




(86) **Date de dépôt PCT/PCT Filing Date:** 2016/11/11
 (87) **Date publication PCT/PCT Publication Date:** 2017/05/18
 (45) **Date de délivrance/Issue Date:** 2020/07/14
 (85) **Entrée phase nationale/National Entry:** 2018/05/09
 (86) **N° demande PCT/PCT Application No.:** US 2016/061698
 (87) **N° publication PCT/PCT Publication No.:** 2017/083772
 (30) **Priorité/Priority:** 2015/11/12 (US62/254,701)

(51) **Cl.Int./Int.Cl. G06N 3/063** (2006.01),
G06N 3/08 (2006.01), *G06N 3/10* (2006.01)
 (72) **Inventeurs/Inventors:**
 MNIH, VOLODYMYR, GB;
 BADIA, ADRIA PUIGDOMENECH, GB;
 GRAVES, ALEXANDER BENJAMIN, GB;
 HARLEY, TIMOTHY JAMES ALEXANDER, GB;
 SILVER, DAVID, GB;
 KAVUKCUOGLU, KORAY, GB
 (73) **Propriétaire/Owner:**
 DEEPMIND TECHNOLOGIES LIMITED, GB
 (74) **Agent:** SMART & BIGGAR LLP

(54) **Titre : APPRENTISSAGE PAR RENFORCEMENT PROFOND ASYNCHRONE**
 (54) **Title: ASYNCHRONOUS DEEP REINFORCEMENT LEARNING**

Neural Network Training System 100



(57) **Abrégé/Abstract:**

Methods, systems, and apparatus, including computer programs encoded on computer storage media, for asynchronous deep reinforcement learning. One of the systems includes one or more computers configured to implement a plurality of workers, wherein each worker is configured to operate independently of each other worker, and wherein each worker is associated with a respective actor that interacts with a respective replica of the environment during the training of the deep neural network. Aspects of the present specification have the technical effect of faster training of a neural network and/or reducing the memory requirements for the training.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(10) International Publication Number
WO 2017/083772 A1

(43) International Publication Date
18 May 2017 (18.05.2017)

- (51) International Patent Classification:
G06N 3/063 (2006.01) *G06N 3/10* (2006.01)
G06N 3/08 (2006.01)
- (21) International Application Number:
PCT/US2016/061698
- (22) International Filing Date:
11 November 2016 (11.11.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/254,701 12 November 2015 (12.11.2015) US
- (71) Applicant: **GOOGLE INC.** [US/US]; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).
- (72) Inventors: **MNIH, Volodymyr**; 7 Pancras Square, Kings Cross, London N1C 4AG (GB). **BADIA, Adria Puigdomenech**; 7 Pancras Square, Kings Cross, London N1C 4AG (GB). **GRAVES, Alexander Benjamin**; 7 Pancras Square, Kings Cross, London N1C 4AG (GB). **HARLEY, Timothy James Alexander**; 7 Pancras Square, Kings Cross, London N1C 4AG (GB). **SILVER, David**; 7 Pancras Square, Kings Cross, London N1C 4AG (GB).

KAVUKCUOGLU, Koray; 7 Pancras Square, Kings Cross, London N1C 4AG (GB).

- (74) Agents: **PORTNOV, Michael** et al.; Fish & Richardson P.C., P.O. Box 1022, 3300 Rbc Plaza, Minneapolis, Minnesota 55440-1022 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU,

[Continued on next page]

(54) Title: ASYNCHRONOUS DEEP REINFORCEMENT LEARNING

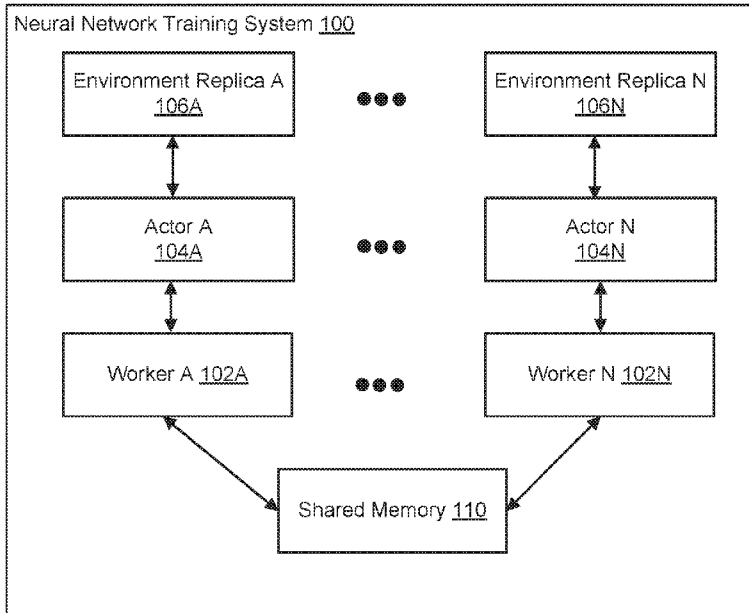


FIG. 1

(57) Abstract: Methods, systems, and apparatus, including computer programs encoded on computer storage media, for asynchronous deep reinforcement learning. One of the systems includes one or more computers configured to implement a plurality of workers, wherein each worker is configured to operate independently of each other worker, and wherein each worker is associated with a respective actor that interacts with a respective replica of the environment during the training of the deep neural network. Aspects of the present specification have the technical effect of faster training of a neural network and/or reducing the memory requirements for the training.

WO 2017/083772 A1

WO 2017/083772 A1 

LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, **Published:**
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, — *with international search report (Art. 21(3))*
GW, KM, ML, MR, NE, SN, TD, TG).

ASYNCHRONOUS DEEP REINFORCEMENT LEARNING

FIELD

This specification relates to reinforcement learning.

BACKGROUND

In a reinforcement learning system, an agent interacts with an environment by performing actions that are selected by the reinforcement learning system in response to receiving observations that characterize the current state of the environment.

5 Some reinforcement learning systems select the action to be performed by the agent in response to receiving a given observation in accordance with an output of a neural network.

Neural networks are machine learning models that employ one or more layers of nonlinear units to predict an output for a received input. Some neural networks are deep neural networks that include one or more hidden layers in addition to an output layer. The output of each hidden layer is used as input to the next layer in the network, i.e., the next hidden layer or the output layer. Each layer of the network generates an output from a received input in accordance with current values of a respective set of parameters.

10

SUMMARY

This specification describes technologies that relate to reinforcement learning. In particular, this specification describes techniques for asynchronously training a deep neural network used by a reinforcement learning system to select actions to be performed by an agent when interacting with an environment. Aspects of the present specification have the technical effect of faster training of a neural network and/or reducing the memory requirements for the training.

15

For a system of one or more computers to be configured to perform particular operations or actions means that the system has installed on it software, firmware, hardware, or a combination of them that in operation cause the system to perform the operations or actions. For one or more computer programs to be configured to perform particular operations or actions means that the one or more programs include instructions that, when executed by data processing apparatus, cause the apparatus to perform the operations or actions.

20

The subject matter described in this specification can be implemented in particular embodiments so as to realize one or more of the following advantages. By parallelizing training, a neural network used by a reinforcement learning system can be trained faster. More particularly, by parallelizing the training using multiple workers operating independently on a single machine, communication costs incurred as a result of the parallelization can be reduced. Additionally, by removing the need for the workers to store experience tuples in a replay memory or other storage as is generally required in parallelization techniques that include multiple workers on multiple machines, memory requirements for the training can be reduced and on-policy reinforcement learning methods can be employed.

