtual Machine Profiling Interface (JVMPI), Java Virtual Machine Tool Interface (JVMTI), and/or other APIs or tools. These APIs can be used for instrumentation of an application, for

notification of Java Virtual Machine (VM) events, and/or other functionalities. One of the tracing or profiling techniques that can be utilized includes bytecode instrumentation. The profiler can insert bytecodes into all or some of the classes. In application execution profiling, for example, these bytecodes may include `methodEntry` and `methodExit` calls.

In a further example, **obtaining instruction sets, data, and/or other information can be implemented through `java.lang.Runtime` class that provides an interface for application tracing or profiling.** Examples of methods provided in `java.lang.Runtime` that can be used to obtain an application's instruction sets, data, and/or other information include `tracemethodcalls`, `traceinstructions`, and/or other methods. These methods prompt the Java Virtual Machine to output trace information for a method or instruction in the virtual machine as it is executed.

In a further example, obtaining instruction sets, data, and/or other information can be implemented through logging tools of the platform and/or operating system on which an application runs. **Some logging tools may include nearly full feature sets of the tracing or profiling tools previously described.** In one example, Visual Basic enables logging of runtime messages through its `Microsoft.VisualBasic.Logging` namespace that

provides a log listener where the log listener may direct logging output to a file and/or other target. In another example, Java enables logging through its `java.util.logging` class. In some aspects, obtaining an application's instruction sets, data, and/or other information can be implemented through logging capabilities of the operating system on which an application runs. For example, Windows NT features centralized log service that applications and operating-system components can utilize to report their events including any messages. Windows NT provides functionalities for system, application, security, and/or other logging. An application log may include events logged by applications.

'449 patent, 74:46-52, 74:64-75:3, 75:64-76:9, 76:53-62, 77:8-27. The '449 patent describes that "tracing" may include simply "adding trace code . . . to an application and/or outputting trace information to a specific target.'"449 patent, 73:22-24.

146. Grotmol discloses that "instruction sets for operating the first physical device are obtained from [the] application." For example, Grotmol discloses that "[m]otor commands associated with a given image may be stored." Ex. 1005, Abstract; *see also id.*, 33:21-24 ("Training of the robot's controller may be based on the user observing the robot's actions and sending one or more target control commands to the robot via the training handset."), *id.*, 39:17-19 ("Motor commands issued by a trainer corresponding to time instances when the images are acquired may be stored.").

147. Grotmol does not expressly disclose that that tracing is used to store the motor commands (“using at least a tracing of the application for operating the first physical device”), but this would have been obvious in view of Cosic 192.

148. Cosic 192 discloses that, in some embodiments of its autonomous system, “the receiving the first, the second, and the new instruction sets includes at least one of: **tracing**, profiling, or instrumentation of an [sic] user input.” Ex. 1006, 3:19-22 (“In further embodiments, the receiving the first, the second, and the new instruction sets includes at least one of: tracing, profiling, or instrumentation of an user input”). Cosic 192 describes the use of tracing as follows:

Tracing or profiling is used in software design as a technique of outputting an application's instructions, data, and/or other information during the application's execution (runtime). Tracing or profiling may include adding trace code (i.e. instrumentation, etc.) to an application and/or outputting trace information to a specific target. The outputted trace information (i.e. Software Application's 120 instructions, data, and/or other information, etc.) can then be provided to or recorded into a file, a table, a database, a DBMS, a data structure, a repository, an application, and/or other system or target that may receive such trace information. As such, Instruction Acquisition Unit 111 can utilize tracing or profiling to obtain Software Application's 120 instructions, data, and/or other information and provide them as input into Artificial Intelligence Unit 130 that may structure them

into knowledge for future autonomous operation of Software Application 120.

Ex. 1006, 73:2-73:19. Cosic 192 also identifies various tracing tools existing in the prior art. *E.g.*, 75:26-77:10.

149. In my opinion, it would have been obvious to combine Cosic 192's use of tracing with Grotmol's autonomous robot system in order to record motor commands input by a user. A POSITA would recognize such combination as the simple use of a known technique to improve a similar device in the same way. In addition, it would have been obvious to try tracing because "there are a finite number of identified, predictable solutions, [and] a person of ordinary skill in the art has good reason to pursue the known options within his or her technical grasp." *KSR* at 402. In my opinion, A POSITA would have had a reasonable expectation of success in combining the teachings of Grotmol and Cosic 192, which describe clearly complementary systems. This is especially true here, because the tracing taught by Cosic 192 is a well-known, routine programming technique. Adding a snippet of code to an application to log received motor commands would have been well within the skill of a POSITA in 2017.

- b. [2.2] wherein the application for operating the first physical device runs on the one or more processor circuits or on another one or more processor circuits.

Limitation [2.2]	Claimed Features
	application for operating 1st device runs on processor circuit(s)

150. In my opinion, Grotmol discloses “an application for operating the first physical device . . . wherein the application for operating the first physical device runs on the one or more processor circuits or on another one or more processor circuits.” *See supra* Claim 2.

151. Grotmol discloses that the robotic device may “comprise a controller operating a software component (e.g., the BrainOS® software platform) configured to enable training.” Ex. 1005, 33:15-18; *see also id.*, 2:14-18 (“One aspect of the disclosure relates to a method of determining a control signal for a robot. The method may be performed by a special purpose computing platform having one or more processors executing instructions stored by a non-transitory computer-readable storage medium.”).

3. Claim 4

Claim 4	Additional Features to Claim 1
	+ Causing processor circuit(s) to execute 1st instructions includes instrumenting an application for 1st or 2nd device with instructions for 1st device + application for operating 1st or 2nd device runs on processor circuit(s)

- a. [4.1] The system of claim 1, wherein the causes the one or more processor circuits to execute the first one or more instruction sets for operating the first physical device includes:

instrumenting an application for operating the first physical device with the first one or more instruction sets for operating the first physical device,

or instrumenting an application for operating the second physical device with the first one or more instruction sets for operating the first physical device, and

Limitation	Claimed Features
[4.1]	Causing processor circuit(s) to execute 1st instructions includes instrumenting an application for 1st or 2nd device with instructions for 1st device

152. In the field of computer science, “instrumentation” generally refers to adding code to a program or modifying existing code. *See* ’449 patent, 73:22-24 (“For instance, tracing or profiling may include adding trace code (i.e., instrumentation, etc.) to an application and/or outputting trace information to a specific target.”). The ’449 patent describes that “instrumentation” may comprise “application modifying code,” for example “selective modifications such as insertion of alternate code (i.e. anticipatory Instruction Sets 526, etc.).” *Id.*, 138:40-139:10; *see also id.*, 131:55-58 (describing instrumentation to comprise “inserting

additional code, overwriting or rewriting existing code, and/or branching to a separate segment of code”).

153. As I discussed above, Grotmol discloses that the artificial intelligence unit “causes the one or more processor circuits to execute the first one or more instruction sets for operating the first physical device.” *See supra* limitation [1.7.3].

154. Grotmol does not expressly disclose that this step is performed by “instrumenting an application for operating the first physical device with the first one or more instruction sets,” but this would have been obvious in view of Cosic 192.

155. Cosic 192 discloses that an “instrumented call” may be made to an “application modifying function,” such as “anticipatory instructions,” that can be placed anywhere in the application's code and executed at any point:

In the above sample code, **instrumented call to Application Modification Unit's 112 function** (i.e. modifyApplication, etc.) can be placed after the web page loading function (i.e. loadPage, etc.) of an application such as web browser, for example. Similar call to application modifying function can be placed after or before some or all functions/routines/subroutines, some or all lines of code, some or all statements, some or all instructions or instruction sets, some or all basic blocks, and/or some or all other code segments of web browser or other application. One or more application modifying function calls **can be placed anywhere in the application's code and can be executed at any points in**

the application's execution. The application modifying function (i.e. modifyApplication, etc.) **may include UAIE determined anticipatory instructions** (i.e. anticipatory Instruction Sets 600 or Operations 610, etc.) that can modify the application (i.e. web browser in this example, etc.) to which UAIE is attached.

For example, when an event occurs and an appropriate event handler is called, **instrumentation can be implemented** in the application's source code to insert (including modifying, rewriting, overwriting, etc. of the source code) alternate instructions **such as anticipatory instructions** (i.e. anticipatory Operations 610 or Instruction Sets 600) from Artificial Intelligence Unit 130 at which point the modified source code can be recompiled and/or executed.

Ex. 1006, 109:26-43, 114:45-53. Cosic 192 also gives examples of existing software “utilized across different platforms” that can be used to instrument an application. *E.g., id.*, 112:41-113:39.

156. In my opinion, it would have been obvious to combine Cosic 192's use of instrumentation with Grotmol's autonomous robot system in order to execute anticipatory motor commands. A POSITA would recognize such combination as the simple use of a known technique to improve a similar device in the same way. It is also my opinion that a POSITA would have had a reasonable expectation of success in combining the teachings of Grotmol and Cosic 192, which describe clearly

complementary systems. This is especially true here, because the instrumentation taught by Cosic 192 is a well-known, routine programming technique. Adding a snippet of code to an application to instrument a call to an anticipatory instruction set would have been well within the skill of a POSITA in 2017.

- b. [4.2] wherein the application for operating the first physical device or the application for operating the second physical device runs on the one or more processor circuits.**

Limitation	Claimed Features
[4.2]	application for operating 1st or 2nd device runs on processor circuit(s)

157. As I discussed above with respect to claim 2, Grotmol discloses “an application for operating the first physical device . . . wherein the application for operating the first physical device runs on the one or more processor circuits or on another one or more processor circuits.” *See supra* Claim 2.

E. Ground 4: Claim 9 Is Obvious Based on Grotmol in View of Hickman

158. As set forth below, Grotmol in view of Hickman renders obvious claim 9.

1. Motivation to Combine Grotmol and Hickman

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 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 Hickman, which, as discussed in Section VI.B above, discloses a shared robot knowledge base that is updated based on robot experiences, allowing any particular robot to benefit from the prior experiences of any other connected robot. *See* Section VI.B.

160. Grotmol discloses that training sets and training configurations, which include pictures and associated instruction sets for operating the robot, can be stored or saved, and then loaded into the robot to provide the trained behavior. Ex. 1005, 33:27-32 (“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.”). Grotmol explains that the saved training configuration can also be loaded into robots other than the

one that did the learning. *Id.*, 33:30-32 (“the trained configuration may be loaded to ***one or more other robots*** in order to provide learned behaviors”).

161. Grotmol does not go into further detail regarding this sharing of learning behaviors. Thus, a POSITA would seek out other references, such as Hickman, for disclosures related to sharing knowledge between robots. A POSITA would have specifically looked to Hickman because, like Grotmol, it is directed to the training and autonomous operation of robots. Hickman is also cited on the face of Grotmol as a cited U.S. patent document. *Id.*, p.2. A POSITA would certainly consider the references listed on the face of Grotmol when looking to expand upon its teachings.

162. In addition, Grotmol and Hickman both address similar robotic training applications, such as learning how to “fetch” based on human input or commands:

Another example application may be called “Fetch.” The “Fetch” application may include robot instructions for playing fetch with a dog... When the robot executes the “Fetch” application, the person may adjust or further configure certain aspects of the individual tasks via a robot user interface on a smartphone, tablet computer, etc. to “fine tune” the application. For example, the person may increase or decrease the force with which the robot throws the ball or vary the praise given to the dog upon returning with the ball, e.g., “Good dog!” Creating the “Fetch” application in this manner is akin to “teaching” the robot a new skill.

After creating the “Fetch” application, the person may upload the application (or perhaps instruct the robot to upload the application) to the application knowledge base 311. Once uploaded to the application knowledge base 311, other robots can download and/or execute the “Fetch” application. Thus, the skill (i.e., playing fetch with a dog) learned by a single robot can be shared with all other robots via the application knowledge base 311 so that each robot does not have to learn the skill on its own.

Ex. 1007, 23:49-24:7.

For example, a task of playing fetch may be quite complex: the robot may be supposed to find an object, grasp it, and bring it back to base. In response to the robot not “seeing” the target, the robot may perform a random exploration behavior avoiding obstacles. In response to the robot being close to the target, the robot may perform a positioning maneuver and grasp the target. If the target is in the robot's gripper, the robot may find its way to the base avoiding obstacles and then release the object. The whole task may be highly nonlinear and/or noisy, so a very powerful classifier may infer the right sequence of actions from just the examples and be resistant to noise. The classifier may need to determine which particular input is important for a particular task. For example, a state of the gripper may not be relevant to target search behavior. Additional training may be needed to show the predictor that the search behavior should be performed independently of the state of the gripper.

According to the hierarchy depicted in FIG. 9, the “Fetch switcher” may be a high level behavior that activates one or more low level behavior. Such activation may depend on the context. The user may train the “Fetch switcher” using supervised learning. For example, the user may demonstrate to the switcher which behavior it should select in a particular situation. The user may use a user interface for such demonstrations, according to some implementations. In some implementations, if the target is far or not visible, the user may send a “select target approach” command. If the target is close, the user may send a “select grasping a target” command.

Ex. 1005, 24:31-48, 25:10-21; *see also id.*, 27:28-40, Fig. 9.

163. Grotmol states that the fetch task can be “quite complex.” *Id.*, 24:31-33. Thus, a POSITA would consider references, such as Hickman, that explicitly teach ways to “optimize” such complex training:

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.

Ex. 1007, 24:8-25.

164. More generally, a POSITA would have understood that pairing Hickman’s shared robot knowledge base with Grotmol’s autonomous robotic system would facilitate wider and faster learning among all robots connected to the knowledgebase and conserve local system resources. Further, A POSITA would expect these advantages to scale proportionally as the number of robots connected to the shared knowledgebase increases. Hickman explains as follows:

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.

[D]ividing application execution and storage between client devices and the cloud computing system allows more processing to be performed by the cloud computing system, thereby taking advantage of the cloud computing system's processing power and capability, for example.

In many situation, cloud computing may free users from certain hardware and software installation and maintenance tasks through the use of simplified hardware on the user's computer. Also, because the user's computer can access a vast network of computing resources (e.g., processors, disk drives, etc.), the user is not limited just to the computing and storage power of his or her local computer. Additionally, the sharing of computing resources across many users may reduce computing costs to individuals. For example, multiple computers connected to the cloud may be able to share the same pool of computing power, applications, and files.

Id., 11:31-41, 5:42-46, 1:38-48.

165. A POSITA familiar with Grotmol during or prior to 2017 would have readily appreciated the advantages of combination with Hickman's shared robot

knowledgebase. For example, in a 2014 Science Article, the CEO of Brain Corporation, the assignee of Grotmol, discussed the “profound” benefits of shared robot learning:

Efficient chips can run complex—and normally power-hungry—algorithms that enable robots to learn. The Internet, in turn, will allow far-flung robots to share those lessons and skills. A robot that learns how to pick a strawberry without crushing it, for instance, could uplink that skill for the benefit of its kind around the globe. That means neuromorphic computing could offer robots something far more profound than enhanced perceptual skills. When humans collectively pass along life lessons, we call that culture.

Ex. 1006, 183.

166. In my opinion, a POSITA would have found the resulting combination of Grotmol’s autonomous robotic system and Hickman’s shared robot knowledgebase to be expected and predictable. Further, it is my opinion that a POSITA would have had a reasonable expectation of success in combining the teachings of Grotmol and Hickman, which describe clearly complementary systems. Grotmol expressly contemplates the sharing of trained configurations among robots, and Hickman’s expansion upon this idea does not interfere with Grotmol’s autonomous robotic system.

2. **Claim 9: The system of claim 1, wherein the memory further stores at least a second one or more digital**

pictures correlated with a second one or more instruction sets for operating the second physical device, 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 second physical device are learned in a second learning process that includes operating the second physical device at least partially by a second user.

Claim 9	Additional Features to Claim 1
	+ 2nd DP correlated with 2nd instructions for 2nd device + Portion of 2nd DP or 2nd instructions learned with 2nd user operating 2nd device

167. In my opinion, Grotmol in view of Hickman renders obvious claim 9, which requires (i) storing a second digital picture correlated with a second instruction set for operating the second physical device; and (ii) the second digital picture or instruction set is learned during an additional learning process that involves operating the second device by a second user.

