

(19) **United States**

(12) **Patent Application Publication**

Beer et al.

(10) **Pub. No.: US 2023/0286158 A1**

(43) **Pub. Date: Sep. 14, 2023**

(54) **AUTONOMOUS SENSE AND GUIDE MACHINE LEARNING SYSTEM**

(71) Applicant: **Lawrence Livermore National Security, LLC**, Livermore, CA (US)

(72) Inventors: **N. Reginald Beer**, Pleasanton, CA (US); **David H. Chambers**, Livermore, CA (US); **Jeffrey Edward Mast**, Loveland, CO (US)

G06T 7/70

(2006.01)

G06V 20/10

(2006.01)

(52) **U.S. Cl.**
CPC *B25J 9/1676* (2013.01); *G05D 1/101* (2013.01); *G05B 13/02* (2013.01); *G05D 1/0088* (2013.01); *B64C 39/024* (2013.01); *G06T 7/70* (2017.01); *G06V 20/10* (2022.01); *G06T 2207/30252* (2013.01); *B64U 10/13* (2023.01)

(21) Appl. No.: **18/316,518**

(22) Filed: **May 12, 2023**

Related U.S. Application Data

(62) Division of application No. 16/697,652, filed on Nov. 27, 2019, now Pat. No. 11,685,050.


Publication Classification

(51) **Int. Cl.**
B25J 9/16 (2006.01)
G05D 1/10 (2006.01)
G05B 13/02 (2006.01)
G05D 1/00 (2006.01)
B64C 39/02 (2006.01)


(57) **ABSTRACT**

A system for generating a machine learning system to generate guidance information based on locations of objects is provided. The system accesses training data that includes training time-of-arrival (“TOA”) information of looks and guidance information for each look. The guidance information is based on a training collection of object locations. The TOA of a look represents, for each object location of a training collection of object locations, times between signals transmitted by transmitters and return signals received by receivers. The return signals represent signals reflected from an object at the object location. The system trains a machine learning system using the training data wherein the machine learning system inputs TOA information and outputs guidance information.