In an aspect, there is provided a system for training a deep neural network used to control an agent that interacts with an environment by performing actions selected from a predetermined set of actions, the deep neural network comprising a policy neural network having a plurality of policy parameters and a baseline neural network having a plurality of baseline parameters, and the system comprising one or more computers configured to implement: a plurality of workers, wherein each worker is configured to operate independently of each other worker, wherein each worker is associated with a respective actor that interacts with a respective replica of the environment during the training of the deep neural network, and wherein each worker is further configured to repeatedly perform operations comprising: determining current values of the parameters of the deep neural network, comprising: receiving, from a shared memory accessible by each of the workers, current values of the baseline parameters of the baseline neural network; and receiving, from the shared memory, current values of the policy parameters of the policy neural network; until the environment replica interacted with by the actor associated with the worker transitions into a state that satisfies particular criteria: receiving observations characterizing states of the environment replica and, in response to each observation, selecting an action to be performed by the actor associated with the worker in response to the observation in accordance with a respective action selection policy for the worker, wherein the selecting comprises, for each observation: processing the observation using the policy neural network and in accordance with the current values of the parameters of the policy neural network to generate a respective score for each action in the predetermined set of actions in accordance with the current values of the parameters of the policy neural network, and selecting the current action to be performed by the actor using

the respective scores; for each observation received before the particular criteria are satisfied: generating, based on the current values of the parameters of the baseline network, a baseline score representing an estimated reward received by the agent starting from a state characterized by the observation; identifying an actual reward resulting from the actor
5 performing the selected action in response to the observation; determining an actual long-term reward from at least the actual reward and a baseline score for the observation characterizing the state that satisfies the particular criteria; and performing an iteration of a reinforcement learning technique to determine respective current gradients for the baseline and policy
10 accumulated gradients for the baseline and policy networks based on respective current gradient updates to determine respective updated accumulated gradients; determining whether criteria for updating the current values of the parameters of the deep neural network have been satisfied; and when the criteria for updating the current values of the parameters of the deep neural network have been satisfied: updating the parameters of the deep neural network,
15 wherein the updating comprises: determining updated values of the parameters using the updated accumulated gradients, and storing the updated values in the shared memory accessible to all of the plurality of workers.

In another aspect, there is provided one or more non-transitory storage media storing instructions that when executed by one or more computers cause the one or more computers
20 to implement a system for training a deep neural network used to control an agent that interacts with an environment by performing actions selected from a predetermined set of actions, the deep neural network comprising a policy neural network having a plurality of policy parameters and a baseline neural network having a plurality of baseline parameters, and the system comprising: a plurality of workers, wherein each worker is configured to
25 operate independently of each other worker, wherein each worker is associated with a respective actor that interacts with a respective replica of the environment during the training of the deep neural network, and wherein each worker is further configured to repeatedly perform operations comprising: determining current values of the parameters of the deep neural network, comprising: receiving, from a shared memory accessible by each of the workers,
30 current values of the baseline parameters of the baseline neural network; and receiving, from the shared memory, current values of the policy parameters of the policy neural network; until

the environment replica interacted with by the actor associated with the worker transitions into a state that satisfies particular criteria: receiving observations characterizing states of the environment replica and, in response to each observation, selecting an action to be performed by the actor associated with the worker in response to the observation in accordance with a
5 respective action selection policy for the worker, wherein the selecting comprises, for each observation: processing the observation using the policy neural network and in accordance with the current values of the parameters of the policy neural network to generate a respective score for each action in the predetermined set of actions in accordance with the current values of the parameters of the policy neural network, and selecting the current action to be performed by
10 the actor using the respective scores; for each observation received before the particular criteria are satisfied: generating, based on the current values of the parameters of the baseline network, a baseline score representing an estimated reward received by the agent starting from a state characterized by the observation; identifying an actual reward resulting from the actor performing the current action in response to the observation; determining an actual long-term
15 reward from at least the actual reward and a baseline score for the observation characterizing the state that satisfies the particular criteria; and performing an iteration of a reinforcement learning technique to determine respective current gradients for the baseline and policy networks based on the baseline score and the actual long-term reward; updating respective accumulated gradients for the baseline and policy networks based on respective current
20 gradient updates to determine respective updated accumulated gradients; determining whether criteria for updating the current values of the parameters of the deep neural network have been satisfied; and when the criteria for updating the current values of the parameters of the deep neural network have been satisfied: updating the parameters of the deep neural network, wherein the updating comprises: determining updated values of the parameters using the
25 updated accumulated gradients, and storing the updated values in the shared memory accessible to all of the plurality of workers.

In another aspect, there is provided a method of training a deep neural network used to control an agent that interacts with an environment by performing actions selected from a predetermined set of actions, the deep neural network comprising a policy neural network
30 having a plurality of policy parameters and a baseline neural network having a plurality of baseline parameters, the method comprising: determining, by a first worker of a plurality of

workers, current values of the parameters of the deep neural network, wherein each worker is configured to operate independently of each other worker, and wherein each worker is associated with a respective actor that interacts with a respective replica of the environment during the training of the deep neural network, the determining comprising: receiving, from a shared memory accessible by each of the workers, current values of the baseline parameters of the baseline neural network; and receiving, from the shared memory, current values of the policy parameters of the policy neural network; until the environment replica interacted with by the actor associated with the worker transitions into a state that satisfies particular criteria: receiving, by the first worker, observations characterizing states of the environment replica and, in response to each observation, selecting, by the first worker an action to be performed by the actor associated with the worker in response to the observation in accordance with a respective action selection policy for the worker, wherein the selecting comprises, for each observation: processing the observation using the policy neural network and in accordance with the current values of the parameters of the policy neural network to generate a respective score for each action in the predetermined set of actions in accordance with the current values of the parameters of the policy neural network, and selecting the current action to be performed by the actor using the respective scores; for each observation received before the particular criteria are satisfied: generating, by the first worker and based on the current values of the parameters of the baseline network, a baseline score representing an estimated reward received by the agent starting from a state characterized by the observation; identifying, by the first worker, an actual reward resulting from the actor performing the current action when the environment replica is in the current state; determining an actual long-term reward from at least the actual reward and a baseline score for the observation characterizing the state that satisfies the particular criteria; and performing, by the first worker, an iteration of a reinforcement learning technique to determine respective current gradients for the baseline and policy networks based on the baseline score and the actual long-term reward; updating, by the first worker, respective accumulated gradients for the baseline and policy networks based on respective current gradient updates to determine respective updated accumulated gradients; determining, by the first worker, whether criteria for updating the current values of the parameters of the deep neural network have been satisfied; and when the criteria for updating the current values of the parameters of the deep neural network have been satisfied: updating,

by the first worker, the parameters of the deep neural network, wherein the updating comprises: determining updated values of the parameters using the updated accumulated gradients, and storing, by the first worker, the updated values in the shared memory accessible to all of the plurality of workers.

5 The details of one or more embodiments of the subject matter of this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the specification. For example, the specification provides methods, systems, and apparatus, including computer programs encoded on computer storage media, for carrying out any of the techniques disclosed herein. In some
10 embodiments, a computer program comprises instructions that when executed by computing apparatus causes it to perform any of the techniques disclosed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows an example neural network training system.

FIG. 2 is a flow diagram of an example process for training a deep neural network.

FIG. 3 is a flow diagram of an example process for performing an iteration of a Q-
15 learning technique.

FIG. 4 is a flow diagram of an example process for performing an iteration of a SARSA technique.

FIG. 5 is a flow diagram of an example process for training a policy neural network.

Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

20 This specification generally describes techniques for asynchronous deep reinforcement learning. In particular, this specification describes how multiple workers can asynchronously train a deep neural network used by a reinforcement learning system to select actions to be performed by an agent interacting with an environment.

In some implementations, the environment is a simulated environment and the agent is implemented as one or more computer programs interacting with the simulated environment. For example, the simulated environment may be a video game and the agent may be a simulated user playing the video game. As another example, the
5 simulated environment may be a motion simulation environment, e.g., a driving simulation or a flight simulation, and the agent is a simulated vehicle navigating through the motion simulation. In these implementations, the actions may be control inputs to control the simulated user or simulated vehicle.

In some other implementations, the environment is a real-world environment and
10 the agent is a mechanical agent interacting with the real-world environment. For example, the agent may be a robot interacting with the environment to accomplish a specific task. As another example, the agent may be an autonomous or semi-autonomous vehicle navigating through the environment. In these implementations, the actions may be control inputs to control the robot or the autonomous vehicle.

Generally, the reinforcement learning system receives observations characterizing
15 the current state of the environment and selects actions to be performed by the agent using the observations. In response to the agent performing a selected action, the environment transitions into a new state and the reinforcement learning system receives a reward. The reward is a numeric value that is a function of the state of the environment.
20 While interacting with the environment, the reinforcement learning system attempts to maximize the long-term reward received by the agent in response to the actions performed by the agent.

In particular, in response to a given observation, the reinforcement learning system selects the action to be performed by the agent in accordance with an output
25 generated by a deep neural network.

In order to allow the reinforcement learning system to effectively select actions to be performed by the agent, a neural network training system trains the deep neural network to adjust the values of the parameters of the deep neural network from initial values of the parameters.

In particular, the neural network training system includes multiple workers that
30 asynchronously train the deep neural network by selecting actions to be performed by actors interacting with replicas of the environment.