168. As I discussed above for claims 1, 6, and 8 in Ground 1, Grotmol discloses (i) training a robot wherein images (digital pictures) detected by a camera during operation of the robot are correlated with motor commands (instructions); (ii) multiple training (learning) processes can be done to train multiple behaviors; and (iii) a robot can be loaded with a first trained configuration learned during a training process involving operation of that robot by a user, and can also be loaded with

another trained configuration from a training process involving operation of a different robot by a user. Ex. 1005, 33:24-35; *see supra* limitations [1.8], claim 6, claim 8. Thus, a POSITA would understand that the memory of a first robot in Grotmol can include both a first training set (first pictures correlated with first instruction sets) from itself and a second training set (second image correlated with second instruction set) from a second device loaded onto its memory.

169. Hickman discloses “a shared robot knowledge bases for use with cloud computing systems” that “collects data from a robot about an object the robot has encountered in its environment, and stores the received data in the shared robot knowledge base ... sends instructions for interacting with an object to a robot ... sends instructions to a robot for executing an application based on information stored in the shared robot knowledge base.” Ex. 1007, Abstract. As noted in Hickman, the “information in the shared robot knowledge bases is updated based on robot experiences so that any particular robot may benefit from prior experiences of other robots.” *Id.*

170. Grotmol discloses that multiple learning processes can be done to train multiple behaviors for the robot, including “simple tasks” and “complex tasks.”

Apparatus and methods for hierarchical learning are disclosed herein, in accordance with one or more implementations. In supervised learning, a user may train **simple tasks** by demonstrating example tasks to the system.

For example, **a task of playing fetch may be quite complex:** the robot may be supposed to find an object, grasp it, and bring it back to base. In response to the robot not “seeing” the target, the robot may perform a random exploration behavior avoiding obstacles. In response to the robot being close to the target, the robot may perform a positioning maneuver and grasp the target. If the target is in the robot's gripper, the robot may find its way to the base avoiding obstacles and then release the object. **The whole task may be highly nonlinear and/or noisy, so a very powerful classifier may infer the right sequence of actions from just the examples and be resistant to noise. The classifier may need to determine which particular input is important for a particular task.** For example, a state of the gripper may not be relevant to target search behavior. Additional training may be needed to show the predictor that the search behavior should be performed independently of the state of the gripper.

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”.

Ex. 1005, 24:17-20, 24:31-48 26:5-8. Grotmol does not limit the number of users that can train the multiple behaviors for the robot. Thus, in my opinion, a POSITA would have looked for other references that explicitly disclose different users training different behaviors. Hickman is one such reference, explicitly stating that robot experiences used to train other robots connected to its shared robot knowledgebase may include coaching from “*humans*” in plural:

By analyzing past experiences from many different robots, the cloud computing system 304 may be configured to synthesize certain rules for robots to follow when executing the “Groceries” application. For example, some rules that could be synthesized or generalized from a large number of robot experiences might include (i) unidentifiable frozen items should be stored in a freezer, (ii) identifiable cold items should be stored in a refrigerator, (iii) certain types of vegetables should be stored in a certain refrigerator drawer, (iv) items should not be placed on top of bread in a pantry, etc. In this manner, **robot experiences (including “coaching” from humans)** can be used to improve the way that robots execute the “Groceries” application.

Ex. 1007, 23:36-48.

171. In combination with Hickman’s shared robot knowledgebase, Grotmol’s system could utilize different trainers to train behaviors involving different levels of complexity in order to take advantage of the varying skills and knowledge of trainers. Thus, in my opinion, it would have been obvious to combine

Hickman's shared robot knowledge base with Grotmol's autonomous robot system such that the shared knowledgebase includes the training of multiple robots by multiple users.

172. Accordingly, the combination of Grotmol and Hickman renders obvious storing in memory a second digital picture correlated with second instructions for operating a second device, where the second digital picture or second instruction are learned in a learning process involving operating the second device by a second user.

IX. The '583 Patent

A. The Claims

173. I understand that the petition challenges claims 1-34 of the '583 patent. Independent claim 1 of the '583 patent is directed to a system in which a device anticipates and autonomously executes instruction(s) corresponding to a newly detected circumstance based on comparison to a previously detected and stored circumstance. As I explain in this Declaration, a POSITA, in and prior to 2017, would have understood that there was nothing inventive about the steps of this method, which were all well known in the art at the time. Claim 1 is reproduced below:

1. A system comprising:
 - (a) one or more processors; and

- (b) one or more non-transitory machine readable storing machine readable code that, when executed by the one or more processors, causes the one or more processors to perform at least:
- (c) 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;
- (d) receiving or generating a first one or more instruction sets for operating the first device; and
- (e) learning the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device.

174. Independent claim 33 is nearly identical to claim 1, but recites a different preamble – “A method implemented using a computing system that includes one or more processors, the method comprising...”

175. Independent claim 34 is nearly identical to claim 1, but requires “means for generating or receiving,” and “means for learning,” as compared to claim 1 which specifies that the generating/receiving and learning steps are performed by “one or more processor circuits.”

176. Dependent claims 2-32, which all depend from claim 1, can be summarized as follows in the table below based on the additional steps or features that they require. As I further explain in this Declaration, the steps and features added by these dependent claims were also well-known at or prior to 2017, and a

POSITA would have understood that there was nothing inventive about the combination of steps and features recited in these claims.

Dependent Claim	Additional Features Claimed
Claim 2	In addition to the limitations of claim 1, claim 2 further requires receiving the first one or more digital pictures from a picture capturing apparatus, receiving the first one or more instruction sets from an application, and a learning process that includes storing a correlation between a first digital picture and a first instruction set. + picture capturing apparatus + application + storing correlation between 1st DP and 1st instructions
Claim 3	In addition to the limitations of claim 1, claim 3 further requires the learning process to include a determination that a first one or more instruction sets temporally corresponds to the first one or more digital pictures. + 1st DP and 1st instructions temporally correspond
Claim 4	In addition to the limitations of claim 1, claim 4 further requires receiving or generating a new digital picture,

Dependent Claim	Additional Features Claimed
	determining an operation defined within an instruction set based on a partial match between a new and old digital picture, and executing a first instruction set. + determining 1st instructions based on a partial match between 1st DP and 2nd DP + causing operation from 1st instructions to perform on 1st device
Claim 5	In addition to the limitations of claim 4, claim 5 further requires a first or second device execute the instruction set + execute 1st instructions on 1st device
Claim 6	In addition to the limitations of claim 4, claim 6 further requires determining the first instruction set from a partial match of portions of new and old digital images, where the portions are objects. + Partial match between 1st DP and 2nd DP based on objects
Claim 7	In addition to the limitations of claim 4, claim 7 further requires determining the 1st instruction set based on objects detected in a first and second digital picture. + Objects detected in 1st DP and 2nd DP + Partial match of objects in 1st DP and 2nd DP

Dependent Claim	Additional Features Claimed
Claim 8	In addition to the limitations of claim 4, claim 8 further requires the first instruction set to include information about the state of the first device. + State information of 1st device
Claim 9	In addition to the limitations of claim 4, claim 9 further requires a learning process that includes operating a first device by a user and a knowledgebase for storing the correlation of a first digital picture and first instruction set and the correlation of a second digital picture and a second instruction set. + Knowledgebase + 2nd DP correlated with 2nd instructions for 1st device + Learning process where user operates 1st device
Claim 10	In addition to the limitations of claim 4, claim 10 further requires a knowledgebase with a first correlation and a second correlation for operating a first device, where the learning process for the first correlation and second correlation are operated by two different users. + Knowledgebase

Dependent Claim	Additional Features Claimed
	+ 2nd DP correlated with 2nd instructions for 1st device + 2nd DP or 2nd instructions learned in process by 2nd user
Claim 11	In addition to the limitations of claim 4, claim 11 further requires a knowledgebase with a first correlation for operating a first device, a second correlation for operating a third device, and a learning process that includes operating the third device. + Knowledgebase + 2nd DP and 2nd instructions for 3rd device + 1st DP or 1st instructions learned on 1st device + 2nd DP or 2nd instructions learned on 3rd device
Claim 12	In addition to the limitations of claim 4, claim 12 further requires a second digital image depicting the first device's surroundings and causing the 1st device to execute an operation defined by the first instruction set. + 2nd DP depicts 1st device's surrounding + 1st device performs operations defined in 1st instructions
Claim 13	In addition to the limitations of claim 4, claim 13 further requires the second digital image to depict a second device's

Dependent Claim	Additional Features Claimed
	surroundings and causing the second device to execute an operation defined by the first instruction set. + 2nd DP depicts 2nd device's surroundings + 2nd device performs operations defined in 1st instruction set
Claim 14	In addition to the limitations of claim 4, claim 14 further requires the second digital picture to depict the second device's surroundings, modifying the first instruction set and applying it to the second device, and causing the second device to perform the modified instruction set. + 2nd DP depicts 2nd device's surroundings + Modifying 1st instructions and applying to 2nd device + Performing operation defined in modified 1st instructions on 2nd device
Claim 15	In addition to the limitations of claim 4, claim 15 further requires modifying the first instruction set and executing the modified first instruction set. + Modifying 1st instructions + Performing modified 1st instructions
Claim 16	In addition to the limitations of claim 4, claim 16 further requires modifying the first digital picture and determining an

Dependent Claim	Additional Features Claimed
	instruction set based a partial match between the modified first digital picture and a second digital picture. + Modifying 1st DP + Partial match between 2nd DP and modified 1st DP
Claim 17	In addition to the limitations of claim 4, claim 17 further requires modifying the first digital picture, learning the correlation of the modified first digital picture with the first instruction set, and determining the first instruction set based on partial match between modified first digital picture and second digital picture. + Modifying 1st DP + Learning correlation modified 1st DP and 1st instructions + Partial match between 2nd DP and modified 1st DP
Claim 18	In addition to the limitations of claim 4, claim 18 further requires modifying the first instruction set, learning the correlation of the first digital picture with the modified first instruction set, determining the modified first instruction set based on a partial match between a first and second digital

Dependent Claim	Additional Features Claimed
	picture, and performing the operation defined by the modified first instruction set. + Modifying 1st instructions + Learning correlated 1st DP and modified 1st instructions + Performing operation defined by modified 1st instructions
Claim 19	In addition to the limitations of claim 4, claim 19 further requires an artificial intelligence system that learns the correlation between a first digital picture and a first instruction set and determines the first instruction set based on a partial match between a first and second digital picture. + AI system
Claim 20	In addition to the limitations of claim 4, claim 20 further requires an artificial intelligence system with inputs for receiving digital pictures, applying the first digital picture to the inputs, and outputting the first instruction set. + AI system + Input to receive DPs + Output instructions

Dependent Claim	Additional Features Claimed
Claim 21	In addition to the limitations of claim 20, claim 21 further requires one input for one digital picture or multiple inputs for multiple digital pictures. + One input for one DP
Claim 22	In addition to the limitations of claim 19, claim 22 further requires a knowledgebase. + Knowledgebase
Claim 23	In addition to the limitations of claim 4, claim 23 further requires connecting the first digital picture with a first instruction set. + Connecting 1st DP and 1st instructions
Claim 24	In addition to the limitations of claim 4, claim 24 further requires the first digital picture is a whole digital image, collection of pixels, or a collection of values. + 1st DP = whole image or collection of pixels
Claim 25	In addition to the limitations of claim 4, claim 25 further requires a server receiving a first digital picture or first instruction set from a first device, a second device receiving

Dependent Claim	Additional Features Claimed
	the first digital picture or first instruction set from a server, and the second device performing the first instruction set. + Server receives 1st DP or 1st instructions from 1st device + 2nd device receives 1st DP or 1st instructions + 2nd device performing operation from 1st instructions
Claim 26	In addition to the limitations of claim 4, claim 26 further requires all of the structures listed in claim 1 to include one of various listed embodiments. + Various embodiments
Claim 27	In addition to the limitations of claim 4, claim 27 further requires a knowledgebase including a first correlation for a first device, a second correlation for operating a third device, and performing the second instruction set on a third device. + Knowledgebase + 2nd instructions for operating 3rd device + Performing 2nd instructions on 3rd device
Claim 28	In addition to the limitations of claim 4, claim 28 further requires a partial match between a first and second digital

Dependent Claim	Additional Features Claimed
	picture based on a product of values representing a portion of the respective images. + Partial match based on product of values
Claim 29	In addition to the limitations of claim 4, claim 29 further requires detecting object(s) in the second digital picture, generating an object representation, and generating a new image including said object representation. + Detecting object in 2nd DP + Generating object representation + Generating 2nd DP image including object representation
Claim 30	In addition to the limitations of claim 4, claim 30 further requires detecting the location of one or more objects and a new digital picture including a top-down view of said object. + Detecting location of object in 2nd DP
Claim 31	In addition to the limitations of claim 4, claim 31 further requires generating the first instruction set as a representation of the second instruction set. + 1st instructions as representation of 2nd instructions

Dependent Claim	Additional Features Claimed
Claim 32	In addition to the limitations of claim 4, claim 32 further requires the first digital picture to depict a representation of the first device's surroundings. + 1st DP includes 1st device's surroundings

B. Ground 1: Claims 1-9, 11-24, 26, 28-29, and 31-34 Are Anticipated by Grotmol

177. As set forth in detail below, claims 1-9, 11-24, 26, 28-29, and 31-34 of the '583 patent are anticipated by Grotmol. The discussion below specifies where each element of the claims is found in Grotmol, and includes a detailed explanation of the significance of the quotations and citations.

1. Claim 1

a. [1.1] A system comprising:

Limitation [1.1]	Claimed Features
	system

178. Grotmol discloses a system implemented as a “self-contained robot apparatus:”

Another aspect of the disclosure relates to **a self-contained robotic apparatus**. The apparatus may comprise a platform, a

first sensor, a second sensor, a memory component, and one or more physical processors.

An apparatus and methods for training and/or operating a **robotic device** to follow a trajectory.

FIG. 2 shows a diagram of the **system**, according to some implementations.

Ex. 1005, 3:27-30, Abstract, 10:47-48.

179. Accordingly, a POSITA would have understood Grotmol to disclose a system.

b. [1.2] one or more processors; and

Limitation [1.2]	Claimed Features
	Processor(s)

180. Grotmol discloses “a self-contained robotic apparatus, the apparatus comprising...one or more physical processors configured to operate [the system].”

Id., claim 12.

A method of determining a control signal for a robot, the method being performed by a special purpose computing platform

having **one or more processors** executing instructions stored by a non-transitory computer-readable storage medium,

A non-transitory computer-readable storage medium having instructions embodied thereon, the instructions being executable **by a processor** to selecting an outcome of a plurality of outcomes by at least:

The apparatus may comprise a platform, a first sensor, a second sensor, a memory component, **and one or more physical processors.**

Id., claim 1, claim 17, 3:28-30.

181. Grotmol discloses a system including a “processing” unit in Figure 7 connected to at least a sensory unit and a memory unit.




FIG. 7

Ex. 1005, Fig. 7

182. Further, Grotmol defines “processor” as follows:

As used herein, the terms “**processor**”, “microprocessor” and “digital processor” are meant generally to include **all types of digital processing devices** including, without limitation, digital signal processors (DSPs), reduced instruction set computers (RISC), general-purpose (CISC) processors, microprocessors, gate arrays (e.g., field programmable gate arrays (FPGAs)), PLDs, reconfigurable computer fabrics (RCFs), array processors, secure microprocessors, and application-specific integrated circuits (ASICs). Such digital processors may be contained on a single unitary IC die, or distributed across multiple components.

Id., 8:5-15.

183. A POSITA would have understood Grotmol's use of the term processor to be consistent with the claimed "one or more processors" of the '583 patent.