110



120



130



140



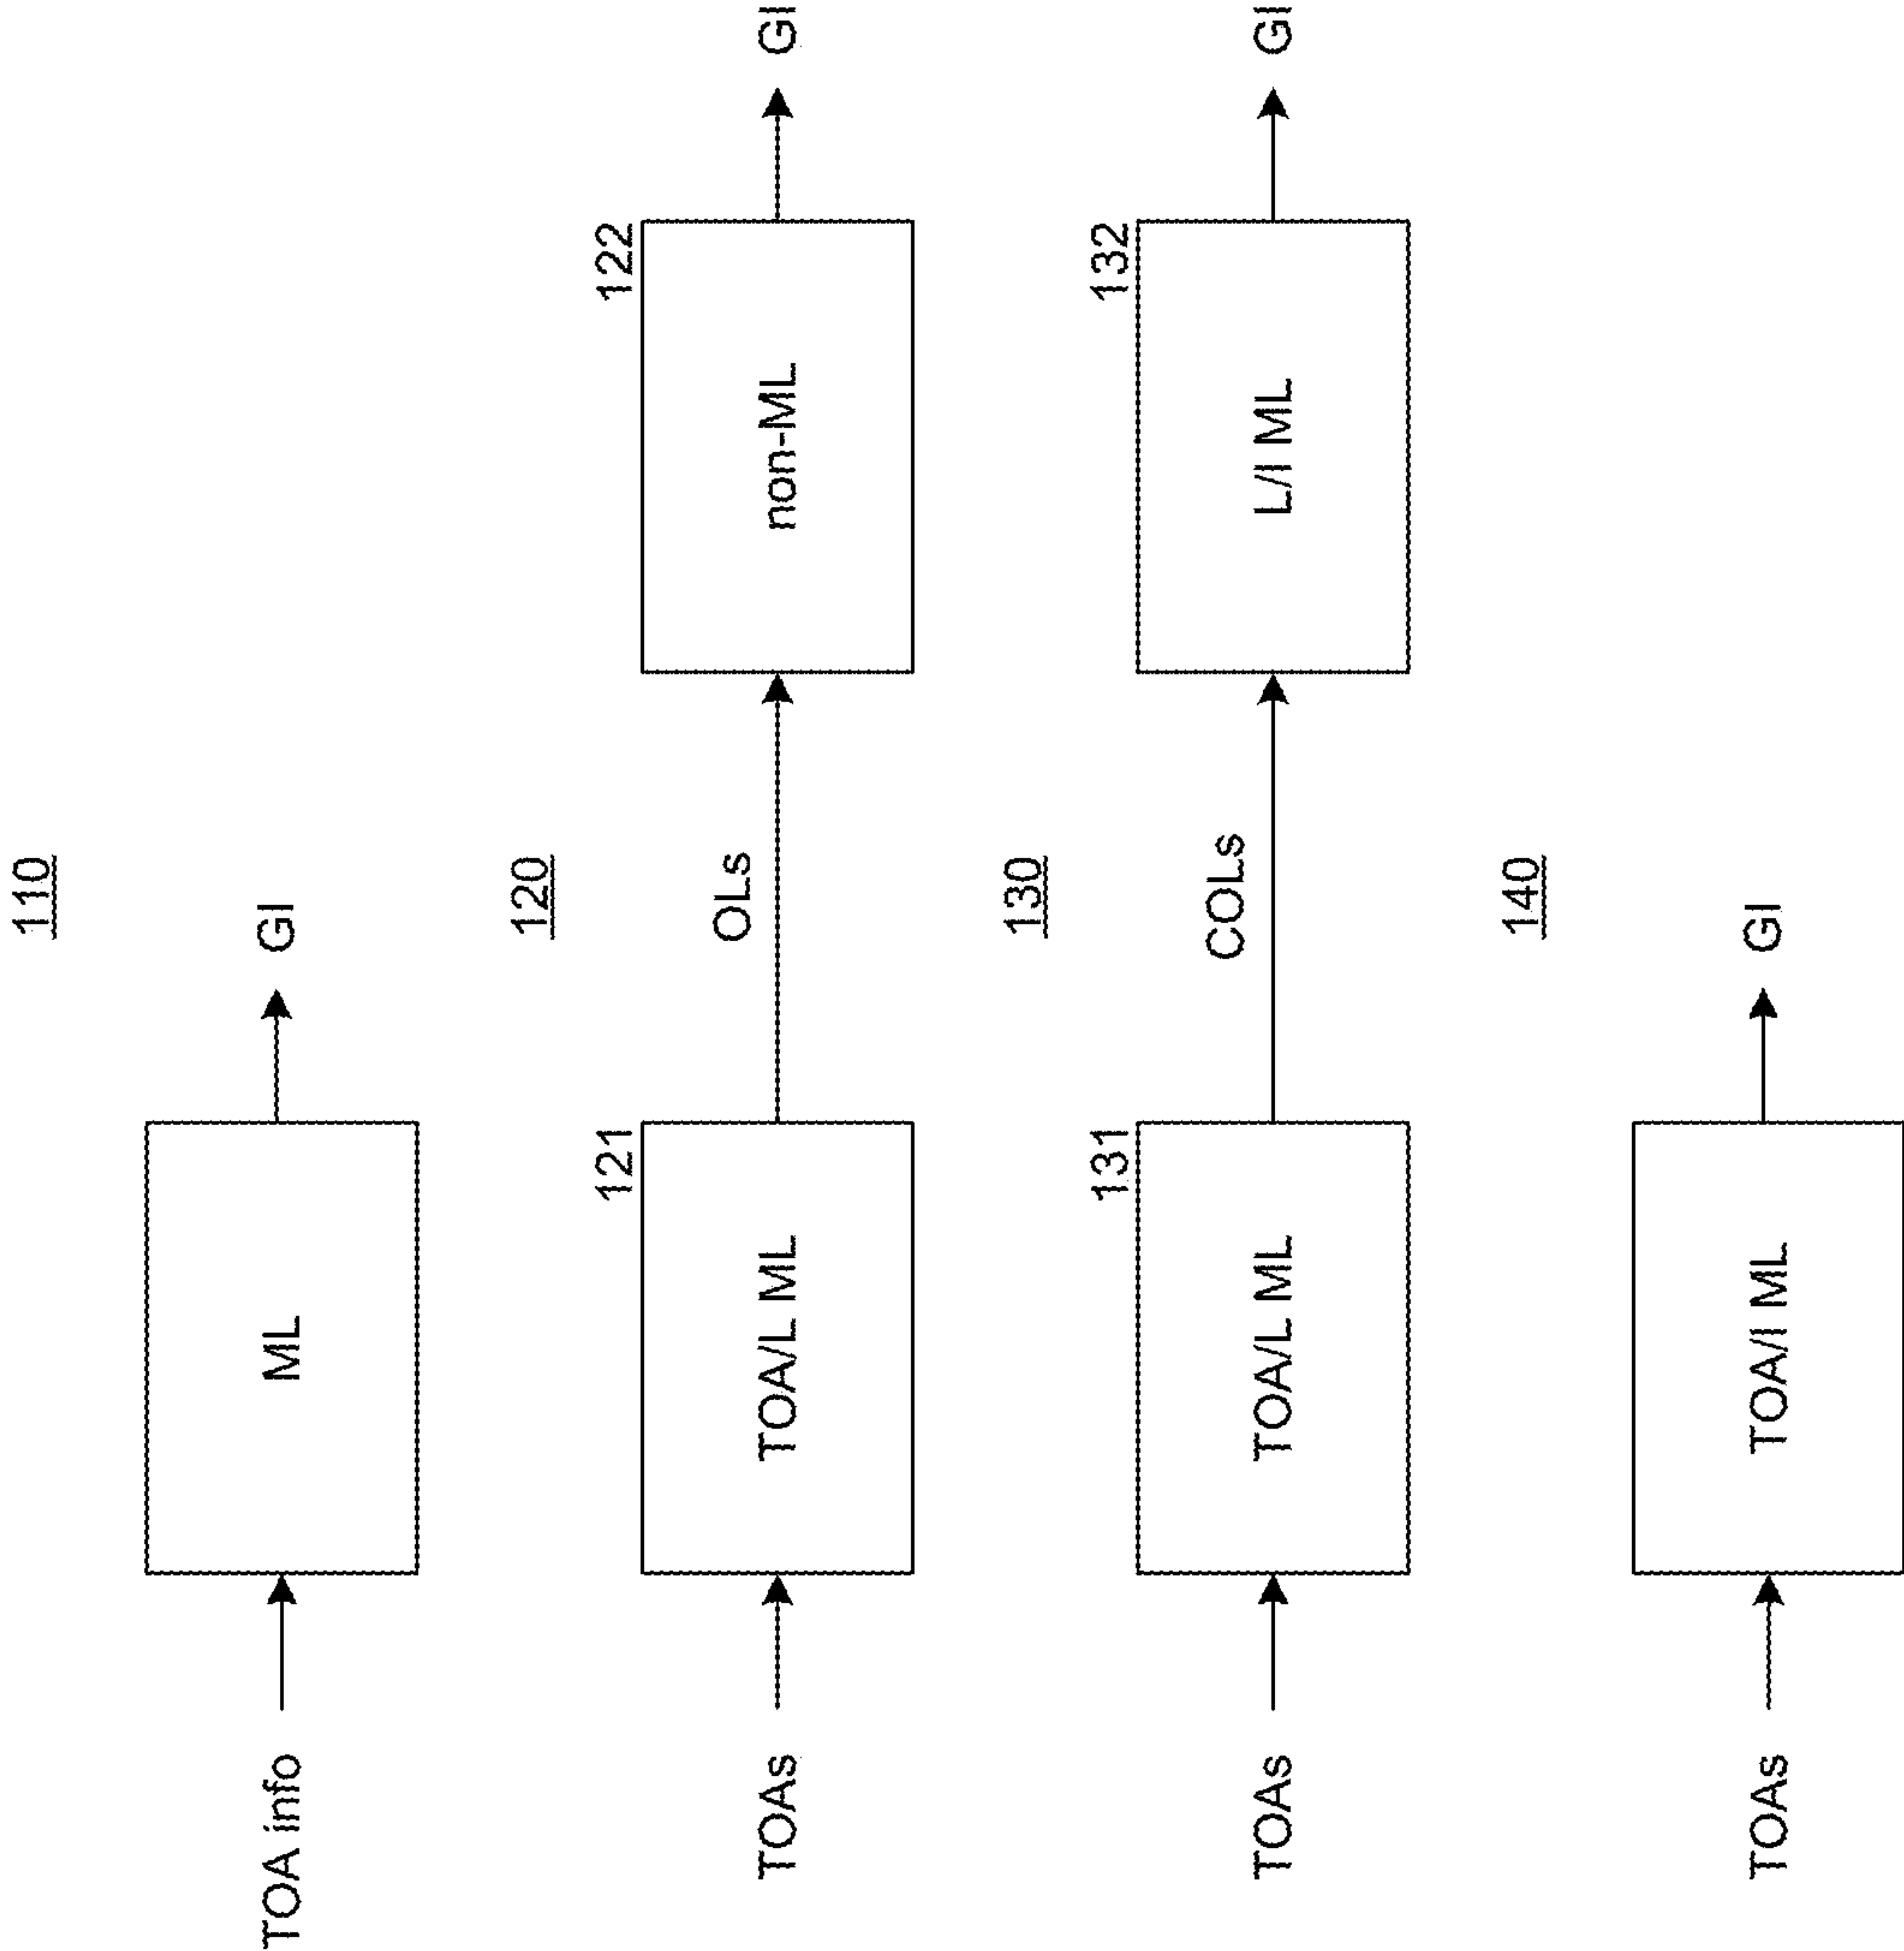