Each actor is an instance of the agent that operates independently from each other instance. For example, when the agent is a mechanical agent, each actor is also a separate

mechanical agent that is the same as or is approximately the same as the agent, e.g., has the same or approximately the same properties and can perform the same actions as the agent, and operates independently from each other actor.

FIG. 1 shows an example neural network training system 100. The neural
5 network training system 100 is an example of a system implemented as computer programs on one or more computers in one or more locations in which the systems, components, and techniques described below are implemented.

The neural network training system 100 includes multiple workers 102A-N and a shared memory 110. Each of the workers 102A-102N is associated with a respective
10 actor 104A-N that interacts with a corresponding environment replica 106A-106N. For example, the worker 102A is associated with the actor 104A that interacts with the environment replica 106A.

In some implementations, each of the workers 110A-N is implemented on the same computer but executes in a separate thread, process or other hardware or software
15 within the computer capable of independently performing the computation for the worker. In these implementations, the shared memory 110 is memory of the computer that is accessible to each of the workers 102A-102N. In implementations where the environment is a virtualized environment, the actor and environment replica associated with a given worker also execute in the same thread or other hardware or software as the
20 given worker.

Each of the workers 102A-102N maintains an instance of the deep neural network that is being trained and accesses the shared memory 110 to obtain current values of the parameters of the deep neural network that are stored in the shared memory 110. Each worker then uses the worker's instance of the deep neural network to select actions to be
25 performed by the actor associated with the worker in accordance with the current values of the parameters of the neural network obtained from the shared memory 110.

In particular, during the training, each worker receives observations characterizing the state of the environment replica associated with the worker and, in response to each observation, selects an action to be performed by the actor associated with the
30 environment replica using the deep neural network. The worker then receives a next observation characterizing the next state of the environment replica, i.e., the state that the environment replica transitions into as a result of the actor performing the selected action, and a reward resulting from the actor performing the selected action. The worker uses the reward and the next observation to determine a gradient and periodically uses

gradients determined by the worker to update the parameter values that are stored in the shared memory 110. Determining gradients and updating parameter values stored in a shared memory is described in more detail below with reference to FIG. 2.

By each of the workers 102A-102N repeatedly and asynchronously updating the parameter values stored in the shared memory 110, the workers 102A-102N collectively train the deep neural network to determine trained values of the parameters of the deep neural network.

FIG. 2 is a flow diagram of an example process 200 for training a deep neural network. For convenience, the process 200 will be described as being performed by a worker, e.g., one of the workers 120A-N of FIG. 1.

The worker determines current values of the parameters of the deep neural network (step 202).

In some implementations, the worker obtains the current values from a memory shared between all of the workers, e.g., the shared memory 110 of FIG. 1.

In some other implementations, the worker only obtains the current values from the memory after updating the shared memory as described below and writes them to a local memory. In these implementations, the worker determines the current values to be the values stored in the local memory.

The worker receives a current observation characterizing a current state of the environment replica being interacted with by the actor associated with the worker (step 204).

The worker selects a current action to be performed by the actor when the environment is in the current state using the deep neural network (step 206).

In particular, in some implementations, the deep neural network is a Q network that is configured to receive as an input an observation and an action and to process the input to generate a neural network output in accordance with the current values of the parameters of the Q network. The neural network output, which will be referred to in this specification as a Q value, is an estimate of a long-term reward that will result from an agent or an actor performing the input action when the environment is in a state characterized by the input observation.

In these implementations, for each action in the predetermined set of actions, the worker processes the action and the current observation using the Q network to generate a respective current Q value for each action.

The worker then selects an action from the predetermined set of actions as the current action to be performed by the actor using the current Q values and in accordance with an action selection policy for the worker.

5 Generally, the action selection policy for a given worker is different from the action selection policy for each other worker. For example, in some implementations, the action selection policy for each worker is an ϵ -greedy policy in which the worker selects an action randomly from the predetermined set of actions with probability ϵ and selects the action having the highest Q value with probability $1 - \epsilon$, but with the value of ϵ being different for each worker. For example, each worker can periodically sample a value for
10 ϵ for the worker from a distribution of possible values of ϵ for the worker to determine the ϵ to be used in the action selection policy. As another example, the worker could select the action by applying a softmax nonlinearity to the Q values and sampling an action from the resulting distribution. Each worker could then use a different temperature parameter in the softmax to ensure that each worker had a different action selection
15 policy.

The worker receives a next observation and an actual reward (step 208). The next observation characterizes a next state of the environment replica, i.e., the state that the environment replica transitioned into as a result of the actor associated with the worker performing the current action. The actual reward is a numeric value that is received from
20 the environment replica as a result of the transition from the current state into the next state.

The worker performs an iteration of a reinforcement learning procedure to determine a current gradient using the actual reward and the next observation (step 210). The gradient is the gradient of a loss function being optimized in order to train the deep
25 neural network.

The loss function and the manner in which the worker determines the gradient are dependent on the reinforcement learning procedure being used to train the deep neural network, which is in turn dependent on how the deep neural network is configured.

30 For example, when the deep neural network is a Q network, the workers may use a Q-learning technique or a SARSA technique. Performing an iteration of a Q-learning technique is described below with reference to FIG. 3. Performing an iteration of a SARSA technique is described below with reference to FIG. 4.

The worker updates an accumulated gradient with the current gradient to determine an updated accumulated gradient (step 212). That is, the worker maintains a

gradient in local memory and, each time a new gradient is computed, updates the maintained gradient with the new gradient, e.g., by adding the new gradient to the accumulated gradient.

The worker determines whether criteria for updating the current values of the parameters of the neural network are satisfied (step 214). For example, the criteria for updating the current values may specify that the worker update the current values after the worker has performed a specified number of reinforcement learning iterations since the most recent time the worker updated the parameter values. As another example, the criteria may specify that the worker update the current values when the total number of reinforcement learning iterations performed by all of the multiple workers since the most recent time the workers updated the parameter values exceeds a specified threshold value.

If the criteria are satisfied, the worker determines updated parameter values using the accumulated gradients (step 216) and writes the updated parameter values to the shared memory (step 218). In particular, the worker determines parameter value updates from the accumulated gradients, accesses the current values in the shared memory, and then updates those values using the parameter value updates. In some examples, the worker multiplies the accumulated gradients by a learning rate parameter and then sums the current parameter values and the parameter value updates to determine the updated parameter values. In some other examples, the worker uses an asynchronous variant of the RMSProp optimization procedure to update the parameter values.

After the worker writes the updated parameter values to the shared memory, the worker clears the accumulated gradient, i.e., so that no gradient is maintained by the worker when the next iteration of the process 200 is initiated.

If the criteria are not satisfied, the worker refrains from writing anything to the shared memory (step 220), i.e., the worker proceeds to perform another iteration of the process 200 without updating the current parameter values.

Each of multiple workers can asynchronously repeatedly perform the process 200 to collectively determine trained values of the parameters of the deep neural network. Once the trained values have been determined, a reinforcement learning system can use the deep neural network to select values to be performed by an agent interacting with the environment, i.e., by processing neural network inputs using the deep neural network in accordance with the trained values of the neural network parameters. In some other cases, the reinforcement learning system uses the deep neural network to select values to be performed by the agent while the workers continue to perform the process 200, i.e., by

periodically obtaining the values stored in the shared memory and using those values to select the actions performed by the agent.

In some cases, once the trained values have been determined, they are stored and then transmitted over a network to another system for use in instantiating a trained
 5 version of the deep neural network, i.e., to allow an agent to effectively select actions in interacting with the environment.

FIG. 3 is a flow diagram of an example process 300 for performing an iteration of a Q-learning technique. For convenience, the process 300 will be described as being performed by a worker, e.g., one of the workers 102A-N of FIG. 1.

10 The worker obtains the current observation, the current action, the actual reward, and the next observation (step 302), i.e., as described above with reference to FIG. 2.

The worker determines a maximum output of a target network when the environment replica is in a state characterized by the next observation (step 304). The target network is the same neural network as the Q network, i.e., has the same neural
 15 network architecture / neural network form, but with possibly different parameter values. In particular, each worker periodically synchronizes the values of the parameters of the target network with the current values of the parameters of the Q network stored in the shared memory. Generally, however, the worker synchronizes the values of the
 20 parameters of the target network less frequently than the worker updates the parameter values stored in the shared memory, so that the target network and the Q network will often have different parameter values.

The worker determines the maximum output of the target network by, for each action in the predetermined set of actions, processing the action in combination with the next observation using the target network to determine a target network output for the
 25 action. The system then selects the maximum of the target network outputs for the actions in the predetermined set of actions as the maximum output of the target network.