- c. **[1.3] 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:**

Limitation [1.3]	Claimed Features
	Non-transitory machine readable media storing machine readable code

184. Claims 1, 12, and 17 of Grotmol recite, in relevant part:

A method of determining a control signal for a robot, the method being performed by a special purpose computing platform having one or more processors executing instructions stored by a **non-transitory computer-readable storage medium...**

A **self-contained robotic apparatus**, the apparatus comprising...one or more physical processors configured to operate [the system].”

A non-transitory computer-readable storage medium...

Ex. 1005, claims 1, 12, and 17.

185. Further, Grotmol may be implemented using non-transitory machine readable media storing machine readable code:

Yet another aspect of the disclosure relates **to a non-transitory computer-readable storage medium having instructions embodied thereon, the instructions being executable by a processor to perform** a method of selecting an outcome of a plurality of outcomes.

In some implementations, **a robotic device may comprise a controller operating a software component** (e.g., the BrainOS® software platform) configured to enable training

Id., 4:13-17, 33:12-30.

186. Thus, a POSITA would have understood Grotmol's system to include non-transitory machine readable media storing machine readable code that causes its processors to perform at least limitations [1.4]-[1.6].

- d. **[1.4] 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:**

Limitation	Claimed Features
[1.4]	Receiving a new digital picture

187. Grotmol discloses processor(s) receiving “training sets” where each training set includes at least a digital picture³:

Another aspect of the disclosure relates to a self-contained robotic apparatus. The **apparatus** may comprise a platform, a **first sensor**, a second sensor, a **memory component**, and **one or more physical processors**.

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**.

Ex. 1005, 3:27-30, 3:35-39

188. Grotmol further describes that the received new images depict at least a portion of a first device’s surroundings:

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.

³ I understand that the ’583 patent describes that a “digital picture” may be a “frame” of “a stream of digital pictures (i.e. motion picture, video, etc.).” 583 patent, 87:45-51.

In some implementations of robotic navigation in an arbitrary environment, **the sensor component 166 may comprise a camera configured to provide an output comprising a plurality of digital image frames refreshed at, e.g., 25 Hz frame rate.** The sensor output may be processed by a learning controller, e.g., as illustrated and described with respect to FIG. 4A.

The sensor component 166 may be characterized by an aperture/field of view 168 (e.g., **an extent of the observable world that may be captured by the sensor at a given moment**). The sensor component 166 may **provide information associated with objects within the field-of-view 168.** In one or more implementations, such as object recognition, and/or obstacle avoidance, the **output provided by the sensor component 166 may comprise a stream of pixel values associated with one or more digital images.**

In some implementations, the camera may **be spanning the entire horizontal view field of 360 degrees.** A 360 degree horizontal view field may provide additional information about robot's environment and enable cyclical shift of images.

Ex. 1005, Abstract, 10:16-22, 8:61-9:3, 38:47-51; *see also* 39:13-22, Fig. 20A. A POSITA would have understood Grotmol's disclosures of following a visual across a trajectory and capturing images with a 360 degree view depicts a portion of a first device's surrounding.

189. Thus, a POSITA would have understood that Grotmol discloses processor(s) receiving a first one or more digital pictures, wherein the first one or more digital pictures depict at least a portion of a first physical device's surrounding.

e. [1.5] receiving or generating a first one or more instruction sets for operating the first device; and

Limitation [1.5]	Claimed Features
	Receiving instruction sets

190. As I described above in limitation [1.4], Grotmol discloses processor(s) receiving training sets where each training set comprises at least a digital image and an instruction set⁴ for operating the first device:

A given training set may comprise an instance of the video frame,
an instance of the binary signal, and an **instance of a motor**

⁴ I understand that the '583 patent describes that "instruction sets" may comprise, for example, instructions "causing [a] vehicle to stop," "causing [a] change of direction," and "causing [a] robotic arm to screw [a] screw," among other examples. '583 patent, 92:65-93:41; *see also id.*, 159:39-160:12 (describing a learning device to "turn or steer right").

control indication configured to cause the robot to execute an action.

Motor commands issued by a trainer corresponding to time instances when the images are acquired may be stored.

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.**

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.**

Ex. 1005, 3:36-39, 39:17-19, Abstract, 33:21-27. Because the training is performed on a first device, a POSITA would have understood that the training sets are for operating the first device.

191. A POSITA would have understood the “instruction sets” of the ’583 patent to correspond to the motor commands disclosed by Grotmol due to the way the term is broadly described in the specification. The ’583 patent discloses instruction sets of any programming language or construct that can be executed on a processor or logic circuit, used by an actuator to operate a device, and can be received by any element of the system:

In some aspects, **an instruction set** may include **any computer command, instruction, signal, or input** used in Application Program 18, Processor 11, Logic Circuit 250, and/or other processing element. Therefore, the terms instruction set, command, instruction, signal, input, or other such terms may be used interchangeably herein depending on context.

In further embodiments, **an instruction set** includes at least one of: a command, a keyword, a symbol, an instruction, an operator, a variable, a value, an object, a data structure, a function, a parameter, a state, a signal, an input, an output, a character, a digit, or a reference thereto. In further embodiments, the one or more instruction sets include a source code, a bytecode, an intermediate code, a compiled code, an interpreted code, a translated code, a runtime code, an assembly code, a structured query language (SQL) code, or a machine code.

’583 patent, 72:51-57, 4:34-43.

192. Thus, a POSITA would have understood that Grotmol discloses receiving or generating a first one or more instruction sets for operating the first physical device.

- f. **[1.6] learning the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device.**

Limitation	Claimed Features
[1.6]	Learning association for first picture and first instruction set

193. As I discussed above for limitations [1.4]-[1.5], Grotmol discloses processors generating and storing training sets that include at least a digital picture and an instruction set for operating a first device. Ex. 1005, 3:27-39, Abstract, 39:17-19, 33:12-30.

194. Grotmol discloses a memory component which may be configured to store training sets, where each training set consists of at least one instance of a video frame and one instance of a motor control indication:

The apparatus may comprise a platform, a first sensor, a second sensor, **a memory component**, and one or more physical processors. The platform may comprise a motor.... 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.

Ex. 1005, 3:28-39.

195. Grotmol discloses that, as a robot is trained, the video frames are “associated with” or “correspond to” (i.e., are the claimed “correlated with⁵”) motor control indications issued around the same time:

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**.

Motor commands **associated with** a given image may be stored.

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.

Motor commands issued by a trainer **corresponding to** time instances when the images are acquired may be stored.

⁵ I understand that the '583 patent describes that a device may “correlate Digital Pictures 525 from the stream of Digital Pictures 525 with temporally corresponding Instruction Sets”). '583 patent, 91:52-56.

Ex. 1005, 33:59-62, Abstract, 43:5-9, 39:17-19. In my opinion, because motor commands are stored corresponding to and/or associated with digital pictures, Grotmol discloses learning correlated motor commands and digital pictures.

196. Grotmol also discloses that a human (“user”)⁶ can, by manually operating the robot, train the robot to learn pictures that are correlated with instructions (i.e., a “learning process”):

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.

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

⁶ Certain challenged claims of the ’583 patent recite “user(s),” which in my analysis I have considered to include humans. However, the ’583 patent does not limit the term “user” to only humans. ’583 patent, 69:11-13 (“User 50 (also referred to simply as user or other similar name or reference) comprises a human user or non-human user.”).

operational parameters) corresponding to time instances when the images are acquired may be stored.

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.

Ex. 1005, 33:59-62, 39:13-22, 43:4-11; *see also id.*, Abstract.

197. For example, Grotmol discloses both a “listen mode,” and an “override mode,” whereby the robot will listen to commands sent by a teacher (user) and learn the association:

In listen mode, the robot may be “listening” to or otherwise observing what the teacher teaches, and may not do anything on its own. However, the robot may **learn an association**. But if the teacher stops sending a command, the robot may stay still and wait for the next command. In some implementations, teaching the robot may begin with the listen mode. Once enough examples have been provided and **the robot has learned something**, the override mode may be used.

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.

According to exemplary implementations, the **user may train a component** of the fetch behavior.

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.

In some implementations of supervised training of robots, control instructions (also referred to as corrections) produced by the **trainer** (e.g., human) may be combined with control

instructions produced by the robot controller instructions (predictions).

Ex. 1005, 23:3-22, 24:49-50, 33:21-27, 34:50-55.

198. Grotmol describes this training as a learning process:

Apparatus and methods **for hierarchical learning** are disclosed herein, in accordance with one or more implementations. **In supervised learning, a user may train** simple tasks by demonstrating example tasks to the system.

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

Id., 24:17-20, 33:59-62; *see also id.*, claim 12 (reciting “one or more physical processors configured to operate a random k-nearest neighbors *learning process*”).

Thus, the correlated digital pictures and motor controls are “learned” in Grotmol.

199. A POSITA would have understood the claimed “learning process” of the ’583 patent to correspond to the learning process described in Grotmol due to the way the term is broadly described in the specification. The ’583 patent does not limit “learning process” to a particular technique:

The disclosed systems, devices, and methods are independent of the artificial intelligence model and/or technique used and **any**

model and/or technique can be utilized to facilitate the functionalities described herein. Examples of these models and/or techniques include deep learning, supervised learning, unsupervised learning

'583 patent, 95:53-64.

200. Thus, a POSITA would have understood that Grotmol discloses learning a first one or more digital pictures correlated with the first one or more instruction sets for operating the first physical device.

2. Claim 2

Claim 2	Additional Features to Claim 1
	+ picture capturing apparatus + application + storing correlation between 1st DP and 1st instructions

- a. **[2.1] The system of claim 1, wherein the receiving the first one or more digital pictures includes receiving the first one or more digital pictures from a picture capturing apparatus,**

201. As I described above in limitation [1.4], Grotmol discloses a system with processors for receiving one or more digital pictures. Ex. 1005, Abstract, 3:27-39, 10:16-22, 8:57-9:3.

202. Grotmol further discloses a robot receiving images from a picture capturing apparatus (camera):

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.

In some implementations of robotic navigation in an arbitrary environment, the sensor component 166 may comprise a **camera configured to provide an output comprising a plurality of digital image frames** refreshed at, e.g., 25 Hz frame rate. The sensor output may be processed by a learning controller, e.g., as illustrated and described with respect to FIG. 4A.

In some implementations, the robot may comprise a robotic vehicle (e.g., 160 in FIG. 1) **comprising a camera** (e.g., 166). The robot may comprise a controller configured to store a sequence of images obtained during training phase.

Images 2000, 2010 may be obtained with a variety of sensors, e.g., a **video camera, IR camera, RADAR, SONAR, LiDAR, ultrasonic imager, and/or other device**. In some implementations, the camera may be spanning the entire horizontal view field of 360 degrees.

Ex. 1005, Abstract, 10:16-22, 38:36-39, 38:44-48. In my opinion, a camera such as that described in Grotmol constitutes a picture capturing apparatus. Grotmol teaches

that the sequence of images obtained from a camera during a training phase are stored, therefore Grotmol's processors must receive digital pictures from the camera.

203. Thus, A POSITA would have understood that Grotmol discloses a system wherein processors receive digital pictures from a picture capturing apparatus.

- b. **[2.2] and wherein the receiving the first one or more instruction sets for operating the first device includes receiving the first one or more instruction sets for operating the first device from: an application for operating the first device, a system for operating the first device, one or more microcontrollers, another one or more processors, or one or more actuators,**

204. As I described above in limitations [1.2] and [1.5], Grotmol discloses a system implemented as a "self-contained apparatus," that includes using one or more processors and a memory component for storing training sets that include digital pictures correlated with instruction motor control commands. Ex. 1005, Abstract, 3:27-39, Fig. 7. Because the robotic vehicle receives instruction sets from the system, and the processor in turn receives the instruction sets from the memory, a POSITA would have understood that Grotmol discloses a system receiving instruction sets for operating the first device.

205. Grotmol further discloses that the robotic device includes an application for operating the first device configured to enable training (e.g., the BrainOS® software platform):

In some implementations, a robotic device may comprise a controller **operating a software component** (e.g., the BrainOS® software platform) **configured to enable training**[.] A user may control/training the robot with a remote device (e.g., comprising a Gamepad® controller and an **iOS® application**, and/or a handset device (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.

One aspect of the disclosure relates to a method of determining a control signal for a robot. The method may be performed by **a special purpose computing platform having one or more processors executing instructions stored by a non-transitory computer-readable storage medium.**

Ex. 1005, 33:12-24, 2:15-19. In my opinion, the BrainOS® described in Grotmol is an application for operating a robotic vehicle. Based on Grotmol's disclosures above, a POSITA would have understood that the robotic vehicle receives instruction sets from said application.

206. Therefore, a POSITA would have understood that Grotmol discloses that the instruction sets are received from an application for operating the first device.

- c. **[2.3] 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 storing the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device into or onto at least one of: at least one non-transitory machine readable medium of the one or more non-transitory machine readable media, another or more non-transitory machine readable media, one or more volatile memories, one or more non-volatile memories, one or more storage devices, or one or more storage systems.**

207. As I described above in limitation [1.6], Grotmol discloses a memory component configured to store training sets comprising at least one digital picture (video frame instance) correlated with at least one instruction set (motor control command). Ex. 1005, 3:27-39 (“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.”); *see also id.*, 33:59-62, Abstract.

208. Grotmol further discloses that its memory component configured to store training sets may comprise a non-volatile memory:

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

Id., 3:27-39, 38:38-42. In my opinion, based on the above, Grotmol discloses storing training sets (comprised of both a digital picture and an instruction set) as a unit. Further, because training sets are learned during training and because the images are stored in a non-volatile memory, in my opinion Grotmol discloses that learned correlations are stored in a non-volatile memory (or it would have been obvious to do so).

209. As I described above in limitation [1.6], the '583 patent does not limit the learning process. '583 patent, 95:53-64. Further, because training sets are learned during training and because the images are stored in a non-volatile memory, in my opinion Grotmol discloses that learned correlations are stored in a non-volatile memory (or it would have been obvious to do so).

210. Thus, a POSITA would have understood that Grotmol discloses storing (e.g., "learning") training sets, comprised of digital images correlated with instruction sets, in a non-volatile memory.

3. **Claim 3: The system of claim 1, 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 determining that the first one or more instruction**

sets for operating the first device temporally correspond to the first one or more digital pictures.

Claim 3	Additional Features to Claim 1
	+ 1st DP and 1st instructions temporally correspond

211. As I described above in limitation [1.6], Grotmol discloses learning associations comprised of a digital picture and motor control indications around the same time. Ex. 1005, 33:59-62 (“Trained behaviors of the robotic device may be configured based on learning of associations between sensory context ... and a respective action ... during training.”), Abstract (“Motor commands associated with a given image may be stored.”).

212. Grotmol discloses learning associations between sensory context and a respective action during training. *Id.*, 33:59-62. Grotmol further discloses that the learning of associations may be based on the timing of the images and the motor commands:

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

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.

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.

Ex. 1005, 39:13-22, 43:5-9. I understand that an “ordered buffer” is a region of memory that is allocated for storing data in a specific order or sequence. Grotmol describes that the training buffer may be ordered according to when the images were acquired. Because Grotmol teaches that the training data stored in memory in a sequence, in my opinion the training data is thus also stored where a digital picture temporally corresponds to an instruction set.

213. Thus, a POSITA would have understood that Grotmol discloses determining that the first one or more instruction sets for operating the device temporally correspond to the first one or more digital pictures.

4. Claim 4

Claim 4	Additional Features to Claim 1
	+ determining 1st instructions based on a partial match between 1st DP and 2nd DP + causing operation from 1st instructions to perform on 1st device

- a. **[4.1] 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:**

214. Grotmol discloses limitation [4.1] for the same reasons I discussed above in Section IX.B.1.c, *supra*, for limitation [1.3].

- b. **[4.2] receiving or generating a new one or more digital pictures;**

215. As I described above in limitation [1.4], Grotmol discloses receiving digital pictures. Ex. 1005, 39:23-28 (“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.”).