FIG. 1

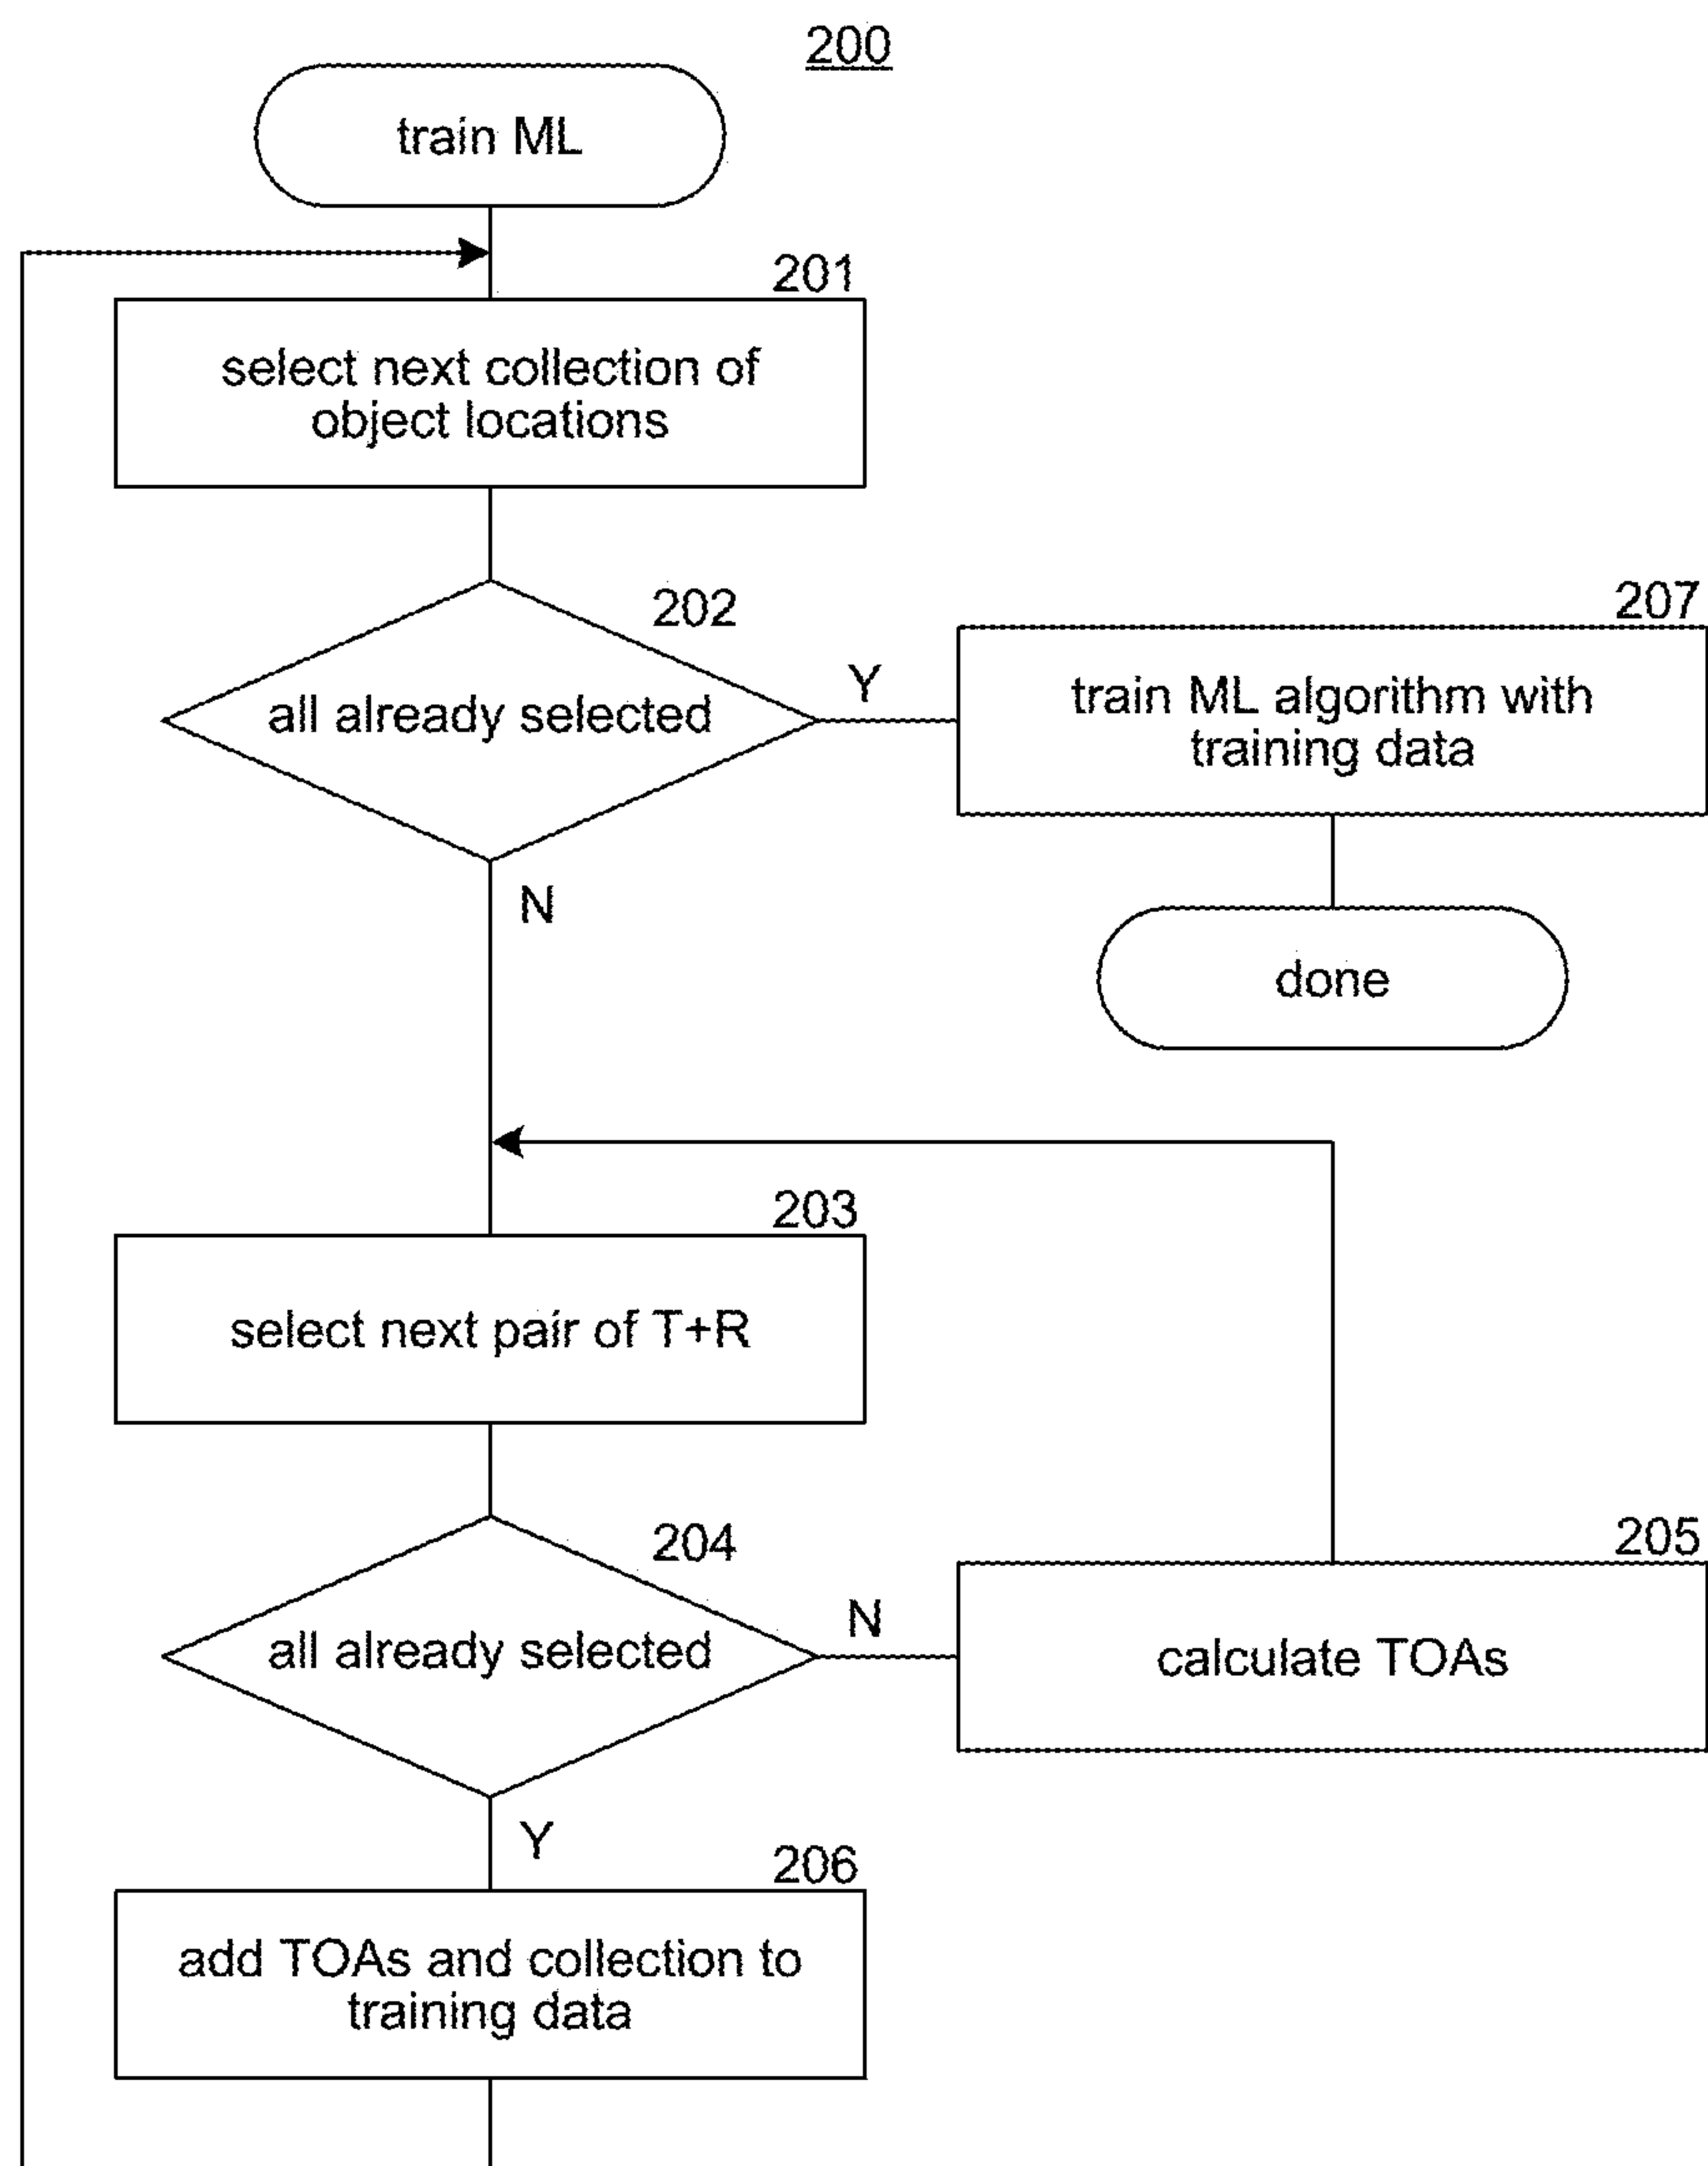