The worker determines an error from the maximum output of the target network, the actual reward, and the current Q value previously determined for the current action when the environment was in the current state, i.e., the current Q value used to select the
 30 current action as the action to be performed when the environment was in the current state (step 306). In some implementations, the error E satisfies:

$$E = r + \gamma \max(Q^{\text{target}}) - Q^{\text{current}},$$

where r is the actual reward, γ is a predetermined discount factor, $\max(Q^{\text{target}})$ is the maximum output of the target network, and Q^{current} is the current Q value previously determined for the current action when the environment was in the current state.

The worker determines the current gradient using the determined error (step 308).
 5 That is, the worker updates the values using backpropagation, with the backpropagated error being the determined error for the selected action and zero for all other actions.

FIG. 4 is a flow diagram of an example process 400 for performing an iteration of a SARSA technique. For convenience, the process 400 will be described as being performed by a worker in a neural network training system, e.g., one of the workers
 10 102A-N in the neural network training system 100 of FIG. 1.

The worker receives the current observation, the current action, the actual reward, and the next observation (step 402).

The worker selects a next action to be performed by the actor when the environment replica is in the next state (step 404). The worker selects the next action as
 15 described above with reference to step 206. That is, for each action in the predetermined set of actions, the worker processes the action and the next observation using the Q network to generate a respective next Q value for each action. The worker then selects an action from the predetermined set of actions as the next action to be performed by the actor using the next Q values and in accordance with the action selection policy for the
 20 worker.

The worker determines a target network output for the next action (step 406). That is, the worker processes the next action and the next observation using the target network to determine the target network output for the next action.

The worker determines an error from the target network output for the next action,
 25 the actual reward, and the current Q value previously determined for the current action (step 406). In some implementations, the error E satisfies:

$$E = r + \gamma Q^{\text{target}} - Q^{\text{current}},$$

where r is the actual reward, γ is a predetermined discount factor, Q^{target} is the target network output for the next action, and Q^{current} is the current Q value previously
 30 determined for the current action when the environment was in the current state.

The worker determines the current gradient using the determined error (step 408).

When performing either the process 300 or the process 400, the worker also determines whether or not criteria for synchronizing the target network parameters with the parameters in the shared memory are satisfied. For example, the criteria can be the

same as the criteria described above with reference to step 216, but with the specified threshold value or the specified number being greater than the value or number used for the determination of whether to update the parameters stored in the shared memory.

Thus, the worker synchronizes the values of the parameters of the target network less frequently than the worker updates the parameter values stored in the shared memory, so that the target network and the Q network will often have different parameter values.

The above description has described implementations where the deep neural network is a Q network. In some other implementations, however, the deep neural network is a policy neural network that receives an input observation and outputs a respective score for each action in the predetermined set of actions, where the scores each represent a likelihood that the corresponding action is the action that should be performed to maximize the long- reward that will result from the environment being in the state characterized by the input observation.

In these implementations, to select an action in response to a given observation, the reinforcement learning system processes the given observation using the policy neural network and selects the action to be performed by the agent by sampling from the actions in accordance with the scores for the actions.

To ensure that each worker has a different action selection policy, each worker can use a different random seed when sampling from the actions or, alternatively, each worker can have a different temperature for the output layer of the policy neural network that generates the scores for the actions.

FIG. 5 is a flow diagram of an example process 500 for training a policy neural network. For convenience, the process 500 will be described as being performed by a worker in a neural network training system, e.g., one of the workers 102A-N in the neural network training system 100 of FIG. 1.

The worker determines current values of the parameters of the policy neural network (step 502). In particular, the worker obtains the current values from a memory shared between all of the workers, e.g., the shared memory 110 of FIG. 1.

The worker also obtains current values of the parameters of a baseline neural network from the shared memory. The baseline neural network is a neural network that is configured to receive an input observation and to generate a baseline score that represents an estimated long-term reward received by the agent starting from the state characterized by the input observation. In some implementations, the baseline and policy networks share some of their parameters.

The worker receives observations characterizing the state of the environment replica and selects actions to be performed by the actor in accordance with the current values of the parameters of the policy neural network until the environment replica transitions into a state that satisfies particular criteria (step 504). That is, for each
5 observation, the processes the observation using the policy neural network in accordance with the current values and selects the action to be performed by the agent by sampling from the actions in accordance with the scores for the actions. For example, the particular criteria may be satisfied after a predetermined number t_{\max} of observations have been received or until a predetermined terminal state has been reached.

10 In particular, for each received observation characterizing the state of the environment replica, the worker processes the observation using the policy neural network to generate a respective score for each action in the predetermined set of actions in accordance with the current values of the parameters of the policy neural network and selects the action to be performed by the actor by sampling from the distribution.

15 For each received observation, the worker also processes the action and the observation using the baseline neural network to generate a baseline score for the observation in accordance with the current values of the parameters of the baseline neural network.

20 In some implementations, one or both of the baseline neural network and the policy neural network may be a recurrent neural network.

The worker determines an actual long-term reward for each received observation (step 506).

That is, for the last observation received, the system sets the actual long-term reward to, e.g., the baseline score predicted by the baseline neural network.

25 For a given observation other than the last observation, the system sets the actual long-term reward for the given observation as the reward received when the environment replica transitioned into the state characterized by the last observation plus a discounted sum of the actual long-term reward for the observations received after the actor performed the action selected in response to receiving the given observation.

30 The worker determines, for each observation, a respective gradient update for the policy neural network and a respective gradient update for the baseline neural network (step 510).

In particular, the gradient update for the policy neural network for the observation s_t satisfies:

$$\nabla_{\theta'} \log \pi(a_t | s_t; \theta') (R_t - b_t) ,$$

where $\pi(a_t | s_t; \theta')$ is the score generated for the action performed in response to the observation s_t , R_t is the actual long-term reward determined for the observation s_t , and b_t is the baseline score for the observation s_t .

5 The gradient update for the baseline neural network for the observation s_t satisfies:

$$\partial (R_t - b(s_t; \theta'_b))^2 / \partial \theta'_b ,$$

where $b(s_t; \theta'_b)$ is the baseline score for the observation s_t .

The worker updates an accumulated gradient for the policy neural network with the gradient updates for the policy neural network to determine an updated accumulated
10 gradient for the policy neural network and updates an accumulated gradient for the baseline neural network with the gradient updates for the baseline neural network to determine an updated accumulated gradient for the baseline neural network (step 512) e.g., by adding the new gradient for each neural network to the accumulated gradient for the neural network.

15 The worker determines whether criteria for updating the current values of the parameters of the baseline neural network and the policy neural network are satisfied (step 514).

For example, the criteria for updating the current values may specify that the worker update the current values after the worker has performed a specified number
20 iterations of step 504 since the most recent time the worker updated the parameter values.

As another example, the criteria may specify that the worker update the current values when the total number of iterations of step 504 performed by all of the multiple workers since the most recent time the workers updated the parameter values exceeds a specified threshold value.

25 If the criteria are satisfied, the worker determines updated parameter values for the baseline neural network and the policy neural network using the corresponding accumulated gradients (step 516) and writes the updated parameter values to the shared memory (step 518), e.g., as described above with reference to steps 216 and 218.

30 After the worker writes the updated parameter values to the shared memory, the worker clears the accumulated gradient, i.e., so that no gradient is maintained by the worker when the next iteration of the process 500 is initiated.

If the criteria are not satisfied, the worker refrains from writing anything to the shared memory (step 520), i.e., the worker proceeds to perform another iteration of the process 500 without updating the current parameter values.

Embodiments of the subject matter and the functional operations described in this specification can be implemented in digital electronic circuitry, in tangibly-embodied computer software or firmware, in computer hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions encoded on a tangible non-transitory program carrier for execution by, or to control the operation of, data processing apparatus. Alternatively or in addition, the program instructions can be encoded on an artificially-generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal, that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus. The computer storage medium can be a machine-readable storage device, a machine-readable storage substrate, a random or serial access memory device, or a combination of one or more of them.

The term “data processing apparatus” refers to data processing hardware and encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can also be or further include special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit). The apparatus can optionally include, in addition to hardware, code that creates an execution environment for computer programs, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

A computer program (which may also be referred to or described as a program, software, a software application, a module, a software module, a script, or code) can be written in any form of programming language, including compiled or interpreted languages, or declarative or procedural languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data, e.g., one or more scripts stored in a markup language

document, in a single file dedicated to the program in question, or in multiple coordinated files, e.g., files that store one or more modules, sub-programs, or portions of code. A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and
5 interconnected by a communication network.