216. Grotmol discloses that during operation, the system acquires a new image. *Id.*, Abstract (“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.”). A POSITA would understand that when the robot “acquires” a new image, it is receiving or generating a new digital picture and that such acquisition would go through the processor (or it would have been obvious to do so).

217. Additionally, Grotmol discloses comparing a recent image to the stored images in the training buffer (the first one or more digital pictures):

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 of robotic navigation in an arbitrary environment, the sensor component 166 may comprise a camera configured to provide an **output comprising a plurality of digital image frames** refreshed at, e.g., 25 Hz frame rate. The sensor output may be processed by a learning controller, e.g., as illustrated and described with respect to FIG. 4A.

In some implementations, the robot may comprise a robotic vehicle (e.g., 160 in FIG. 1) comprising a camera (e.g., 166). The robot may comprise a controller configured **to store a sequence**

of images obtained during training phase. In some implementations, the images **may be stored in an ordered buffer in a non-volatile memory** (e.g., memory 814 in FIG. 8).

Ex. 1005, 39:23-26, 10:15-21, 38:36-42. A POSITA would have understood that “the most recent, current” image corresponds to the claimed “new one or more digital pictures,” at least because Grotmol discloses that images are captured and stored sequentially.

218. Thus, a POSITA would have understood that Grotmol discloses that its processor receives new digital pictures.

- c. **[4.3] 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**

219. Grotmol discloses that selecting a motor command based on a best match between a new image with stored images:

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**.

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.**

Ex. 1005, 43:4-9, Abstract.

220. Grotmol discloses that the best matching step comprises calculating a similarity metric between a new image and an old image. A selection is determined when a threshold is breach:

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

Ex. 1005, 2:30-38. Because a “best match” is selected after a comparison between a stored digital picture (first digital picture) and a new digital picture, or portions thereof, Grotmol discloses using a partial match. Further, after the threshold is breached, Grotmol discloses selecting a corresponding training set. In my opinion, the selecting of a motor control based on a best match comparison corresponds to the claimed “determining” step of the ’583 patent.

221. As used in Grotmol, the term feature may describe, among other things, values of pixels in an image:

[R]esponsive to the **similarity measure breaching a threshold, selecting one or more training sets** from the plurality of **training feature sets**

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**

Ex. 1005, claim 1, 28:19-40.

222. A POSITA would understand Grotmol's matching process as consistent with the claimed partial match step of the '583 patent at least because both may comprise pixel-wise difference analysis between two pictures:

In some implementations, the control signal may be configured to cause the robot to execute the action. **The first input type may comprise a digital image comprising a plurality of pixel values.**

A similarity metric may be determined between the shifted/trimmed frames (e.g., 2020, 2040 in FIG. 20B). In some

implementations, **the similarity may be determined by performing pixel-wise difference between the overlapping regions** (e.g., the frames 2020, 2040 in FIG. 20B) of the given image and the matching image. A norm of the image difference in the overlapping region may be determined for individual images from the matching search set. An image corresponding to the lowest norm may be referred to as the best match.

Compare Ex. 1005, 3:12-15, 41:8-17, with

The comparing the at least one portion of the new digital picture with the at least one portion of the first digital picture may include **comparing at least one feature of the new digital picture with at least one feature of the first digital picture**. The comparing at least one portion of the new digital picture with the at least one portion of the first digital picture may include **comparing at least one pixel of the new digital picture with at least one pixel of the first digital picture**.

In any of the comparisons involving digital pictures or pixels, **Similarity Comparison 125 can utilize a threshold for acceptable number or percentage difference in determining a match for each compared pixel**.

Therefore, a pixel encoded as R130, G240, B50 matches or is sufficiently similar to a compared pixel encoded as R135,

G231, B57 because the differences in all three colors fall within the acceptable difference threshold (i.e. 10 in this example, etc.).

'583 patent, 21:39-46, 112:38-41, 112:49-53.

223. A POSITA would have understood that Grotmol discloses determining the first one or more instruction sets for operating the first physical device based on at least partial match between the new one or more digital pictures and the first one or more digital pictures.

d. [4.4] 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.

224. As I described above in limitations [1.6] and [4.3], Grotmol discloses a system that determines the learned instruction sets for operating the first device by comparing a new picture with a previously learned picture. Ex. 1005, 2:24-30, 43:4-9, Abstract. This determination is made at least by matching portions of the new picture and the old picture and finding a best match once a similarity metric breaches a threshold. *Id.*

225. Grotmol further discloses that, after breaching a similarity metric threshold, the robot (first device) determines and executes⁷ one or more potential

⁷ I understand that the '583 patent describes that this step may comprise simply causing an application to execute an instruction. '583 patent, 11:1-7 ("In certain embodiments, the causing the processor circuit to execute the one or more

control signals. The training set is obtained during a training operation of the robot (the instruction sets correlated with the first digital picture):

The method may comprise, **responsive to the similarity measure breaching a threshold, selecting one or more training sets from the plurality of training sets.** The method may comprise determining one or more potential control signals for the robot. Individual ones of the one or more potential control signals may be associated with a corresponding training set of the plurality of training sets.

In some implementations, the control signal may be configured to **cause the robot to execute the action.**

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.**

In some implementations, motor commands at a given time step may be obtained by **taking the stored motor commands from**

instruction sets for operating the device correlated with the first digital picture includes causing an application for operating the device to execute the one or more instruction sets for operating the device correlated with the first digital picture, the application running on the processor circuit.”).

the training buffer that may correspond to the best matching stored image.

Ex. 1005, 2:34-41, 3:12-13, Abstract, 43:5-9. Grotmol teaches that receiving or issuing a motor command causes the robotic vehicle to act. Because Grotmol discloses executing the motor commands, it is my opinion that following the selection of a motor control from a best matching image (e.g. determining) causes a first or second device to perform one or more operations defined by the instruction set.

226. A POSITA would understand that Grotmol also discloses that a second robot may execute actions learned on a first robot. Specifically, Grotmol discloses that training sets learned on a first robot may be loaded into a second robot:

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**.

See Id., 33:24-32. In such a case, autonomous operation of the second robot's processor executes actions defined by the training set learned on the first robot.

227. A POSITA would have understood that Grotmol discloses that at least in response to the determining, causing the first device or a second device (Grotmol discloses both) to perform one or more operations defined by the first one or more instruction sets for operating the first device.

5. **Claim 5: The system of claim 4, 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 device includes executing the first one or more instruction sets for operating the first device.**

Claim 5	Additional Features to Claim 4
	+ execute 1st instructions on 1st device

228. As I described above in limitations [4.3]-[4.4], Grotmol discloses determining a motor control command (instruction set) when a similarity metric threshold is breached when comparing a new and old digital image and further discloses executing instruction sets learned on a first robot on both that first robot and a second robot. Ex. 1005, 33:21-32; *see also id.*, 2:33-41, 3:12-13, 43:4-11.

229. Grotmol further discloses that, after a similarity measure is breached and a motor control command is determined, one or more potential control signals corresponding to the best matching stored image are executed.

The method may comprise, responsive to the similarity measure breaching a threshold, selecting one or more training sets from

the plurality of training sets. **The method may comprise determining one or more potential control signals for the robot.** Individual ones of the one or more potential control signals may be associated with a corresponding training set of the plurality of training sets.

In some implementations, the control signal may **be configured to cause the robot to execute the action.**

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.

Ex. 1005, 2:34-41, 3:12-13, 43:5-11.

230. A POSITA would understand that the claimed “causing” step in the ’583 patent is consistent with Grotmol because it may comprise simply causing an application to execute an instruction:

In certain embodiments, the **causing the processor circuit to execute the one or more instruction sets** for operating the device correlated with the first digital picture includes **causing an application** for operating the **device to execute the one or**

more instruction sets for operating the device correlated with the first digital picture, the application running on the processor circuit.

'583 patent, 11:1-7. Because Grotmol discloses executing the motor commands, it is my opinion that following the selection of a motor control from a best matching image (e.g. determining) causes a first or second device to execute the instruction set.

231. Accordingly, a POSITA would have understood that Grotmol discloses wherein the causing the first physical device or the second physical device to perform the one or more operations defined by the first one or more instruction sets for operating the first physical device includes executing the first one or more instruction sets for operating the first physical device.

- 6. **Claim 6: 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 portions of the new one or more digital pictures that represent one or more objects and one or more portions of the first one or more digital pictures that represent one or more objects.**

Claim 6	Additional Features to Claim 4
	+ Partial match between 1st DP and 2nd DP based on objects

232. As I described above in limitation [4.3], Grotmol discloses determining the learned instruction sets for operating the device by comparing a new image with a stored image, which includes matching portions of the images. Ex. 1005, 43:5-9, Abstract, 2:30-37, claim 1, 28:19-40; *see supra* limitation [4.3].

233. As an initial matter, the '583 patent provides examples of objects and object recognition in digital pictures such as:

In some exemplary embodiments, object recognition techniques and/or tools such as OpenCV (Open Source Computer Vision) library, CamFind API, Kooaba, 6px API, Dextro API, and/or others can be utilized for detecting or recognizing objects (**i.e. objects, animals, people, etc.**) in digital pictures. In some aspects, object recognition techniques and/or tools involve **identifying and/or analyzing object features** such as lines, edges, ridges, corners, blobs, regions, and/or their relative positions, sizes, shapes, etc., which may then be used to search for pictures with matching features. For example, OpenCV library can detect an object (**i.e. car, pedestrian, door, building, animal, person, etc.**) in one or more digital pictures captured by Picture Capturing Apparatus 90 or stored in an electronic repository, which can then be utilized in VSADO Unit 100, Artificial Intelligence Unit 110, and/or other elements.

'583 patent, 88:22-37.

234. A POSITA would have understood that Grotmol's object recognition disclosures are consistent with the usage of "objects" in the '583 patent:

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.




FIG. 20A

Ex. 1005, Fig. 20A

Ex. 1005, 38:56-61, Fig. 20A.

235. Grotmol further discloses that the camera used to capture these images may perform object recognition and provide information associated with the objects:

The sensor component 166 may be characterized by an aperture/field of view 168 (e.g., an extent of the observable world that may be captured by the sensor at a given moment). **The sensor component 166 may provide information associated with objects within the field-of-view 168.** In one or more implementations, **such as object recognition, and/or obstacle avoidance, the output provided by the sensor component 166 may comprise a stream of pixel values associated with one or more digital images.**

In one or more implementations, **object recognition and/or classification** may be implemented using a spiking neuron classifier comprising conditionally independent subsets

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).

Id., 8:64-9:3, 9:41-52, 10:23-26. Grotmol teaches that the autonomous operation of a robot may be based, in part, on recognized objects and their characteristics. In my opinion, Grotmol's object detection features and ability to operate based on said object detection, Grotmol identifies objects within a picture.

236. Grotmol further describes calculating the similarity measurement between one or more portions of new and stored images representing one or more objects and then determining whether a similarity metric threshold has been breached:

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.

[R]esponsive to the similarity measure breaching a threshold, **selecting one or more training sets from the plurality of training feature sets**

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**

Id., 2:30-37, claim 1, 28:19-40. A POSITA would have understood that because the similarity metric may compare features, and features may comprise objects and/or values of pixels in an image, Grotmol discloses comparing portions of images comprising objects to select a best match and corresponding instruction set.

237. Thus, a POSITA would have understood that Grotmol discloses determining at least a partial match between one or more objects and one or more portions of the first one or more digital pictures that represent one or more objects.

7. **Claim 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.**

Claim 7	Additional Features to Claim 4
	+ Objects detected in 1st DP and 2nd DP + Partial match of objects in 1st DP and 2nd DP

238. As I described above in limitation [4.3] and claim 6, Grotmol discloses calculating a similarity measure between one or more portions of new and stored images where the similarity measure may characterize similarity between features (e.g., objects) of the subset:

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.

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**

Images 2000, 2010 may comprise representations of visual field that may be obtained at times $t_1 < t_2$, respectively. **The image 2000 may comprise representations one or more objects**, e.g., a tree 2002, a rectangle, a ball 2008, and/or other objects.

Ex. 1005, 2:30-37, 28:19-40, 38:54-58. For the same reasons I described above in claim 6, Grotmol discloses determining an instruction set based on a partial match between one or more objects detected in a first digital picture and one or more objects in a new digital picture. *See supra* Section IX.B.6 (claim 6).

239. A POSITA would have understood that Grotmol discloses determining a partial match based on objects detected in new and stored images because Grotmol describes comparing features (e.g., objects) of digital images as part of the matching step.

8. **Claim 8: The system of claim 4, wherein the first one or more instruction sets for operating the first device include one or more information about one or more states of: the first device, or a portion of the first device.**

Claim 8	Additional Features to Claim 4
	+ State information of 1st device

240. As I described above in limitations [1.4]-[1.6], Grotmol discloses a memory component configured to store training sets comprising a digital picture and an instruction set (motor control command). Ex. 1005, 3:27-39, Abstract, 33:59-62.

241. Additionally, Grotmol discloses that the training set may include information about the state of the device:

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**.

Ex. 1005, 39:17-22. In my opinion, vehicle motion information, ambient environment information, and vehicle operational parameters comprise state information of a robotic vehicle. Grotmol teaches that such state information is stored with the training data. Accordingly, a POSITA would have understood that this “additional sensory information” corresponds to the claimed “information about one or more states” of the first device of the ’583 patent.

242. Accordingly, a POSITA would have understood that Grotmol discloses instruction sets including information about the state of the first device.

9. **Claim 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.

Claim 9	Additional Features to Claim 4
	+ Knowledgebase + 2nd DP correlated with 2nd instructions for 1st device + Learning process where user operates 1st device

243. As I described above in limitations [1.4]-[1.6], Grotmol discloses learning and storing associations between digital images and instruction sets. Ex. 1005, Abstract, 3:27-39, 39:13-22, 33:12-30. Further, said learning comprises a process where a human (“user”) manually operates a robot to learn the associations. *Id.*, 33:59-62, 43:5-9. Accordingly, a POSITA would have understood that Grotmol’s learning process includes operating a first device at least partially by a user.

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:

A robotic vehicle may utilize a camera and **stores the sequence of images** of a visual scene seen when following a trajectory

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.

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.

The robot may comprise a controller configured to store **a sequence of images** obtained during training phase.

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.**

Ex. 1005, Abstract, 38:38-39, 39:13-19; *see also id.*, 3:31-39, 8:66-9:3, 10:23-26.

Because Grotmol discloses sequentially storing images and corresponding motor control commands, it discloses storing multiple correlations. Therefore, a POSITA would have understood that Grotmol discloses a memory that stores a knowledgebase including correlations (training sets), where each correlation includes a digital picture correlated with an instruction set for operating the first device.

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.

246. Thus, a POSITA would have understood Grotmol discloses that the knowledgebase includes a first correlation learned in a first learning process and a second correlation learned in a second learning process that includes operating the first device at least partially by a first user.

10. **Claim 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.

Claim 11	Additional Features to Claim 4
	+ Knowledgebase + 2nd DP and 2nd instructions for 3rd device + 1st DP or 1st instructions learned on 1st device + 2nd DP or 2nd instructions learned on 3rd device

247. As I described above in limitation [1.6] and claim 9, Grotmol discloses that multiple learning processes (e.g., a user training a first device) can be done to train multiple behaviors that are stored in a knowledgebase on a first device. Ex. 1005, 3:27-39, 33:59-62, Abstract, 23:3-29, claim 12, 24:17-67, 26:1-20; *see supra* Sections IX.B.1.f, IX.B.1.9.

248. Grotmol further teaches that the training sets and training configurations from a first robot (including the instruction sets for operating the first device) can be “stored,” or “saved.” The saved configurations can then be loaded onto one or more robots, providing the training sets to a second or more robots (second and third device):

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.

Ex. 1005, 33:24-35. A POSITA would have understood that when a training configuration is loaded into a second one or more robots, the training configuration operates said second one or more robots. Further, once a training configuration is loaded into another robot (e.g. third device), a POSITA would understand that further training may be needed. In this case, any new correlations would be additional (e.g. second digital picture and second instruction set) learned in a separate process on another device.