The processes and logic flows described in this specification can be performed by one or more programmable computers executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special
10 purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit).

Computers suitable for the execution of a computer program include, by way of example, can be based on general or special purpose microprocessors or both, or any other kind of central processing unit. Generally, a central processing unit will receive
15 instructions and data from a read-only memory or a random access memory or both. The essential elements of a computer are a central processing unit for performing or executing instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic,
20 magneto-optical disks, or optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device, e.g., a universal serial bus (USB) flash drive, to name just a few.

Computer-readable media suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash
25 memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.
30

To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball,

by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user's client device in response to requests received from the web browser.

Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back-end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front-end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back-end, middleware, or front-end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network ("LAN") and a wide area network ("WAN"), e.g., the Internet.

The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any invention or of what may be described, but rather as descriptions of features that may be specific to particular embodiments of particular inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially described as such, one or more features from a claimed combination can in some cases be excised from the combination, and the described combination may be directed to a subcombination or variation of a subcombination.

Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system modules and components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

Particular embodiments of the subject matter have been described. Other embodiments are within the scope of the subject matter presented herein. For example, the actions described in the present disclosure can be performed in a different order and still achieve desirable results. As one example, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain implementations, multitasking and parallel processing may be advantageous.

CLAIMS:

1. A system for training a deep neural network used to control an agent that interacts with an environment by performing actions selected from a predetermined set of actions, the deep neural network comprising a policy neural network having a plurality of policy parameters and a baseline neural network having a plurality of baseline parameters, and the system comprising one or more computers configured to implement:

a plurality of workers,

wherein each worker is configured to operate independently of each other worker,

wherein each worker is associated with a respective actor that interacts with a respective replica of the environment during the training of the deep neural network, and

wherein each worker is further configured to repeatedly perform operations comprising:

determining current values of the parameters of the deep neural network, comprising:

receiving, from a shared memory accessible by each of the workers, current values of the baseline parameters of the baseline neural network; and

receiving, from the shared memory, current values of the policy parameters of the policy neural network;

until the environment replica interacted with by the actor associated with the worker transitions into a state that satisfies particular criteria:

receiving observations characterizing states of the environment replica and, in response to each observation, selecting an action to be performed by the actor associated with the worker in response to the observation in accordance with a respective action selection policy for the worker, wherein the selecting comprises, for each observation:

processing the observation using the policy neural network and in accordance with the current values of the

parameters of the policy neural network to generate a respective score for each action in the predetermined set of actions in accordance with the current values of the parameters of the policy neural network, and

selecting the current action to be performed by the actor using the respective scores;

for each observation received before the particular criteria are satisfied:

generating, based on the current values of the parameters of the baseline network, a baseline score representing an estimated reward received by the agent starting from a state characterized by the observation;

identifying an actual reward resulting from the actor performing the selected action in response to the observation;

determining an actual long-term reward from at least the actual reward and a baseline score for the observation characterizing the state that satisfies the particular criteria; and

performing an iteration of a reinforcement learning technique to determine respective current gradients for the baseline and policy networks based on the baseline score and the actual long-term reward;

updating respective accumulated gradients for the baseline and policy networks based on respective current gradient updates to determine respective updated accumulated gradients;

determining whether criteria for updating the current values of the parameters of the deep neural network have been satisfied; and

when the criteria for updating the current values of the parameters of the deep neural network have been satisfied:

updating the parameters of the deep neural network, wherein the updating comprises:

determining updated values of the parameters using the updated accumulated gradients, and
storing the updated values in the shared memory accessible to all of the plurality of workers.

2. The system of claim 1, wherein each worker executes independently of each other worker on the same computer.
3. The system of claim 1, wherein the respective action selection policy for each worker is different from the respective action selection policy for each other worker.
4. The system of claim 1, the operations further comprising, when the criteria for updating the current values of the parameters of the deep neural network have been satisfied, clearing the updated accumulated gradient.
5. The system of claim 1, wherein determining whether the criteria for updating the current values of the parameters of the deep neural network have been satisfied comprises determining that a specified number of iterations of the reinforcement learning technique have been performed since a preceding update of the values of the parameters of the deep neural network.
6. The system of claim 1, wherein determining the actual long-term reward comprises:
for the last observation received before the criteria are satisfied, setting the actual long term reward based on the actual reward for the observation and the baseline score for the observation characterizing the state that satisfies the particular criteria; and
for each other observation received before the criteria are satisfied, setting the actual long-term reward for the observation as the actual reward plus a discounted sum of the actual long-term rewards for observations received after the observation.

7. The system of claim 1, wherein the gradient update for the policy neural network is determined by:

$$\nabla_{\theta'} \log \pi(a_t | s_t; \theta') (R_t - b_t)$$

Wherein $\pi(a_t | s_t; \theta')$ represents the score for the action performed in response to the observation s_t determined in accordance with the current values of the parameters of the policy neural network, R_t is the actual reward and b_t is the baseline score.

8. The system of claim 7, wherein the gradient update for the baseline neural network is determined by:

$$\partial (R_t - b(s_t; \theta'_b))^2 / \partial \theta'_b$$

wherein $b(s_t; \theta'_b)$ is the baseline score for the observation s_t .

9. One or more non-transitory storage media storing instructions that when executed by one or more computers cause the one or more computers to implement a system for training a deep neural network used to control an agent that interacts with an environment by performing actions selected from a predetermined set of actions, the deep neural network comprising a policy neural network having a plurality of policy parameters and a baseline neural network having a plurality of baseline parameters, and the system comprising:

a plurality of workers,

wherein each worker is configured to operate independently of each other worker,

wherein each worker is associated with a respective actor that interacts with a respective replica of the environment during the training of the deep neural network, and

wherein each worker is further configured to repeatedly perform operations comprising:

determining current values of the parameters of the deep neural network, comprising:

receiving, from a shared memory accessible by each of the workers, current values of the baseline parameters of the baseline neural network; and

receiving, from the shared memory, current values of the policy parameters of the policy neural network;

until the environment replica interacted with by the actor associated with the worker transitions into a state that satisfies particular criteria:

receiving observations characterizing states of the environment replica and, in response to each observation, selecting an action to be performed by the actor associated with the worker in response to the observation in accordance with a respective action selection policy for the worker, wherein the selecting comprises, for each observation:

processing the observation using the policy neural network and in accordance with the current values of the parameters of the policy neural network to generate a respective score for each action in the predetermined set of actions in accordance with the current values of the parameters of the policy neural network, and

selecting the current action to be performed by the actor using the respective scores;

for each observation received before the particular criteria are satisfied:

generating, based on the current values of the parameters of the baseline network, a baseline score representing an estimated reward received by the agent starting from a state characterized by the observation;

identifying an actual reward resulting from the actor performing the current action in response to the observation;

determining an actual long-term reward from at least the actual reward and a baseline score for the observation characterizing the state that satisfies the particular criteria; and performing an iteration of a reinforcement learning technique to determine respective current gradients for the baseline and policy networks based on the baseline score and the actual long-term reward;

updating respective accumulated gradients for the baseline and policy networks based on respective current gradient updates to determine respective updated accumulated gradients;

determining whether criteria for updating the current values of the parameters of the deep neural network have been satisfied; and when the criteria for updating the current values of the parameters of the deep neural network have been satisfied:

updating the parameters of the deep neural network, wherein the updating comprises:

determining updated values of the parameters using the updated accumulated gradients, and storing the updated values in the shared memory accessible to all of the plurality of workers.

10. The computer-storage media of claim 9, wherein each worker executes independently of each other worker on the same computer.

11. The computer-storage media of claim 9, wherein the respective action selection policy for each worker is different from the respective action selection policy for each other worker.

12. The computer-storage media of claim 9, the operations further comprising, when the criteria for updating the current values of the parameters of the deep neural network have been satisfied, clearing the updated accumulated gradient.

13. The computer-storage media of claim 9, wherein determining whether the criteria for updating the current values of the parameters of the deep neural network have been satisfied comprises determining that a specified number of iterations of the reinforcement learning technique have been performed since a preceding update of the values of the parameters of the deep neural network.

14. The non-transitory storage media of claim 9, wherein determining the actual long-term reward comprises:

for the last observation received before the criteria are satisfied, setting the actual long-term reward based on the actual reward for the observation and the baseline score for the observation characterizing the state that satisfies the particular criteria; and

for each other observation received before the criteria are satisfied, setting the actual long-term reward for the observation as the actual reward plus a discounted sum of the actual long-term rewards for observations received after the observation.

15. The non-transitory storage media of claim 9, wherein the gradient update for the policy neural network is determined by:

$$\nabla_{\theta'} \log \pi(a_t | s_t; \theta') (R_t - b_t)$$

Wherein $\pi(a_t | s_t; \theta')$ represents the score for the action performed in response to the observation s_t determined in accordance with the current values of the parameters of the policy neural network, R_t is the actual reward and b_t is the baseline score.

16. The non-transitory storage media of claim 15, wherein the gradient update for the baseline neural network is determined by:

$$\partial (R_t - b(s_t; \theta'_b))^2 / \partial \theta'_b$$

wherein $b(s_t; \theta'_b)$ is the baseline score for the observation s_t .

17. A method of training a deep neural network used to control an agent that interacts with an environment by performing actions selected from a predetermined set of actions, the deep neural network comprising a policy neural network having a plurality of policy parameters and a baseline neural network having a plurality of baseline parameters, the method comprising:

determining, by a first worker of a plurality of workers, current values of the parameters of the deep neural network, wherein each worker is configured to operate independently of each other worker, and wherein each worker is associated with a respective actor that interacts with a respective replica of the environment during the training of the deep neural network, the determining comprising:

receiving, from a shared memory accessible by each of the workers, current values of the baseline parameters of the baseline neural network; and

receiving, from the shared memory, current values of the policy parameters of the policy neural network;

until the environment replica interacted with by the actor associated with the worker transitions into a state that satisfies particular criteria:

receiving, by the first worker, observations characterizing states of the environment replica and, in response to each observation, selecting, by the first worker an action to be performed by the actor associated with the worker in response to the observation in accordance with a respective action selection policy for the worker, wherein the selecting comprises, for each observation:

processing the observation using the policy neural network and in accordance with the current values of the parameters of the policy neural network to generate a respective score for each action in the predetermined set of actions in accordance with the current values of the parameters of the policy neural network, and

selecting the current action to be performed by the actor using the respective scores;

for each observation received before the particular criteria are satisfied:

generating, by the first worker and based on the current values of the parameters of the baseline network, a baseline score representing an

estimated reward received by the agent starting from a state characterized by the observation;

identifying, by the first worker, an actual reward resulting from the actor performing the current action when the environment replica is in the current state;

determining an actual long-term reward from at least the actual reward and a baseline score for the observation characterizing the state that satisfies the particular criteria; and

performing, by the first worker, an iteration of a reinforcement learning technique to determine respective current gradients for the baseline and policy networks based on the baseline score and the actual long-term reward;

updating, by the first worker, respective accumulated gradients for the baseline and policy networks based on respective current gradient updates to determine respective updated accumulated gradients;

determining, by the first worker, whether criteria for updating the current values of the parameters of the deep neural network have been satisfied; and

when the criteria for updating the current values of the parameters of the deep neural network have been satisfied:

updating, by the first worker, the parameters of the deep neural network, wherein the updating comprises:

determining updated values of the parameters using the updated accumulated gradients, and

storing, by the first worker, the updated values in the shared memory accessible to all of the plurality of workers.

18. The method of claim 17, wherein each worker executes independently of each other worker on the same computer.

19. The method of claim 17, wherein the respective action selection policy for each worker is different from the respective action selection policy for each other worker.

20. The method of claim 17, wherein determining the actual long-term reward comprises:
for the last observation received before the criteria are satisfied, setting the actual long-term reward based on the actual reward for the observation and the baseline score for the observation characterizing the state that satisfies the particular criteria; and
for each other observation received before the criteria are satisfied, setting the actual long-term reward for the observation as the actual reward plus a discounted sum of the actual long-term rewards for observations received after the observation.

21. The method of claim 17, wherein the gradient update for the policy neural network is determined by:

$$\nabla_{\theta'} \log \pi(a_t | s_t; \theta') (R_t - b_t)$$

wherein $\pi(a_t | s_t; \theta')$ represents the score for the action performed in response to the observation s_t determined in accordance with the current values of the parameters of the policy neural network, R_t is the actual reward and b_t is the baseline score.

22. The method of claim 21, wherein the gradient update for the baseline neural network is determined by:

$$\partial (R_t - b(s_t; \theta'_b))^2 / \partial \theta'_b$$

wherein $b(s_t; \theta'_b)$ is the baseline score for the observation s_t .