249. Thus, a POSITA would have understood that Grotmol discloses that the knowledgebase of a first robot can store additional correlations, including a second correlation (second digital pictures correlated to second instruction set for operating another device) learned in a second learning process that includes operating another device by the user (i.e., first user).

- 11. Claim 12: The system of claim 4, wherein the new one or more digital pictures depict at least a portion of the first device's surrounding, and wherein the first one or more instruction sets for operating the first device are applied to the first device, and wherein the first device is caused to perform the one or more operations defined by the first one or more instruction sets for operating the first device.**

Claim 12	Additional Features to Claim 4
-----------------	---------------------------------------

	+ 2nd DP depicts 1st device's surrounding + 1st device performs operations defined in 1st instructions
--	---

250. As I described above for limitations [1.4]-[1.5], and [4.2], Grotmol discloses receiving new digital pictures (e.g. video frame instances) capturing the first device's surroundings and receiving instruction sets for operating a first device. Ex. 1005, Abstract, 3:27-39, 8:57-9:3, 38:28-60.

251. As I described above for limitations [4.3]-[4.4], claim 5, and claim 6, Grotmol discloses the first one or more instruction sets for operating the first device are applied to the first device, and the first 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 method may comprise, based on the comparison, **determining a similarity measure** for a given training set of the plurality of training feature sets. The similarity measure may characterize similarity between features of the subset and features of the given training set. The method may comprise, responsive to the similarity measure **breaching a threshold**, selecting one or more training sets from the plurality of training sets. The method may comprise determining one or more potential control signals for the robot.

[R]esponsive to the similarity measure breaching a threshold, selecting one or more training sets from the plurality of training feature sets

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**

In some implementations, the control signal may be configured to **cause the robot to execute the action.**

Ex. 1005, 2:30-38, claim 1, 28:19-40, 3:12-13. While Grotmol does not explicitly state that instruction sets are “applied,” a POSITA would have understood that the instruction sets are applied to a first device when they are selected to be performed by the first device.

252. Thus, a POSITA would have understood that Grotmol discloses applying instruction sets to a first device and causing operation of an instruction set on a first device.

12. **Claim 13: 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 are applied to**

the second 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.

Claim 13	Additional Features to Claim 4
	+ 2nd DP depicts 2nd device's surroundings + 2nd device performs operations defined in 1st instruction set

253. As I described above in claim 11, Grotmol discloses that training sets and training configurations can be loaded onto one or more robots to provide additional behaviors. Ex. 1005, 33:12-58.

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 partial match is found, meaning the threshold is breached, one or more potential control signals corresponding to the image from the training set that were obtained during a training operation of the robot (first one or more instruction sets for operating the first device correlated with the first digital picture) are selected (e.g. determined). Thus, the first instruction set is applied to the second device, causing the second device to perform the one or more operations defined by the first one or more instruction sets for operating the first device.

255. Thus, a POSITA would have understood that Grotmol discloses a second physical device that can receive training configurations (by loading stored configurations learned on a first device) and autonomously perform one or more operations defined by one or more instruction sets for operating the first physical. This saving, storing, and loading provides learned behaviors to one or more other robots.

13. **Claim 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.**

Claim 14	Additional Features to Claim 4
	+ 2nd DP depicts 2nd device's surroundings + Modifying 1st instructions and applying to 2nd device + Performing operation defined in modified 1st instructions on 2nd device

256. As I described above in claims 11 and 13, Grotmol discloses loading a saved training configuration (including instruction sets) into one or more robots:

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.

Ex. 1005, 33:24-35. Thus, a POSITA would understand that when the second device performs the matching step (*see supra* limitation [4.3]), the first instruction set is applied to the second device.

257. As an initial matter, the '583 patent describes that the instruction set may be modified by modifying the underlying source code, the hardware, or by modifying the application itself:

In further embodiments, the causing the processor circuit to execute the one or more instruction sets for operating the device correlated with the first digital picture includes modifying one or more instruction sets of an application, the application running on the processor circuit. In further embodiments, the causing the processor circuit to execute the one or more instruction sets for operating the device correlated with the first digital picture includes **modifying a source code, a bytecode, an intermediate code, a compiled code, an interpreted code, a translated code, a runtime code, an assembly code, or a machine code.** In

further embodiments, the causing the processor circuit to execute the one or more instruction sets for operating the device correlated with the first digital picture includes **modifying at least one of: the memory unit, a register of the processor circuit, a storage, or a repository where instruction sets are stored or used**. In further embodiments, the causing the processor circuit to execute the one or more instruction sets for operating the device correlated with the first digital picture includes **modifying one or more instruction sets for operating an application or an object of the application, the application running on the processor circuit**.

See, e.g., '583 patent, 11:28-49.

258. Grotmol further discloses that autonomous operation of a robot may include modifying training sets to compensate for, e.g., “position mismatch” and “orientation mismatch”:

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.**

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.**

Ex. 1005, 42:46-58, Abstract.

259. A POSITA would understand that when a training configuration is loaded into a second physical device—*see supra* claim 13, and the second robot encounters “position mismatch” or “orientation mismatch” from the images stored from the first robot, the instruction set is shifted. In execution, this would require the second robotic vehicle modifying the instruction set to accommodate said shift. A POSITA would understand that this shifting corresponds to the “modified instruction set.”

260. Thus, a POSITA would have understood that Grotmol discloses that the “first one or more instruction sets” are “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” —*see supra* claim 4.

14. Claim 15

Claim 15	Additional Features to Claim 4
	+ Modifying 1st instructions + Performing modified 1st instructions

- a. **[15.1] 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 at least:**

261. Grotmol discloses limitation [15.1] for the same reasons I discussed above in Section IX.B.1.c, *supra*, for limitation [1.3].

- b. **[15.2] modifying: the first one or more instruction sets for operating the first device, or ...**

262. Grotmol further teaches that autonomous operation of a robot may include modifying training sets to compensate for, e.g., “position mismatch” and “orientation mismatch”:

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.**

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.**

Ex. 1005, 42:46-58, Abstract. Further, the “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 image,” and “[t]hose motor commands may be combined with the output from the PID controller in order to stabilize operation of the robot.” *Id.*, 43:9-11

263. In execution, adjusting for a position mismatch or orientation mismatch would require modifying the instruction set to accommodate said shift. A POSITA

would understand that this shifting corresponds to the “modified instruction set.” Thus, a POSITA would have understood that Grotmol discloses processors modifying the first one or more instruction sets for operating the first device.

- c. **[15.3] and wherein the determining ... 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 ...**

264. As I described above for limitations [4.3] and [15.2] and claim 14, 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 may correspond to the best matching stored image and the motor commands may be modified to compensate for, e.g., position mismatch and orientation mismatch. Ex. 1005, 43:5-9, Abstract, 2:30-37, 42:46-43:3; *see supra* limitations [14.3], [15.2], claim 14. 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.

265. Accordingly, a POSITA would have understood that Grotmol discloses 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.

- d. **[15.4] and wherein the causing ... 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**

266. As I described above for limitations [4.3], [4.4], [15.2], and [15.3], 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 may correspond to the best matching stored image and the motor commands may be modified to compensate for, e.g., position mismatch and orientation mismatch, and the robot (i.e., first device) can autonomously perform the motor commands, including the modified motor commands. Ex. 1005, 43:5-9, Abstract, 2:30-37, 42:46-43:3. A POSITA would understand that selecting a motor control command adjusting for said position mismatch or orientation mismatch (i.e. executing a shift to adjust for said mismatch) would correspond to performing one or more operations defined by a modified instruction set.

267. Thus, a POSITA would have understood that Grotmol discloses causing the first device to perform one or more operations defined by the modified the first one or more instruction sets for operating the first device.

- 15. **Claim 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 ... and wherein the at least partial match between the new one or more digital pictures and the first one or more digital pictures includes ... (iii) at least partial match between the new one or more digital pictures and the modified 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.

Claim 16	Additional Features to Claim 4
	+ Modifying 1st DP + Partial match between 2nd DP and modified 1st DP

268. Grotmol discloses that captured and stored digital pictures may undergo preprocessing operations:

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.**

Ex. 1005, 39:13-17.

269. Grotmol further discloses that pixels from individual image instances are selected, scaled, normalized, and/or equalized:

The method may comprise applying a transformation to an instance of the sensory input. The transformation may be

configured to produce **scaled input** based on analysis of the history.

In order to facilitate contributions from different types of signals for determining a distance measure between features in a metric space (e.g., Euclidian distance), **data** from highly redundant input (e.g., the RGB pixel values) **may be normalized**.

Equalization may be applied so that contribution of individual sensory modality on expected distances may be comparable relative contribution from another modality data.

[S]caling individual pixels of the set of pixels by a scaling factor; **scaling features of the subset by a scaling factor**

Ex. 1005, 4:18-22, 30:47-51, 31:61-64, 3:44-58; *see also id.*, 4:58-62, 31:3-8. A 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.

270. Accordingly, a POSITA would have understood that Grotmol discloses determining whether there is a partial match between the new digital picture and modified first digital picture.

16. Claim 17

Claim 17	Additional Features to Claim 4
	+ Modifying 1st DP + Learning correlation modified 1st DP and 1st instructions + Partial match between 2nd DP and modified 1st DP

- a. **[17.1] 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:**

271. Grotmol discloses limitation [17.1] for the same reasons I discussed above in Section IX.B.1.c, supra, for limitation [1.3].

- b. **[17.2] modifying: the first one or more digital pictures, or ...**

272. Grotmol discloses that captured and stored digital pictures may undergo preprocessing operations:

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.**

Ex. 1005, 39:13-17.

273. Grotmol further discloses that pixels from individual image instances are selected, scaled, normalized, and/or equalized:

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

In order to facilitate contributions from different types of signals for determining a distance measure between features in a metric space (e.g., Euclidian distance), **data** from highly redundant input (e.g., the RGB pixel values) **may be normalized**.

Equalization may be applied so that contribution of individual sensory modality on expected distances may be comparable relative contribution from another modality data.

[S]caling individual pixels of the set of pixels by a scaling factor; **scaling features of the subset by a scaling factor**

Id., 4:18-22, 30:47-51, 31:61-64, 3:44-58; *see also id.*, 4:58-62, 31:3-8. In my opinion, pre-processing, normalizing, and equalizing of digital images comprises the

claimed modifying digital pictures. Accordingly, stored digital pictures are stored with said modifications.

274. Thus, a POSITA would have understood that Grotmol discloses modifying captured or stored digital pictures.

- c. **[17.3] and wherein the learning ... 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 ...**

275. As I described above in limitation [1.6], Grotmol discloses a learning process that includes learning correlations between digital pictures and instruction sets for operating the robot (first device):

Apparatus and methods **for hierarchical learning** are disclosed herein, in accordance with one or more implementations. **In supervised learning, a user may train** simple tasks by demonstrating example tasks to the system.

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

Ex. 1005, 24:17-20, 33:59-62. *See also id.*, claim 12 (reciting “one or more physical processors configured to operate a random k-nearest neighbors *learning process*”).

276. As I described above in limitation [17.2], Grotmol discloses modifying digital pictures, including the first one or more digital pictures that are received during the training process:

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.**

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

In order to facilitate contributions from different types of signals for determining a distance measure between features in a metric space (e.g., Euclidian distance), **data** from highly redundant input (e.g., the RGB pixel values) **may be normalized.**

Equalization may be applied so that contribution of individual sensory modality on expected distances may be comparable relative contribution from another modality data.

[S]caling individual pixels of the set of pixels by a scaling factor; **scaling features of the subset by a scaling factor**

Ex. 1005, 39:13-17, 4:18-22, 30:47-51, 31:61-64, 3:44-58; *see also id.*, 4:58-62, 31:3-8. A POSITA would have understood that the pre-processing, normalizing, and equalizing of digital images comprises modifying images. Because Grotmol discloses modifying images to be stored in the training buffer, and using images in the training buffer to select a best matching image and correlated instruction set, in my opinion Grotmol's training requires learning the correlation of a modified digital picture with the instruction set.

277. Accordingly, a POSITA would have understood that Grotmol discloses learning the modified the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device.

- d. **[17.4] and wherein the determining ... 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**

278. As I described above for limitation [4.3], Grotmol discloses motor commands at a given step may be obtained from the training buffer based on best matching image:

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

[R]esponsive to the similarity measure breaching a threshold, selecting one or more training sets from the plurality of training feature sets

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**

Ex. 1005, 2:30-37, claim 1, 28:19-40.

279. A POSITA would have understood that where the first one or more digital pictures stored during the learning process had been modified, the partial match would be determined using the current image (i.e., new one or more digital

pictures) and the modified stored image (i.e., modified first one or more digital pictures). Accordingly, in my opinion, the instruction set is determined based on a partial match between a modified first digital picture and a new digital picture.

280. Thus, a POSITA would have understood that Grotmol discloses 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.

17. Claim 18

Claim 18	Additional Features to Claim 4
	+ Modifying 1st instructions + Learning correlated 1st DP and modified 1st instructions + Performing operation defined by modified 1st instructions

- a. **[18.1] 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:**

281. Grotmol discloses limitation [18.1] for the same reasons I discussed above in Section IX.B.1.c, *supra*, for limitation [1.3].

b. [18.2] modifying: the first one or more instruction sets for operating the first device, or ...

282. Grotmol discloses that autonomous operation of a robot may include processors modifying training sets to compensate for, e.g., “position mismatch” and “orientation mismatch”:

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.**

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.

Ex. 1005, 42:46-58, Abstract.

283. For example, Grotmol discloses that instruction sets may be modified to stabilize 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.**

The best match image obtained using image match methodology (e.g., such as described herein) may be used to determine **changes (corrections) to motor commands** during path navigation by a robotic vehicle.

Id., 43:4-11, 42:37-40. A POSITA would have understood that combining motor commands with additional instructions (e.g., output from the PID controller) corresponds to the claimed “modifying: the first ... instruction sets.” In execution, adjusting for a position mismatch or orientation mismatch would require modifying

the instruction set to accommodate said shift. A POSITA would understand that this shifting corresponds to the “modified instruction set.”

284. Further, Grotmol discloses a fine tuning process in which the robot operates autonomously until a teacher issues a command, in which case the robot will apply the teacher’s command and attempt to learn the correlation:

In listen mode, the robot may be “listening” to or otherwise observing what the teacher teaches, and may not do anything on its own. However, the robot may learn an association. But if the teacher stops sending a command, the robot may stay still and wait for the next command. In some implementations, teaching the robot may begin with the listen mode. Once enough examples have been provided and the robot has learned something, the override mode may be used.

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.

Ex. 1005, 23:13-22. In my opinion, because the robot previously was operating autonomously, when the robot receives a corrective command and remembers the new command it is also modifying the instruction set.

285. Thus, a POSITA would have understood that Grotmol discloses modifying the first one or more instruction sets.

- c. **[18.3] and wherein the learning ... 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 ...**

286. As I described above for limitations [1.6], [4.3]-[4.4], and [18.2], Grotmol discloses that motor commands at a given step may be obtained from the training buffer based on best matching image, the motor commands may be modified, and the motor commands are executed to perform autonomous operation of the robotic vehicle. Ex. 1005, 3:27-39, 33:59-62, Abstract, 43:5-9, 42:46-58, 42:37-40, 34:3-27, 24:55-25:5, 34:49-55, 23:12-22. Further, because Grotmol discloses fine tuning behaviors through further training, it is my opinion that the system would learn the new association of a first digital picture and a modified instruction set.

287. Accordingly, a POSITA would have understood that Grotmol discloses learning a first one or more digital pictures correlated with modified instruction sets for operating a device.

- d. **[18.4] and wherein the determining ... 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 ...**

288. As I described above for limitations [4.3], [18.2], and [18.3], Grotmol discloses that motor commands at a given step may be obtained from the training buffer based on best matching image, and motor commands can be modified to compensate for, e.g., “position mismatch” and “orientation mismatch” between the new and stored images and/or by receiving a corrective instruction. Ex. 1005, 3:27-39, 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 determining the modified instruction set.