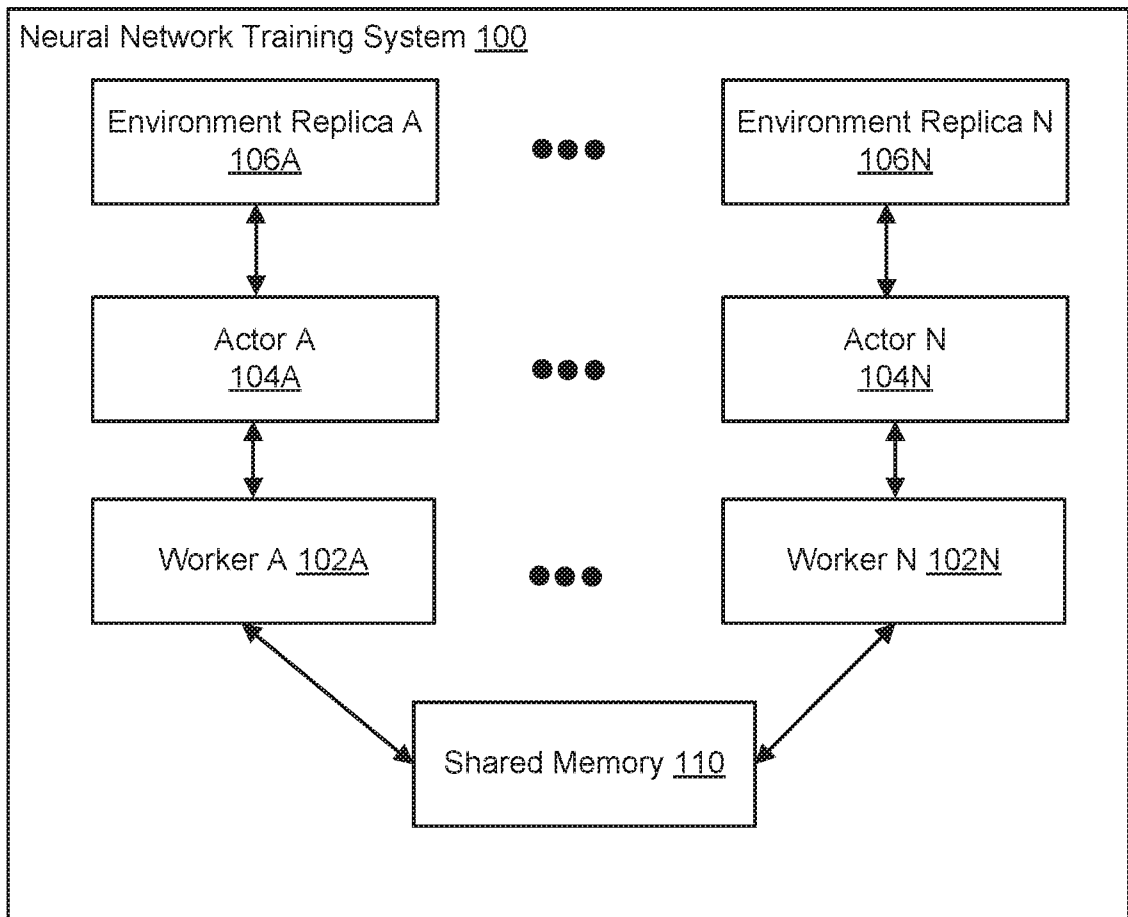


FIG. 1

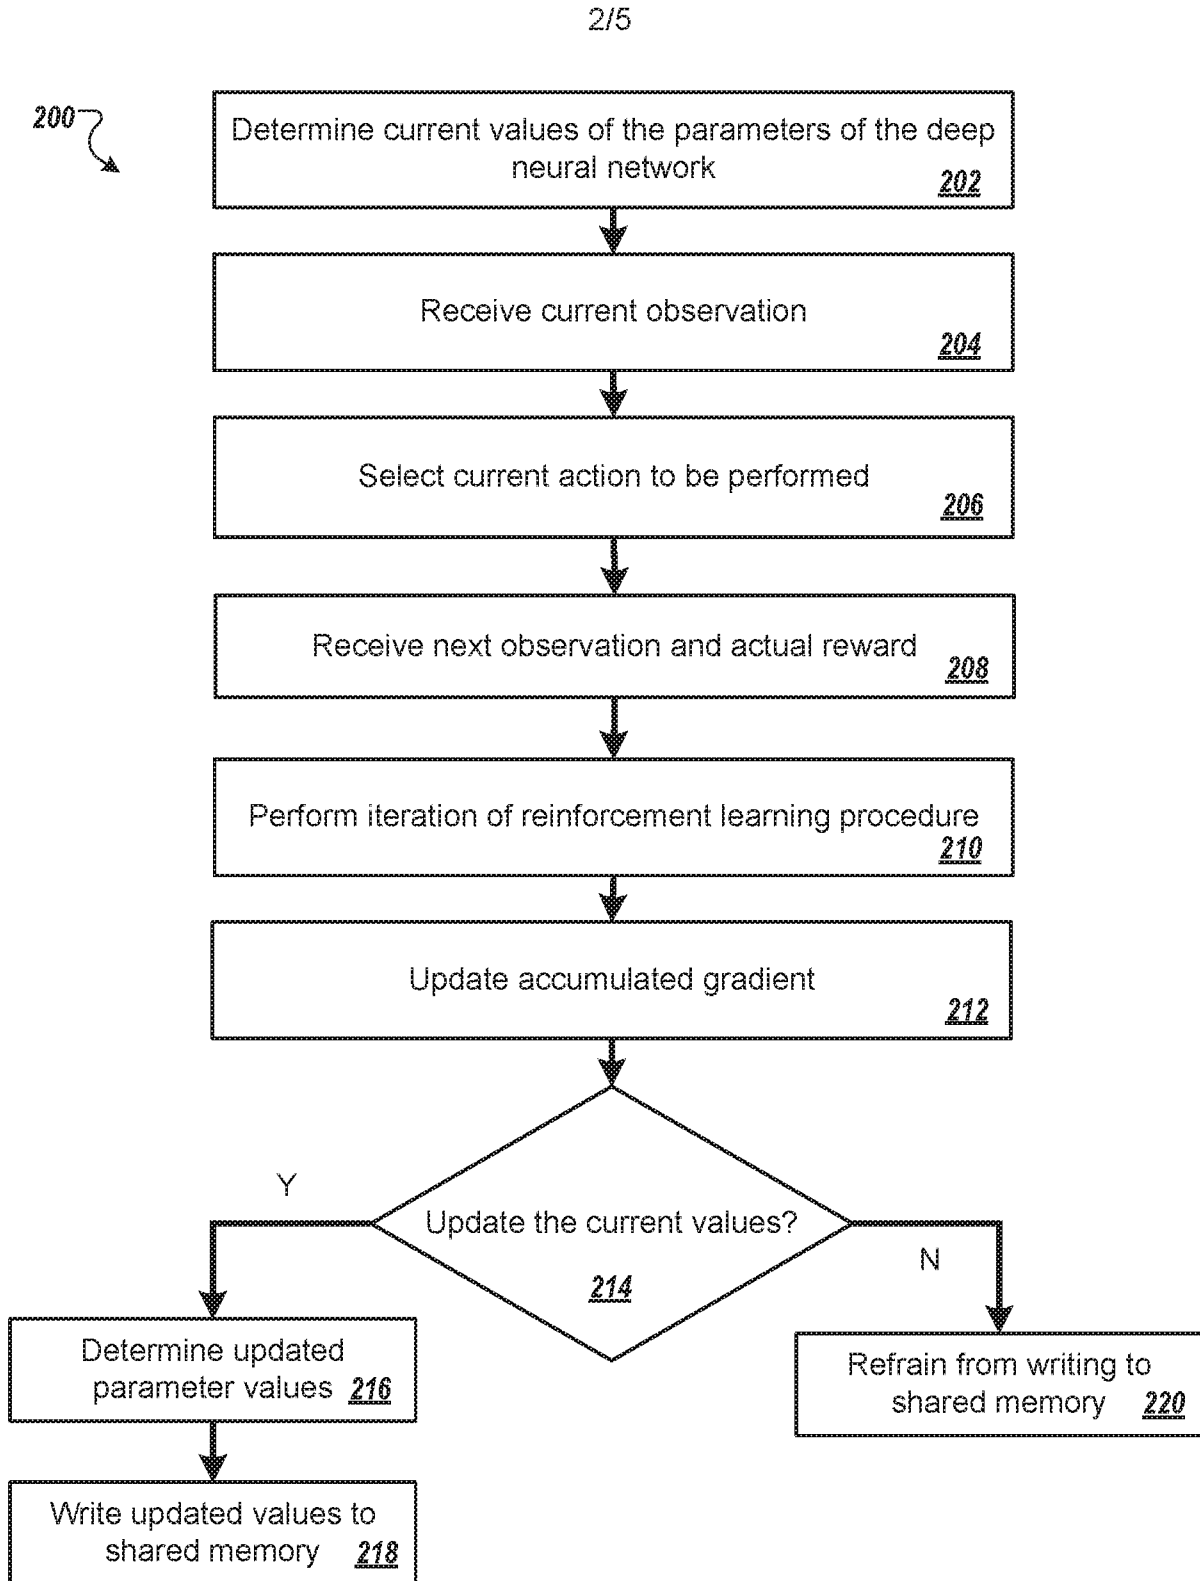


FIG. 2

3/5

300

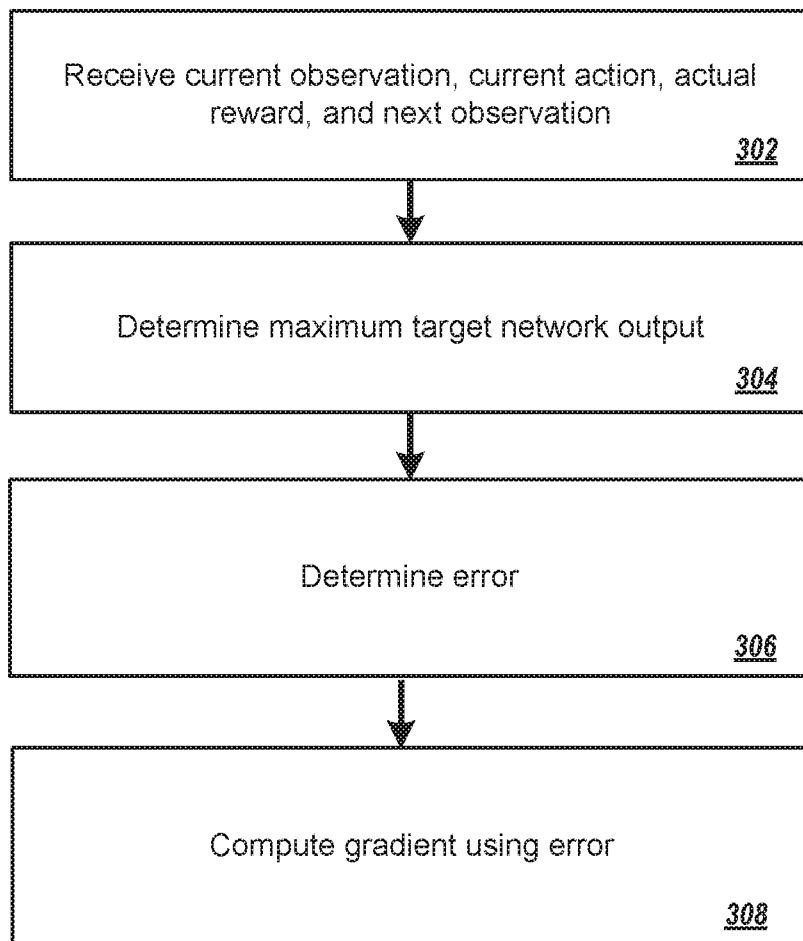