289. Thus, a POSITA would have understood that Grotmol discloses determining a modified instruction set based on a partial match between a new and old digital picture and the first one or more digital pictures.

- e. **[18.5] and wherein the causing ... 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**

290. As I described above for limitations [4.3], [18.2], and [18.3], Grotmol discloses that motor commands at a given step may be obtained from the training

buffer based on best matching image, and motor commands can be modified to compensate for, e.g., “position mismatch” and “orientation mismatch” between the new and stored images and/or by receiving a corrective instruction. Ex. 1005, 3:27-39, 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.

291. Thus, a POSITA would have understood that Grotmol discloses causing the first or second device to perform the modified instruction set.

18. **Claim 19: The system of claim 4, wherein the system further comprising: an artificial intelligence system, 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, at least in part by the artificial intelligence system, 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, at least in part by the artificial intelligence system, 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.**

Claim 19	Additional Features to Claim 4
-----------------	---------------------------------------

	+ AI system
--	-------------

292. I understand that claim 19 depends from claim 4 and requires: (i) the “learning” of the first correlation (digital image with instruction set); and (ii) the determining of the first instruction set are done at least in part by an artificial intelligence system. ’583 patent, claim 19.

293. As I described above for limitation [1.6], Grotmol discloses a system in which processors perform the operation of learning correlations between digital pictures and instruction sets for operating the first device. Ex. 1005, 33:59-62, 39:17-19. Further, as I described above for limitation [4.3], Grotmol discloses a system in which one or more processors perform the operation of determining an instruction by calculating a similarity measure between a current image and a stored image, then determining whether that similarity measure breaches a threshold. *Id.*, Abstract, 43:5-9, 2:30-37, 28:19-40.

294. As an initial matter, the ’583 patent describes an “artificial intelligence unit” as a generic piece of hardware or software. *See generally*, ’583 patent, 2:47-3:22. I understand that neither claim 19 nor the specification of the ’583 patent limits the artificial intelligence system to a specific artificial intelligence model or

technique, indeed stating that “any model and/or technique can be utilized to facilitate the functionalities described herein”:

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.** The disclosed systems, devices, and methods are independent of the artificial intelligence model and/or technique used and any model and/or technique can be utilized to facilitate the functionalities described herein. **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.**

’583 patent, 95:49-64.

295. The ’583 patent further identifies supervised learning and unsupervised learning as examples of artificial intelligence models or techniques that can be used for, *inter alia*, determining a partial match between new and old digital images:

The disclosed systems, devices, and methods are independent of the **artificial intelligence model** and/or technique used and any

model and/or technique can be utilized to facilitate the functionalities described herein. Examples of these models and/or techniques include deep learning, **supervised learning, unsupervised learning**

Similarity Comparison 125 comprises the functionality for comparing or matching Knowledge Cells 800 or portions thereof, and/or other functionalities. Similarity Comparison 125 comprises the functionality for comparing or matching Digital Pictures 525 or portions thereof. Similarity Comparison 125 comprises the functionality for comparing or matching streams of Digital Pictures 525 or portions thereof. Similarity Comparison 125 comprises the functionality for comparing or matching Instruction Sets 526, Extra Info 527, text (i.e. characters, words, phrases, etc.), pictures, sounds, data, and/or other elements or portions thereof. **Similarity Comparison 125 may include functions, rules, and/or logic for performing matching or comparisons and for determining that while a perfect match is not found, a similar or partial match has been found.** In some aspects, Similarity Comparison 125 may include determining substantial similarity or substantial match of compared elements. In other aspects, a partial match may include a substantial or otherwise similar match, and vice versa. Although, substantial similarity or substantial match is frequently used herein, it should be understood that any level of similarity, however high or low, may be utilized as defined by

the rules (i.e. thresholds, etc.) for similarity. **The rules for similarity or similar match can be defined by a user, by VSADO system administrator, or automatically by the system based on experience, testing, inquiry, analysis, synthesis, or other techniques, knowledge, or input. In some designs, Similarity Comparison 125 comprises the functionality to automatically define appropriately strict rules for determining similarity of the compared elements.**

'583 patent, 95:53-64, 102:29-58.

296. When anticipating instruction sets based on a partial match, the '583 patent states that “any level of similarity, however high or low, may be utilized” and the comparison can compare, e.g. regions, pixels, objects, or portions of images. *Id.*, 102:47-63, 105:58-107:14. Given these broad disclosures, in my opinion, a POSITA would not consider the '583 patent's instruction anticipation process to require any specific artificial intelligence model or technique.

297. While Grotmol does explicitly disclose an “artificial intelligence system,” a POSITA would have understood that Grotmol discloses an “artificial intelligence unit” as a part of, or operating on, its processor circuit. *See, e.g.*, Ex. 1005, claim 12 (reciting “one or more physical processors configured to operate a random k-nearest neighbors learning process to determine a motor control indication”). In my opinion, the inclusion of processors configured to operate a k-

nearest neighbors learning process is consistent with the artificial intelligence system description in the '583 patent above.

298. Grotmol further describes that, like in the '583 patent, either supervised or unsupervised learning may be utilized:

In some implementations, the control entity may comprise a **computerized agent** such as a multifunction adaptive controller operable using reinforcement **and/or unsupervised learning and capable of training other robotic devices for one and/or multiple tasks**. In one such implementation, the control entity and the corrector 412 may comprise a single computerized apparatus.

Apparatus and methods for hierarchical learning are disclosed herein, in accordance with one or more implementations. **In supervised learning**, a user may train simple tasks by demonstrating example tasks to the system. To achieve an acceptable performance of a complex task by the system, the user may provide additional information to the system, in addition to examples.

Ex. 1005, 16:8-14, 24:17-30. A POSITA would understand that the learning processes described in Grotmol, akin to the learning processes described in the '583 patent, correspond to claimed “artificial intelligence system” learning the correlation of digital picture and instruction sets and determining the instruction set to be

executed. Accordingly, in my opinion, Grotmol discloses that its artificial intelligence system (e.g. processors) learns the correlation of a digital picture and motor control system and further causes the execution of the instruction set.

299. Thus, a POSITA would have understood that Grotmol discloses an artificial intelligence system that performs the learning and determining steps.

- 19. Claim 20: The system of claim 19, wherein the artificial intelligence system includes: one or more inputs for receiving one or more digital pictures, and one or more outputs for providing one or more instruction sets, and wherein the learning ... includes applying the first one or more digital pictures to the inputs of the artificial intelligence system, and wherein the determining ... includes: applying the new one or more digital pictures to the inputs of the artificial intelligence system; and receiving the first one or more instruction sets for operating the first device from the outputs of the artificial intelligence system.**

Claim 20	Additional Features to Claim 19
	+ AI system + Input to receive DPs + Output instructions

300. As shown below in Figure 8 of Grotmol below, the processor is connected to the “sensor” functionality, which may include camera input:




FIG. 8

Ex. 1005, Fig. 8

In one or more implementations, such as object recognition, and/or obstacle avoidance, **the output** provided by the sensor component 166 **may comprise** a stream of pixel values associated with **one or more digital images**.

Ex. 1005, Fig. 8, 8:66-9:3. Both Grotmol and the '583 patent disclose that the digital picture data may come from, for example, a CCD image sensor. *Compare* '583 patent, 70:2-13, *with* Ex. 1005, 9:8-22.

301. As I described above in limitation [2.1], Grotmol discloses receiving digital pictures from a camera. Ex. 1005, 10:16-22 (“In some implementations of robotic navigation in an arbitrary environment, the sensor component 166 may comprise a camera configured to provide an output comprising a plurality of digital

image frames ...”). Grotmol further discloses an adaptive control system which receives sensory input (402 in Figure 4A below). It is my opinion that the sensory input corresponds to the claimed input(s) for receiving one or more digital pictures. The adaptive control system’s corrector, predictor, and combiner produce to output the motor control command:

FIG. 4A illustrates an implementation of **adaptive control system 400**. The adaptive control system 400 of FIG. 4A may comprise a corrector 412, an adaptive predictor 422, and a combiner 414 cooperating to control a robotic platform 430. The learning process of the adaptive predictor 422 may comprise a supervised learning process, a reinforcement 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.




FIG. 4A

Ex. 1005, Fig. 4A

The corrector 412 may receive a control input 428 from a control entity. The control input 428 may be determined based on one or more of (i) **sensory input 402** and (ii) feedback from the platform (not shown).

Ex. 1005, 15:19-33, Fig. 4A, 15:63-66. The output to the adaptive control system (e.g. artificial intelligence system) is a control signal, in my opinion this corresponds to the claimed “output.” Thus, a POSITA would have understood that Grotmol discloses one or more inputs for receiving one or more digital pictures, and one or more outputs for providing one or more instruction sets.

302. Second, Grotmol discloses “learning” by applying the inputs to the digital pictures and “determining” by receiving instruction sets from the outputs of the artificial intelligence system. As I described above in limitations [1.2]-[1.4] and [4.3]-[4.4], Grotmol discloses that the learning process includes receiving new digital pictures that are compared to stored digital pictures, and further that the processor includes an input for receiving images, and that instruction sets are determined and cause the robot to perform. Ex. 1005, Abstract, 3:27-39, 39:23-28, 43:5-9, 2:30-37, 3:12-13. In my opinion, Grotmol’s selection (e.g. determining) of the motor control signal requires inputting a digital picture (e.g., applying an input to the artificial intelligence system) and the artificial intelligence system producing a motor control command (e.g. receiving the instruction set as an output from the artificial intelligence system).

303. Thus, a POSITA would have understood that Grotmol discloses an artificial intelligence system applying one or more inputs for receiving one or more digital pictures, and one or more outputs for providing one or more instruction sets.

- 20. Claim 21: The system of claim 20, wherein the one or more inputs for receiving the one or more digital pictures include: one input for one digital picture, multiple inputs for multiple digital pictures, one input for one portion of one digital picture, multiple inputs for multiple portions of one digital picture, multiple inputs for multiple portions of multiple digital pictures, an input for one representation of one digital picture, multiple inputs for multiple representations of multiple digital pictures, one input for one representation of one portion of one digital picture,**

multiple inputs for multiple representations of multiple portions of one digital picture, or multiple inputs for multiple representations of multiple portions of multiple digital pictures.

Claim 21	Additional Features to Claim 20
	+ One input for one DP

304. Claim 21 requires one input for one digital picture, or multiple inputs for multiple digital pictures. Grotmol discloses claim 21 for the same reasons I discussed above in Section IX.B.20, *supra*, for claim 20.

21. Claim 22: The system of claim 19, wherein the artificial intelligence system includes: a knowledgebase, or a neural network.

Claim 22	Additional Features to Claim 19
	+ Knowledgebase

305. Claim 22 requires an artificial intelligence system which includes a knowledge base. As I described above in claim 9, Grotmol discloses a knowledgebase. Grotmol discloses claim 22 for the same reasons I discussed above in Section IX.B.9, *supra*, for claim 9.

22. Claim 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.

Claim 23	Additional Features to Claim 4
	+ Connecting 1st DP and 1st instructions

306. As I described above in claim 1, Grotmol discloses storing correlated digital pictures and instruction sets as training sets, which may be configured based on associations between the image and the instruction set. Ex. 1005, 33:24-27 (“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 training sets “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.” *Id.*, 33:59-62.

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

⁹ I understand that the ’583 patent describes that the connection between the digital picture and the instruction set “may include or is associated with at least one of: an occurrence count, a weight, a parameter, or a data.” ’583 patent, 15:12-14.

in memory, for example, would comprise “one or more connections” as recited in the claim:

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.

Motor commands associated with a given image may **be stored**.

Motor commands issued by a trainer corresponding to time instance when the images are acquired may be **stored**.

Ex. 1005, 3:35-39, Abstract, 39:13-19. As I noted above, I understand that an ordered buffer is a region of memory that is allocated for storing data in a specific order or sequence. A POSITA would have understood Grotmol to describe that the training buffer may be ordered temporally, according to when the images were acquired, due to its disclosures of images in the training buffer that precede and follow each other in time:

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

time in the training buffer, and one or more (e.g., 5) images **following the IO image in time** in the training buffer.

Motor commands issued by a trainer corresponding to time instances when the images are acquired may be stored.

Ex. 1005, 39:65-40:4, 39:17-19. Thus, a POSITA would understand that Grotmol describes connected knowledge cells because it describes knowledge cells stored in temporal order – i.e., order one after another based on when the image-instruction correlation was established.

308. Thus, a POSITA would have understood that Grotmol discloses a system “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.”

23. Claim 24: The system of claim 4, wherein the first one or more digital pictures are: one or more whole digital pictures, ... one or more collections of pixels, or one or more collections of values.

Claim 24	Additional Features to Claim 4
	+ 1st DP = whole image or collection of pixels

309. As I described above in limitations [1.4] and [4.3], Grotmol discloses digital pictures, collections of pixels, or other values. Ex. 1005, 8:61-66, 3:14-16, 41:8-17. Thus, a POSITA would have understood that Grotmol requires that the first digital picture is a whole digital picture, a collection of pixels, or a collection of values. *See supra* limitations [1.4], [4.3].

24. Claim 26

Claim 26	Additional Features to Claim 4
	+ Various embodiments

- a. **[26.1] The system of claim 4, wherein at least some elements of the system are included in: a single device, or multiple devices,**

310. Grotmol discloses limitation [26.1] for the same reasons I discussed above in Section IX.B.1.c, *supra*, for limitation [1.3].

- b. **[26.2] and wherein the one or more processors include: one or more microcontrollers, one or more computing devices, or one or more electronic devices,**

311. A POSITA would understand that Grotmol discloses several processor circuits which would include one or more microcontrollers, one or more computing devices, or one or more electronic devices.

312. Grotmol states that the terms “‘processor’, ‘microprocessor’ and ‘digital processor’ are meant generally to include all types of digital processing

devices.” Ex. 1005, 8:5-15. A POSITA would understand that Grotmol gives examples of “digital processing devices” such as digital signal processors (DSPs), field programmable gate arrays (FPGAs), programmable logic devices (PLDs), and application-specific integrated circuits (ASICs). *Id.* Grotmol further states that “the terms ‘computer’, ‘computing device’, and ‘computerized device’, include . . . literally any other device capable of executing a set of instructions and processing an incoming data signal.” *Id.*, 7:30-42. In my opinion, a microcontroller is both a “digital processing device” and a “device capable of executing a set of instructions and processing an incoming data signal.” Thus, a POSITA would understand that Grotmol discloses a system in which the “one or more processors” includes “one or more microcontrollers.”

313. For example, Grotmol discloses adaptive controllers that are designed to process and execute specific tasks. Grotmol discloses an “Adaptive controller” of type “1” that may learn to approach the target autonomously generating correction torques on robot’s wheels and may not control the gripper (“output N”). *See, e.g.*, Ex. 1005, 25:30-42. Grotmol also discloses an “‘Adaptive controller’ of type ‘2’” that is specifically designed to learn to turn the robot and appropriately grasp targets by autonomously effectuating correct torques on the robot’s wheels and gripper while not controlling another wheel. *See, e.g., id.*, 26:5-33, 24:31-42. The inclusion

of multiple adaptive controllers indicates at least one or more processors or at least one or more microcontrollers.

314. Thus, a POSITA would have understood Grotmol's system to be implemented with processors including one or more microcontrollers.

c. [26.3] and wherein the first device includes a robot, a vehicle, an appliance, an electronic device, or a mechanical machine,

315. Grotmol discloses the first device includes a robotic vehicle (e.g., a "self-contained robotic apparatus"):

Another aspect of the disclosure relates to a **self-contained robotic apparatus**.

In some implementations, the robot may comprise a **robotic vehicle** (e.g., 160 in FIG.1) comprising a camera (e.g., 166).

Ex. 1005, 3:27-28, 38:36-37.

d. [26.4] and wherein the second device includes a robot, a vehicle, an appliance, an electronic device, or a mechanical machine,

316. A POSITA would understand that Grotmol discloses a second device that includes a robotic vehicle. *See supra* limitation [26.3]; Ex. 1005, 3:27-28, 38:36-37.

317. Grotmol discloses that the training sets and training configurations from a first robot (including the instruction sets for operating the first device) can be

“stored,” or “saved.” The saved configurations can then be loaded onto one or more robots, providing the training sets to a second or more robots (second and third device):

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.

Ex. 1005, 33:24-35. When a training configuration is saved on a first device and then loaded into a second device, in my opinion, Grotmol is disclosing a system contemplating a second robotic vehicle.

318. Thus, a POSITA would understand that Grotmol discloses a second device (“one or more other robots”) that can autonomously perform one or more operations defined by one or more instruction sets for operating the first device training configuration from first robot (e.g., training sets with circumstance representations correlated to instruction sets) that is loaded into a second one or more other robots to “provide learned behaviors.”

- e. **[26.5] and wherein the first one or more digital pictures include: one or more still digital pictures, or one or more motion digital pictures, and wherein the new one or more digital pictures include: one or more still digital pictures, or one or more motion digital pictures,**

319. As I described above in limitations [1.4], [1.6], [2.1], and [4.2], Grotmol discloses storing digital pictures, including still pictures or video frame instances.

Ex. 1005, 3:27-39, Abstract, 10:15-19, 38:52-58.

- f. **[26.6] and wherein an instruction set of the first one or more instruction sets for operating the first device includes at least one of: only one instruction, a plurality of instructions, one or more inputs, one or more commands, one or more keywords, one or more symbols, one or more operators, one or more variables, one or more parameters, one or more characters, one or more numbers, one or more values, one or more signals, one or more binary bits, one or more functions, one or more function references, one or more objects, one or more object references, one or more data structures, one or more data structure references, one or more states, one or more representations of one or more states, one or more representations of one or more user inputs, one or more codes, one or more data, or one or more information.**

320. As I described above for limitations [1.5], [1.6], [2.2], and [4.3], Grotmol discloses an instruction set for operating the first device, where the instruction set includes, for example, one or more instructions, one or more commands, and one or more information. Ex. 1005, 3:27-39, 39:17-19, 33:12-30.

Specifically, Grotmol discloses using the system “issu[ing]” “motor commands.”

Id., Abstract.

25. **Claim 28:** The system of claim 4, wherein the at least partial match between the new one or more digital pictures and the first one or more digital pictures is at least in part based on a product of one or more values that represent one or more portions of the new one or more digital pictures and one or more values that represent one or more portions of the first one or more digital pictures.

Claim 28	Additional Features to Claim 4
	+ Knowledgebase + 2nd instructions for operating 3rd device + Performing 2nd instructions on 3rd device

321. As I described above in limitation [4.3], a POSITA would understand Grotmol’s matching process as consistent with the claimed partial match step in the ’583 patent because both may comprise pixel-wise difference analysis between two pictures. *Compare*, Ex. 1005, 3:13-15, 41:8-17, *with* ’583 patent, 21:29-46, 112:38-113:2.

322. Grotmol further discloses that the similarity between the new and stored images can be calculated using spatial frequency representation. Each image is processed using a windowing function and a Fast Fourier Transform (FFT), and from the resulting frequency representations, a normalized cross-power spectrum is obtained.

In some implementations, the **cross-correlation between two images may be determined by utilizing the spatial frequency domain**. A windowing function (e.g., Hann, Gaussian, cosine, Hamming, and/or other windowing function) may be applied to individual images to produce windowed image and reduce edge effects. A fast-Fourier transform (FFT) may be performed on the windowed images to obtain a spatial frequency representation of the images.

Ex. 1005, 40:34-41. A POSITA would understand that the cross-power spectrum is calculated by multiplying the FFT of one image with the complex conjugate of the FFT of the other image.

323. Grotmol further discloses that the inverse FFT is then applied to the cross-power spectrum to change it back from the frequency domain to the spatial domain, resulting in a cross-correlation function that shows the degree of similarity between the two images for various amounts of shift. The maximum value (argmax) of the cross-correlation function is then found, which corresponds to the shift (x,y values) that gives the best match between the two images:

Normalized cross-power spectrum may be determined from the two spatial frequency representations. An inverse FFT may be applied to transform the cross spectrum to x,y domain and to obtain the cross-correlation. The argmax of the cross-correlation may be determined in order to obtain x,y coordinates (shift values) corresponding to maximum cross-correlation. In some

implementations wherein x,y dimension may correspond to integer values (e.g., 1 pixel), the cross-correlation matrix may be interpolated onto a grid with greater resolution (e.g., 0.5 or 0.25 pixel grid).

Ex. 1005, 40:42-51.

324. A POSITA would understand the spatial frequency representation calculations (i.e., when calculating the cross-power spectrum by multiplying the fast-Fourier transform of one image with the complex conjugate of the fast-Fourier transform of the other image) described in Grotmol disclose the matching process based on a product of value representing portions of new and old digital image claimed in the '583 patent.

325. Thus, Grotmol discloses a partial match based on a product of values representing portions of a first and stored digital picture.

26. Claim 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.

Claim 29	Additional Features to Claim 4
-----------------	---------------------------------------

	+ Detecting object in 2nd DP + Generating object representation + Generating 2nd DP image including object representation
--	---

326. As I described above in claim 6, the '583 patent provides examples of objects in pictures such as “car, pedestrian, door, building, animal, person, etc.” '583 patent, 88:22-37. The '583 patent further describes various approaches to object recognition, such as “identifying and/or analyzing object features such as lines, edges, ridges, corners, blobs, regions, and/or their relative positions, sizes, shapes, etc.” *Id.*

327. Consistent with the disclosures of the '583 patent, Grotmol discloses that images may comprise representations of objects:

The image 2000 may comprise **representations one or more objects**, e.g., a tree 2002, a rectangle, a ball 2008, and/or other objects.




FIG. 20A
Ex. 1005, Fig. 20A

Ex. 1005, 38:56-58, Fig. 20A.

328. Further consistent with the disclosures in the '583 patent, Grotmol discloses that the sensor used to capture these images may provide information about objects in the frame, including object recognition and edge detection.

The sensor component 166 may be characterized by an aperture/field of view 168 (e.g., an extent of the observable world that may be captured by the sensor at a given moment). **The sensor component 166 may provide information associated with objects within the field-of-view 168.** In one or more implementations, **such as object recognition, and/or obstacle avoidance**, the output provided by the sensor component 166 may comprise a stream of pixel values associated with one or more digital images.

In one or more implementations, **object recognition and/or classification may be implemented** using a spiking neuron classifier comprising conditionally independent subsets

As used herein the term **feature** may refer to a representation of an **object edge**, determined by change in color, luminance, brightness, transparency, texture, and/or curvature. The object features may comprise, inter alia, **individual edges, intersections of edges (such as corners), orifices, and/or curvature.**

Ex. 1005, 8:61-9:3, 9:41-52, 7:60-65. A POSITA would understand that object recognition requires detecting one or more objects in the robotic vehicle's surroundings.

329. Grotmol describes that the autonomous operation of a robot may be based in part on recognized objects and their characteristics and that the images may include representations of the objects detected in the device's surroundings:

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.

The user may have an opportunity to select which input is relevant for a particular behavior. For example, in some implementations, **only the target and obstacle detection may be relevant** to train a target approach, and/or only base and obstacle detection may be relevant to train a base approach. The user may select which actuators are controlled by a particular behavior. For example, navigation tasks may not need to learn to control the gripper. The user may have an opportunity to use a specialized predictor for different low level tasks. For example, a non-linear classifier may be used for target approach and a linear classifier may be used for grasping.

Ex. 1005, 10:23-34, 24:61-25:5. Because Grotmol discloses performing behaviors based on object detection, Grotmol requires storing object data within a total digital picture representation.

330. Thus, a POSITA would understand that Grotmol's object detection corresponds to "detecting one or more objects in the first device's surrounding or

the second device's surrounding," and that Grotmol requires generating a new digital picture including one or more representations of one or more objects.

27. **Claim 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.**

Claim 31	Additional Features to Claim 4
	+ 1st instructions as representation of 2nd instructions

331. As an initial matter, 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* '583 patent, 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.").

332. Grotmol discloses the ability to UNDO, or unlearn behaviors, whereby the instruction set is overwritten:

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.

A user may select a given bookmark in order to restore (undo) the learning configuration to the state corresponding to time of the bookmark.

Ex. 1005, 34:18-21, 34:43-45. A POSITA would understand that re-loading a previous instruction set (second instruction set) is a representation of that instruction set (second instruction set). Further, when the system then next selects that instruction set (i.e. generating the first instruction set) during the matching step, a POSITA would understand that the system is generating a representation of the second instruction set.

333. Accordingly, a POSITA would understand that Grotmol discloses 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.

28. Claim 32: The system of claim 4, wherein the first one or more digital pictures that depict the at least the portion of the first device's surrounding include one or more digital

pictures that depict a representation of the at least the portion of the first device's surrounding.

Claim 32	Additional Features to Claim 4
	1st DP includes 1st device's surroundings

334. As I described above in limitation [1.4] and claim 29, Grotmol discloses digital pictures that depict a representation of at least a portion of the first device's surrounding. *See supra* limitation [1.4] and claim 29; Ex. 1005, 8:57-9:3, 9:41-52, 7:60-65.

335. Additionally, Grotmol discloses that images may be shifted in order to perform the matching step:

The image 2000 may comprise representations one or more objects, e.g., a tree 2002, a rectangle, a ball 2008, and/or other objects.

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.




FIG. 20A
Ex. 1005, Fig. 20A

Ex. 1005, 38:56-58, 40:52-63, Fig. 20A. A POSITA would understand that when a picture is pre-processed, normalized, scaled, or shifted, the image becomes a representation of a first device's surroundings.

29. Claim 33

336. I understand that independent claim 33 is a method claim that corresponds to system claim 1. My analysis above explaining how Grotmol discloses the limitations of claim 1 applies to the limitations of claim 33 as set forth in the chart below. In my opinion, Grotmol anticipates claim 33.

Claim 33	Corresponding Analysis For Claim 1
[33.1] A method implemented using a computing system that includes one or more processors, the method comprising:	To the extent the preamble is limiting, Grotmol discloses a method. Ex. 1005, claim 1 (“A method of determining a control signal for a robot, the method being performed by a special purpose computing platform having one or more processors executing instructions stored by a non-transitory computer-readable storage medium”); <i>See</i> analysis for limitations [1.1]-[1.3] above.
[33.2] 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;	<i>See</i> analysis for limitation [1.4].
[33.3] receiving or generating a first one or more instruction sets for operating the first device; and	<i>See</i> analysis for limitation [1.5].
[33.4] learning the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device.	<i>See</i> analysis for limitation [1.6].

30. Claim 34

337. Independent claim 34 is a system claim that corresponds to claim 1, except that claim 34 requires “means for receiving or generating” and “means for learning.” As I explained in the claim construction section above, a corresponding structure for “means for receiving or generating” and “means for learning” is “one or more processor circuits.” For the remainder of claim 34, my analysis above explaining how Grotmol discloses the limitations of claim 1 applies to the limitations

of claim 34 as set forth in the chart below. In my opinion, Grotmol anticipates claim 34.

Claim 33	Corresponding Analysis For Claim 1
[34.1] A system comprising:	To the extent the preamble is limiting, Grotmol discloses a system implemented as a “self-contained robotic apparatus.” Ex. 1005, 3:27-30; <i>see also id.</i> , Abstract, 10:45-48. <i>See</i> analysis for limitations [1.1]-[1.3] above.
[33.2] 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;	<i>See</i> analysis for limitation [1.4].
[33.3] means for receiving or generating a first one or more instruction sets for operating the first device; and	<i>See</i> analysis for limitation [1.5].
[33.4] means for learning the first one or more digital pictures correlated with the first one or more instruction sets for operating the first device.	<i>See</i> analysis for limitation [1.6].

C. Ground 2: Claims 10 and 11 are Obvious Based on Grotmol in View of the Knowledge of a POSITA

338. As set forth below, claims 10 and 11 of the ’583 patent are rendered obvious by Grotmol in view of the knowledge of a POSITA.

- 1. Claim 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.

Claim 10	Additional Features to Claim 4
	+ Knowledgebase + 2nd DP correlated with 2nd instructions for 1st device + 2nd DP or 2nd instructions learned in process by 2nd user

339. First, as I described above in Ground 1, Grotmol discloses a memory component configured to store training sets (e.g., “knowledgebase”), where a training set comprises at least a digital image and a motor control command (instruction set). Ex. 1005, 3:27-39, Abstract. Thus, a POSITA would understand that Grotmol discloses a memory that stores a knowledgebase including multiple correlations (training sets). *See supra* claim 9.

340. Second, as I also described above in Ground 1, Grotmol further discloses a human trainer (“user”) can operate a robot during a training process for learning correlations and that multiple learning processes can be done to train multiple behaviors for the robot, including “simple tasks” and “complex tasks.”

Apparatus and methods for hierarchical learning are disclosed herein, in accordance with one or more implementations. In supervised learning, a user may train **simple tasks** by demonstrating example tasks to the system.

For example, **a task of playing fetch may be quite complex:** the robot may be supposed to find an object, grasp it, and bring it back to base. In response to the robot not “seeing” the target, the robot may perform a random exploration behavior avoiding obstacles. In response to the robot being close to the target, the robot may perform a positioning maneuver and grasp the target. If the target is in the robot's gripper, the robot may find its way to the base avoiding obstacles and then release the object. **The whole task may be highly nonlinear and/or noisy, so a very powerful classifier may infer the right sequence of actions from just the examples and be resistant to noise. The classifier may need to determine which particular input is important for a particular task.** For example, a state of the gripper may not be relevant to target search behavior. Additional training may be needed to show the predictor that the search behavior should be performed independently of the state of the gripper.

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”.

Ex. 1005, 24:17-20, 24:31-48, 26:5-8; *see also id.*, 23:3-11, 39:13-22.

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 reason a POSITA would have understood Grotmol’s disclosure to be limited to a single user training all possible tasks, especially due to the inefficiency this limitation would impose on learning across all interoperable robots. Thus, it would be obvious based on Grotmol in view of the knowledge of a POSITA to have a second user perform an additional training resulting in an additional correlation stored in the knowledgebase.

2. **Claim 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.

Claim 11	Additional Features to Claim 1
	+ Knowledgebase + 2nd DP and 2nd instructions for 3rd device + 1st DP or 1st instructions learned on 1st device + 2nd DP or 2nd instructions learned on 3rd device

342. As I described above in Ground 1, Grotmol discloses that the training sets and training configurations from a first robot (including the instruction sets for operating the first device) can be “stored,” or “saved.” Ex. 1005, 33:12-58. The saved configurations can then be loaded onto one or more robots, providing the training sets to a second or more robots (second and third device):

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.

Id., 33:24-35.

343. Based on Grotmol's disclosure that a robot can be loaded with a trained configuration, it would have been obvious to a POSITA that Grotmol's robot can be loaded with a second, third, or more configurations trained on other robots. *Id.* For example, a POSITA would have recognized that it would have been desirable to replace a robot's training with a second, third, or more configurations to customize the robot to a particular task—e.g., a first trained configuration may be appropriate for certain tasks while a second or third configuration is more suited for another task. 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.

344. Thus, it would have been obvious based on Grotmol in view of the knowledge of a POSITA to have a knowledgebase that includes multiple configurations learned by a user partially operating different devices.

D. Ground 3: Claims 25 and 27 Are Obvious Based on Grotmol in View of Hickman

345. Grotmol in view of Hickman renders obvious claims 25 and 27.

1. Motivation to Combine Grotmol and Hickman

346. For the same reasons as I discussed in Section VIII.E.2 above, 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 learning across multiple devices with multiple users, increasing efficiency, and conserving local system resources.