FIG. 3

4/5

400

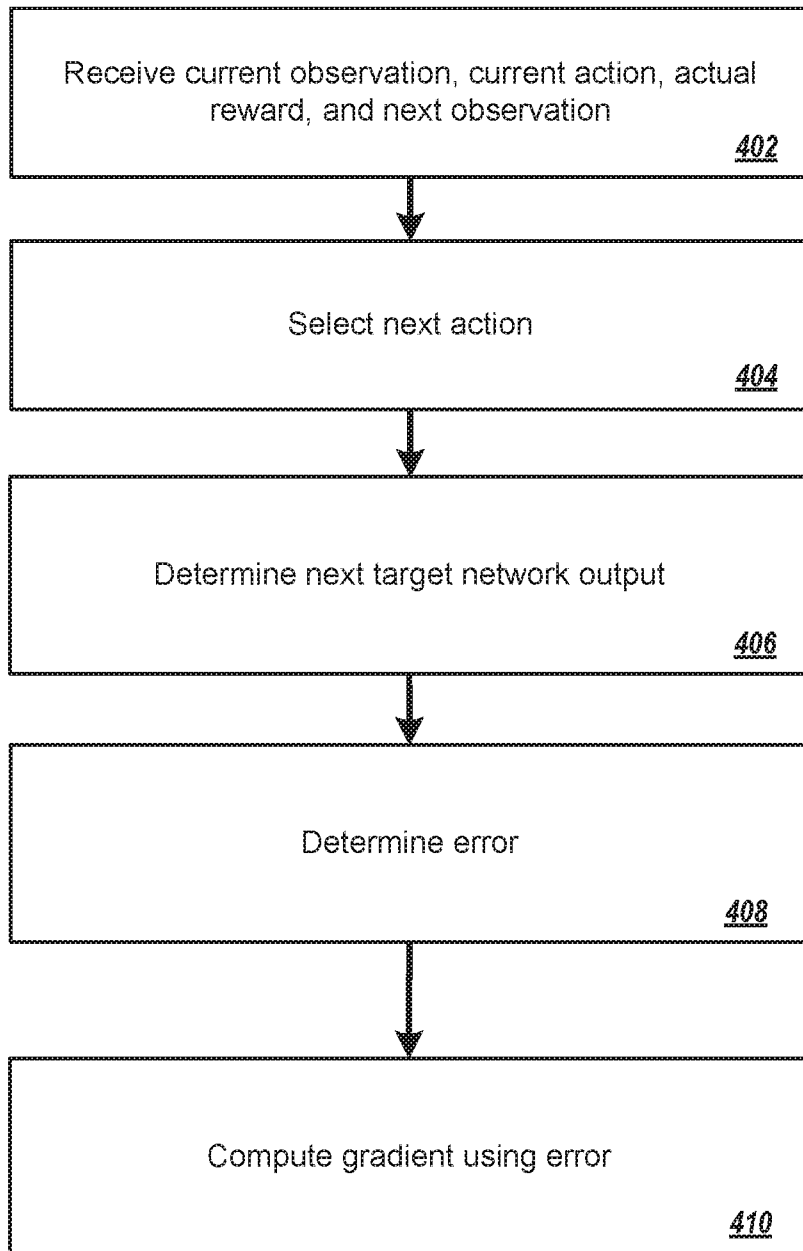


FIG. 4

5/5

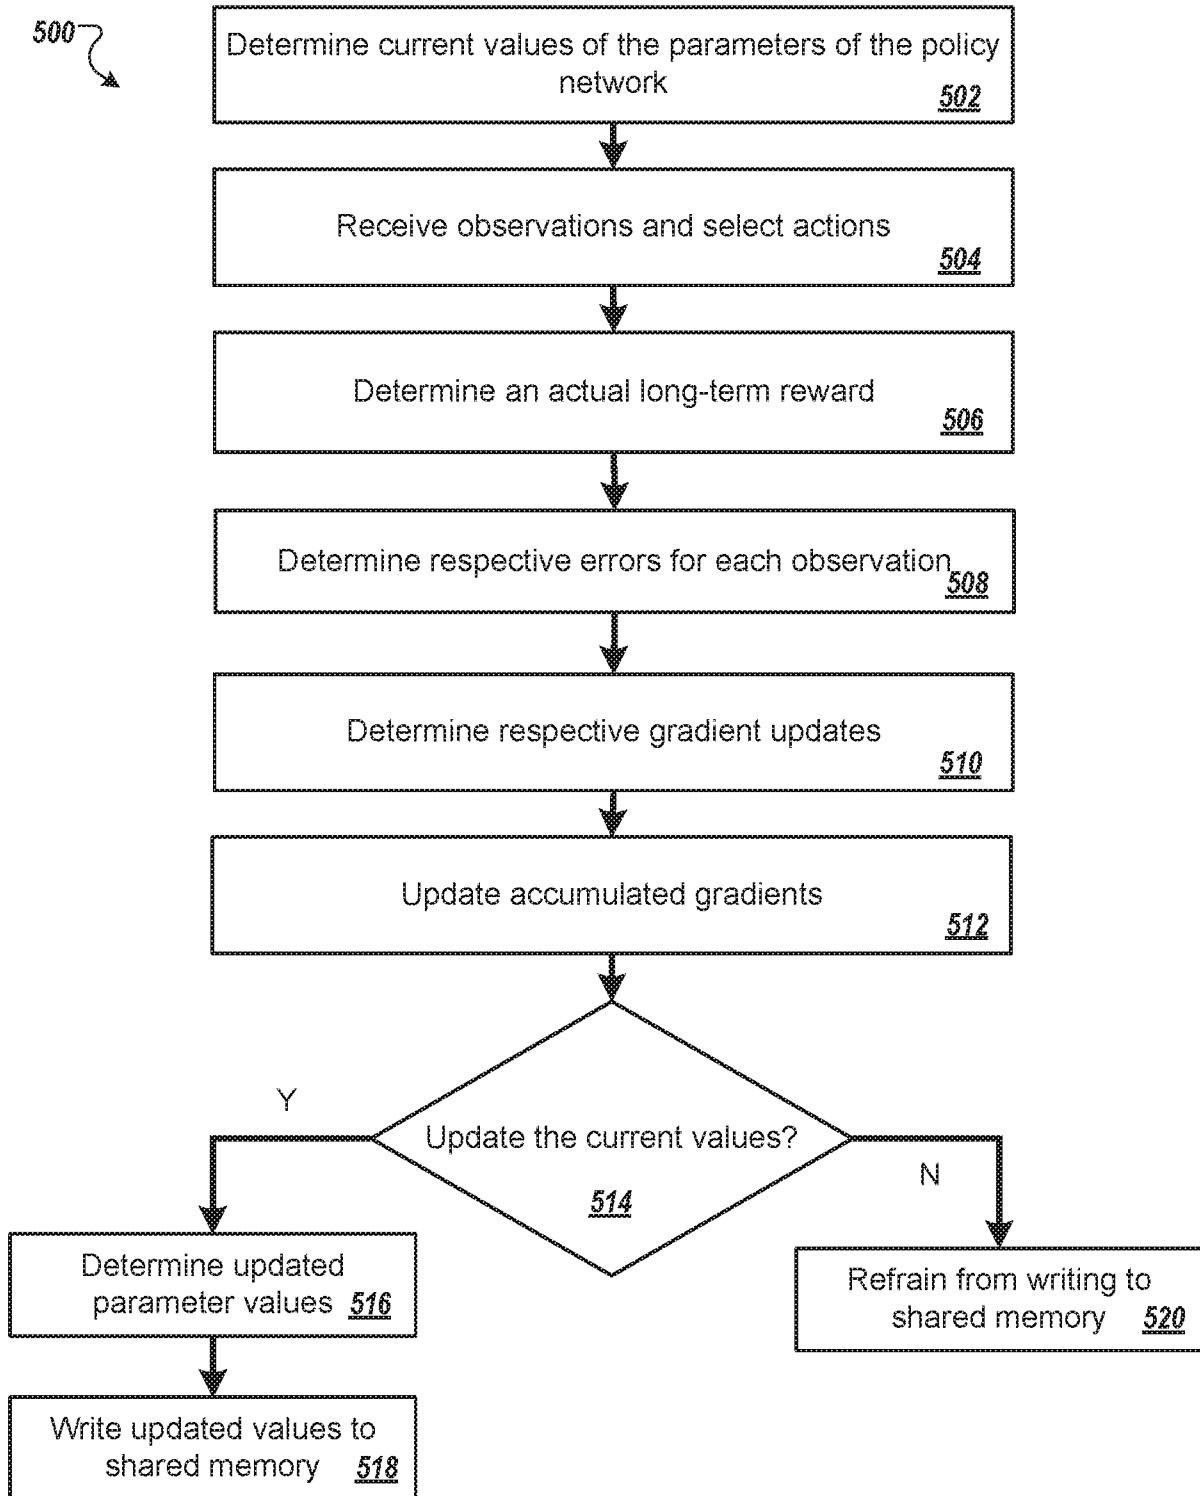


FIG. 5

Neural Network Training System 100