- 2. Claim 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 % [sic] 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.**

Claim 25	Additional Features to Claim 4
	+ Server receives 1st DP or 1st instructions from 1st device + 2nd device receives 1st DP or 1st instructions + 2nd device performing operation from 1st instructions

347. As I described above in Grounds 1 and 2, Grotmol discloses that the training sets and training configurations from a first robot (including the instruction sets for operating the first device) can be “stored,” or “saved.” Ex. 1005, 33:12-58, *see supra* limitations [2.2], [4.4], and claim 11. The saved configurations can then

be loaded onto one or more robots, providing the training sets to a second or more robots (second and third device):

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.

Ex. 1005, 33:24-35. Thus, a POSITA would understand that Grotmol discloses that the memory of a first robot can include a first training set (first pictures correlated with first instruction set) and can have a second training set (second image correlated with second instruction set) from a second device loaded onto its memory.

348. Grotmol also discloses that training sets are used for autonomous operation of the robots by comparing acquired images with images from the stored training set to determine the most likely match and the corresponding instruction to perform. Ex. 1005, 33:24-35, Abstract.

349. Grotmol discloses, for example, situations in which the user may wish to train the robot to “find an object, grasp it, and bring it back” (i.e., fetch). *Id.*,

24:31-25:5. Grotmol describes that this “task of playing fetch may be quite complex,” and outlines numerous user training steps that would be needed in order to teach this skill to the robot. *Id.* For operational efficiency, in my opinion, implementing a system wherein learned robot behaviors could be uploaded to a shared knowledge base for use in numerous other robots around the world, such as the system described in Hickman would have been useful.

350. Hickman discloses a shared robot knowledgebase (including a server) that collects data from a robot, stores the received data, and sends instructions to a robot:

The present application discloses **shared robot knowledge bases for use with cloud computing systems**. In one embodiment, the cloud computing system **collects data from a robot about an object the robot has encountered in its environment, and stores the received data in the shared robot knowledge base**. In another embodiment, the cloud computing system **sends instructions for interacting with an object to a robot**, receives feedback from the robot based on its interaction with the object, and updates data in the shared robot knowledge base based on the feedback. In yet another embodiment, **the cloud computing system sends instructions to a robot for executing an application based on information stored in the shared robot knowledge base**.

Ex. 1007, Abstract.

351. Hickman discloses using a cloud computing system, where “[t]he ‘cloud’ may refer to a service or a group of services accessible over a network (e.g., the Internet) by clients, **server devices**, and cloud computing systems” *Id.*, 5:20-23. Moreover, “[t]he cloud computing system may be configured to (i) send and receive data and queries to and from multiple robots, (ii) send data retrieved from the shared robot knowledge base to robots, and (iii) update data in the shared robot knowledge base based on data received from robots.” *Id.*, 5:7-12; *see also* Fig. 6 (below).




FIGURE 6

Ex. 1007, Fig. 6

352. Hickman notes that shared robot knowledge base can update based on individual robot experiences. *Id.*, Abstract (“[I]nformation in the shared robot knowledge bases is updated based on robot experiences so that any particular robot may benefit from prior experiences of other robots.”). Accordingly, a POSITA

would understand that Hickman discloses a system where learned robot behaviors are uploaded to a shared knowledge base for use in other robots. In this system, a first one or more robots upload digital pictures and/or instruction sets to a server and other robots may receive digital pictures and/or instruction sets from a server.

353. Further, similar to the training described in Grotmol, Hickman discloses that robot experiences may include “coaching” from humans:

By analyzing past experiences from many different robots, the cloud computing system 304 may be configured to synthesize certain rules for robots to follow when executing the “Groceries” application.... In this manner, robot experiences **(including “coaching” from humans)** can be used to improve the way that robots execute the “Groceries” application.

Another example application may be called “Fetch.” The “Fetch” application may include robot instructions for playing fetch with a dog. The “Fetch” application may be created by a human with limited or novice-level robot programming skills. To create the “Fetch” application, **the person could instruct the robot to record a series of tasks (similar to a macro) including** (i) grasping a ball from the ground, (ii) throwing the ball, (iii) waiting for a dog to return with the ball, (iv) praising the dog, and (v) repeating tasks (i)-(iv). When the robot executes the “Fetch” application, the person may adjust or further configure certain aspects of the individual tasks via a robot user interface

on a smartphone, tablet computer, etc. to “fine tune” the application.

Ex. 1007, 23:36-61.

354. For example, Hickman describes a robot that is trained to properly grasp a cup, uploading and storing the information it learned onto the shared knowledge base, and then having another robot obtain these instructions from the memory and executing the instructions:

In the example shown in FIG. 4, the robot 404 may be performing an application that includes moving the cup 405 from one location to another location

In this example interaction, the grasping force specified in the instructions was too high, and as a result, the robot crushed the cup and caused a mess. When an accident happens, **a human can “coach” the robot on how much force to use** by, for example, manually controlling the robot's hand to grasp a new cup. The robot can capture the grasping force that it used to successfully grasp the cup under the manual control of the human, and then **send feedback 411 to the cloud processing system 402 with the modified grasping force**. The cloud processing system 402 can then update 412 task and object data in the shared robot knowledge base 403 to improve how the “grasp” task is applied this particular “cup” object.

Later, a second robot 413 may encounter an object 414 while performing the same (or a similar) application that requires grasping the same type of cup

After determining the object is a cup 414 and obtaining information about the cup 414, the cloud computing system 401 can send object data 418 associated with the cup 414 to the robot 413 in response to the identification query 415 received from the robot 413. The object data 418 may include both the identity of the cup 414 and instructions for interacting with the cup 414 based on the application or task that the robot 413 has been instructed to execute or perform

Importantly, the cloud computing system 401 updated the amount of force to use when grasping the particular type of cup that the robot 413 must now grasp (i.e., cup 414) based on the feedback 411 previously received from robot 404 regarding the same type of cup. As a result, when the robot 413 applies the “grasp” task to the type of “cup” object here (i.e., when robot 413 grasps cup 414), the robot 413 successfully grasps the cup without crushing the cup 419. Thus, **the second robot 413 has in effect learned from the experience of the earlier robot 404.**

Id., 26:29-27:21. Hickman teaches that when a robot queries the cloud processing engine, “the cloud computing system sends data associated with the object of the second query to the second robot in response to the second query received from the second robot at block 606. The data may include ... instructions for interacting with

the object of the second query.” *Id.*, 29:33-41. Thus, a POSITA exposed to Grotmol would be inclined to look to Hickman because both systems describe robots learning from human coaching and using said learning to train other robots. Further, a POSITA would be so inclined because it would have been obvious based on Grotmol’s disclosures of saving training configurations and loading them onto additional robots that a central storage repository of training configurations would be useful.

355. Accordingly, a POSITA would have understood Hickman to disclose a second device autonomously performing the one or more operations defined by the first one or more instruction sets.

356. It would have been obvious to a POSITA to combine Hickman’s shared robot knowledge base with Grotmol’s autonomous robotic system such that training configurations (e.g., first pictures correlated with first instruction set) could be uploaded to a server (i.e., shared robot knowledge base), which could be retrieved and implemented by a second robot. In my opinion, such combination would facilitate the training of complex sequences of actions such as fetch and the sharing of experiences and learning to benefit all robots connected to the knowledge base (e.g., the internet). In my opinion, a POSITA would have had a reasonable expectation of success in combining the teachings of Grotmol and Hickman, which describe complimentary systems.

357. Thus, the combination of Grotmol and Hickman could result in the training configuration learned by operating a robot described in Grotmol, uploaded to the server containing the “shared robot knowledge base” described in Hickman, and a second robot would receive the circumstance representations and instructions stored in the knowledge base in order to autonomously perform the learned behaviors. Such combination would have been obvious to a POSITA.

3. Claim 27

Claim 27	Additional Features to Claim 4
	+ Knowledgebase + 2nd instructions for operating 3rd device + Performing 2nd instructions on 3rd device

- a. [27.1] 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,

358. As I described above in Grounds 1 and 2, Grotmol discloses: (i) training a robot wherein digital pictures detected by a camera during operation of the robot are correlated with motor commands (instructions); (ii) multiple training (learning) processes can be done to train multiple behaviors; and (iii) a robot can be loaded with a first trained configuration learned during a training process involving

operation of that robot by a user, and can also be loaded with another trained configuration from a training process involving operation of a different robot. Ex. 1005, 33:24-35, Abstract, 3:27-39, 8:57-9:3, 10:23-26, 38:36-48, 39:13-22; *see supra* limitations [2.2], [4.4], and claim 11. Thus, in my opinion the knowledgebase described in Grotmol can store additional correlations, including a second correlation (second digital pictures correlated to second instruction set for operating another device). A POSITA would therefore understand that a robot can be loaded with a first trained configuration learned during a training process involving operation of that robot by a user, and the robot can also be loaded subsequently with another trained configuration learned during the training process involving operation of a different (second, third, etc.) robot by a user.

359. As noted in Hickman, the “information in the shared robot knowledge bases is updated based on robot experiences so that any particular robot may benefit from prior experiences of other robots.” Ex. 1007, Abstract. Hickman discloses that “robot experiences” may include “‘coaching’ from humans” (*Id.*, 23:36-48), similar to the training described in Grotmol. For example, similar to the “fetch” behavior described in Grotmol, Hickman describes a human coach training a robot how to properly grasp objects such as a cup, uploading and storing this information on a shared knowledge base (memory), then having another robot obtain these instructions from the memory and executing the instructions. Ex. 1007, 26:21-27:21.

360. As explained above, it would have been obvious to a POSITA to combine Hickman's shared robot knowledge base with Grotmol's autonomous robotic system such that the shared knowledgebase includes the training of multiple robots by multiple users. The combination discloses a knowledgebase that includes digital pictures correlated with instructions for operating a first robot, and digital pictures correlated with instructions for operating a second robot. Such combination would facilitate the training of complex sequences of actions such as fetch. A POSITA would have had a reasonable expectation of success in combining the teachings of Grotmol and Hickman, which describe complementary systems.

361. Thus, Grotmol in view of Hickman discloses a knowledgebase with a first digital image correlated with a first instruction set and a second digital picture correlated with a second instruction set for operating a third device.

- b. [27.2] 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;**

362. As I described above in Ground 1, Grotmol discloses a system with machine readable code executed by one or more processors that receives or generates one or more digital pictures. Ex. 1005, Abstract, 10:15-19, 38:52-58; *see supra* limitations [1.4] and [4.2]. Grotmol may be implemented using non-transitory machine readable media storing machine readable code:

Yet another aspect of the disclosure relates to a non-transitory computer-readable storage medium having instructions embodied thereon, the instructions being executable by a processor to perform a method of selecting an outcome of a plurality of outcomes.

Ex. 1005, 4:13-17. A POSITA would understand that Grotmol, when combined with Hickman, discloses processors receiving additional images throughout operation. Similarly, Hickman discloses a robot, including a processor, that receives additional pictures from a camera:

The **processor 202** may also be configured to receive and interpret **data from the various sensors 206** on the robot 200. Examples of sensors 206 that may be configured for use on the robot 200 include, for example, smoke sensors, light sensors, radio sensors, infrared sensors, microphones, gyroscopes, accelerometers, **cameras**, radars, capacitive sensors, touch sensors, or any other type of sensor now known or later developed.

Ex. 1007, 8:28-35. In my opinion, Grotmol, when combined with Hickman, discloses processors receiving additional images throughout operation.

- c. **[27.3] 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.

363. As I described above in claim 25, in my opinion in the combination of Grotmol and Hickman, the knowledgebase would include training sets learned from operating multiple robot devices (i.e., digital pictures correlated with instructions for operating a first device, second device, third device, etc.). *See supra* claim 25, Ex. 1005 33:12-58, 24:31-25:5. Further, Grotmol discloses that autonomous operation of a robot involves determining an instruction set based on at least a partial match between a new picture and a stored picture in the knowledgebase. *See supra* limitation [4.3]; Ex. 1005, Abstract, 43:5-9, 4:30-38, claim 1.

364. In my opinion, the combination of Grotmol and Hickman renders obvious that any number of robots could be autonomously operated using the shared knowledgebase, such as a fourth robot. *See, e.g.*, Ex. 1007, Abstract (“[I]nformation in the shared robot knowledge bases is updated based on robot experiences so that any particular robot may benefit from prior experiences of other robots.”). In my opinion, autonomous operation, using said combined system, would determine an instruction set for operating a first, second, or third device correlated to a digital picture. This determination would be made based on a partial match between a new picture and stored picture. This determination would cause the fourth robot to perform the operation defined by the determined instruction. Thus, a POSITA would understand Grotmol in view of Hickman to disclose limitation [27.3].

E. Ground 4: Claim 30 Is Obvious Based on Grotmol in view of Hoffman

365. As set forth below, claim 30 of the '583 patent is rendered obvious by Grotmol in view of Hoffman.

1. Motivation to Combine Grotmol and Hoffman

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.

367. Grotmol teaches using a control screen, such as on an iPhone App, that can be used by a user to send commands to the robot. Ex. 1005, 22:55-23:2; *see also id.*, 25:30-37 (“The system may be configured to learn from user commands to control the environment (e.g., robot) by paying attention to features provided by the first and the last feature extractor. For example, the user may control the robot using a gamepad (or other input device used to provide ‘user commands’) to approach the target while avoiding the obstacles (‘Behavior 1’) using input from target tracker (‘FE 1’) and from obstacle detector (‘FE 1’).”); *see also id.*, 33:18-21 (“A user may

control/training the robot with a remote device (e.g., comprising a Gamepad® controller and an iOS® application, and/or a handset device (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."); *see also id.*, 23:31-38, 25:10-25, 26:21-33. Grotmol also discloses that its control screen can include a "touchscreen interface" which can include a "touchfader combiner [that] may comprise overriding control methods" such that "the user can implement 'virtual additive' function by touching the screen just a bit to the left or to the right of the slider's current position." *Id.*, 36:62-37:7

368. Grotmol further describes using its camera to capture its surroundings in a 360 degree view field which can "provide additional information about robot's environment." *Id.*, 38:48-51. Grotmol further describes implementations whereby the user interface includes a head-mounted device such as an Oculus Rift. *Id.*, 45:37-51

369. Hoffman also teaches controlling a robot remotely using a handheld device (controller unit) that can receive and display images from the robot, including augmented images that overlay a representation of the robot on the images. *See, e.g., Ex. 1009*, 2:62-67 ("Another aspect of the disclosure provides an operator controller unit for controlling a robot. The operator controller unit includes a screen and a processor in communication with the screen. The processor is configured to

electronically receive images and augment the images by overlaying a representation of the robot on the images.”); *see also id.*, 3:2-10. “The robot representation includes user-selectable portions corresponding to movable portions of the robot,” which allows the user to remotely manipulate the robot from the controller unit. *Id.*, 2:67-3:2; *see also id.*, 3:2-24. The augmented images can include processed data set of images from the robot used to generate a map layout of the area surrounding the robot. Ex. 1009, 8:1-8 (“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).”

370. In my opinion, a POSITA would have been familiar with the teachings of both Grotmol and Hoffman and would have readily appreciated the advantages of their combination, at least because the two references are directed to the same field of semi-autonomous and autonomous operation of robots. In my opinion, the resulting combination of Grotmol’s autonomous robotic system and Hoffman’s handheld device and augmented display would have been expected and predictable. Grotmol expressly discloses controlling and training the robot via a handheld device,

such as a gamepad, and Hoffman's expansion upon this feature does not interfere with Grotmol's autonomous robotic system.

371. For at least these reasons, it would have been obvious to a POSITA to combine Grotmol and Hoffman.

2. **Claim 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.**

Claim 30	Additional Features to Claim 29
	+ Detecting location of object in 2nd DP

372. As I described above with respect to claim 29, Grotmol discloses detecting one or more objects in the first device's surrounding, which would include detecting the locations of the one or more objects. *See supra* claim 29; Ex. 1005, 38:56-65, Fig. 20A, 8:57-9:3, 9:41-52, 7:60-65.

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:

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).




FIG. 6F

Ex. 1009, Fig. 6F

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.

375. Thus, Grotmol in view of Hoffman discloses claim 30.

X. Declaration

376. I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were 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.

Dated: June 30, 2023

By: 
Jason Janet