FIG. 2




FIG. 3

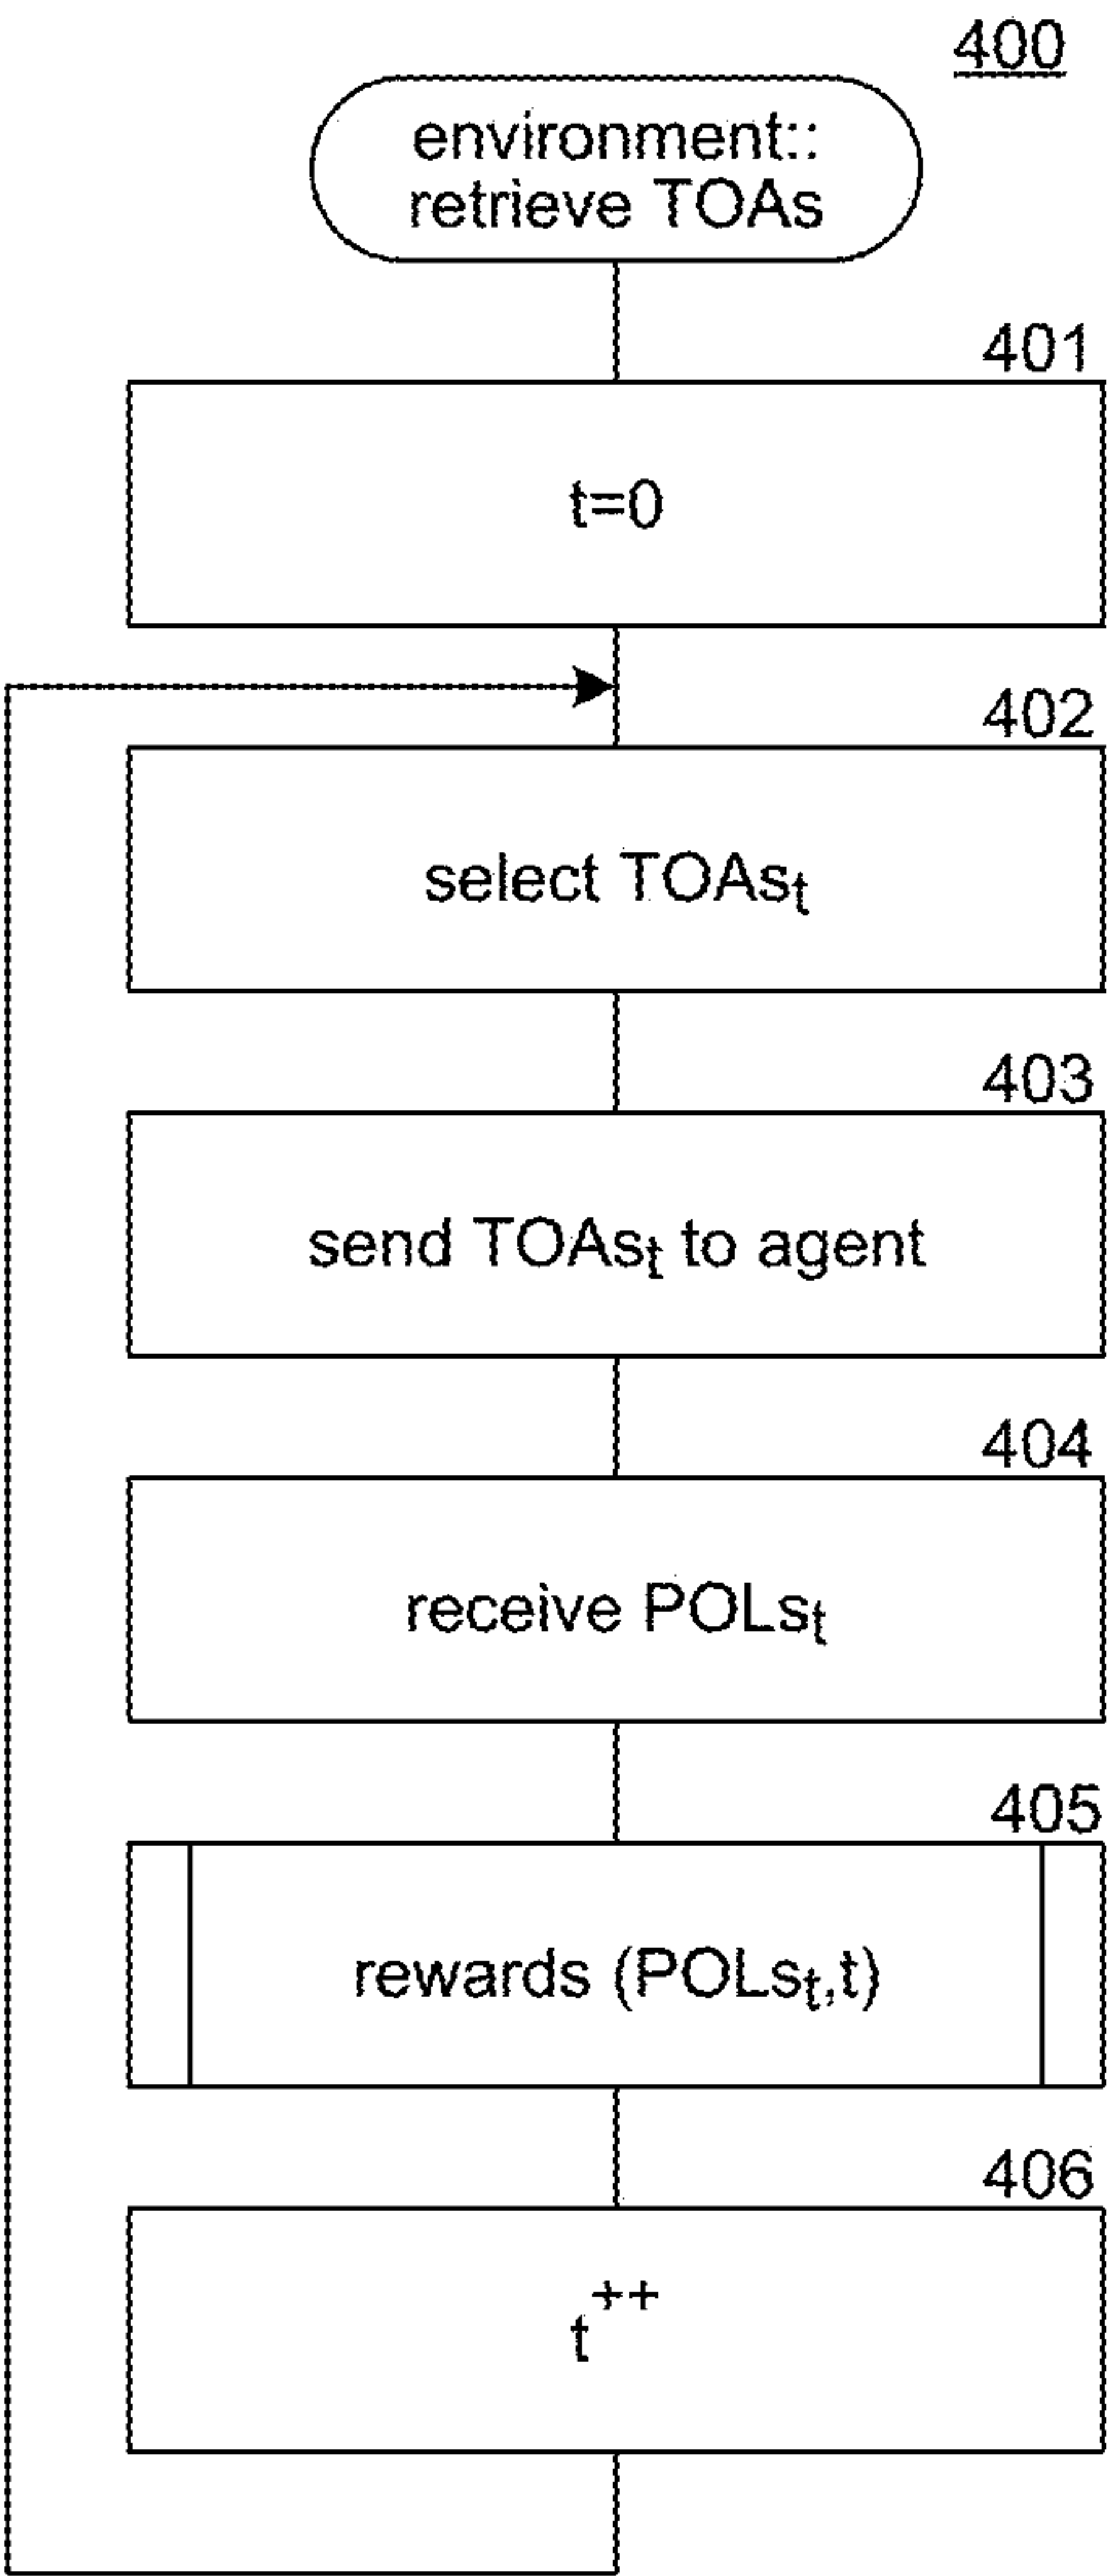


FIG. 4

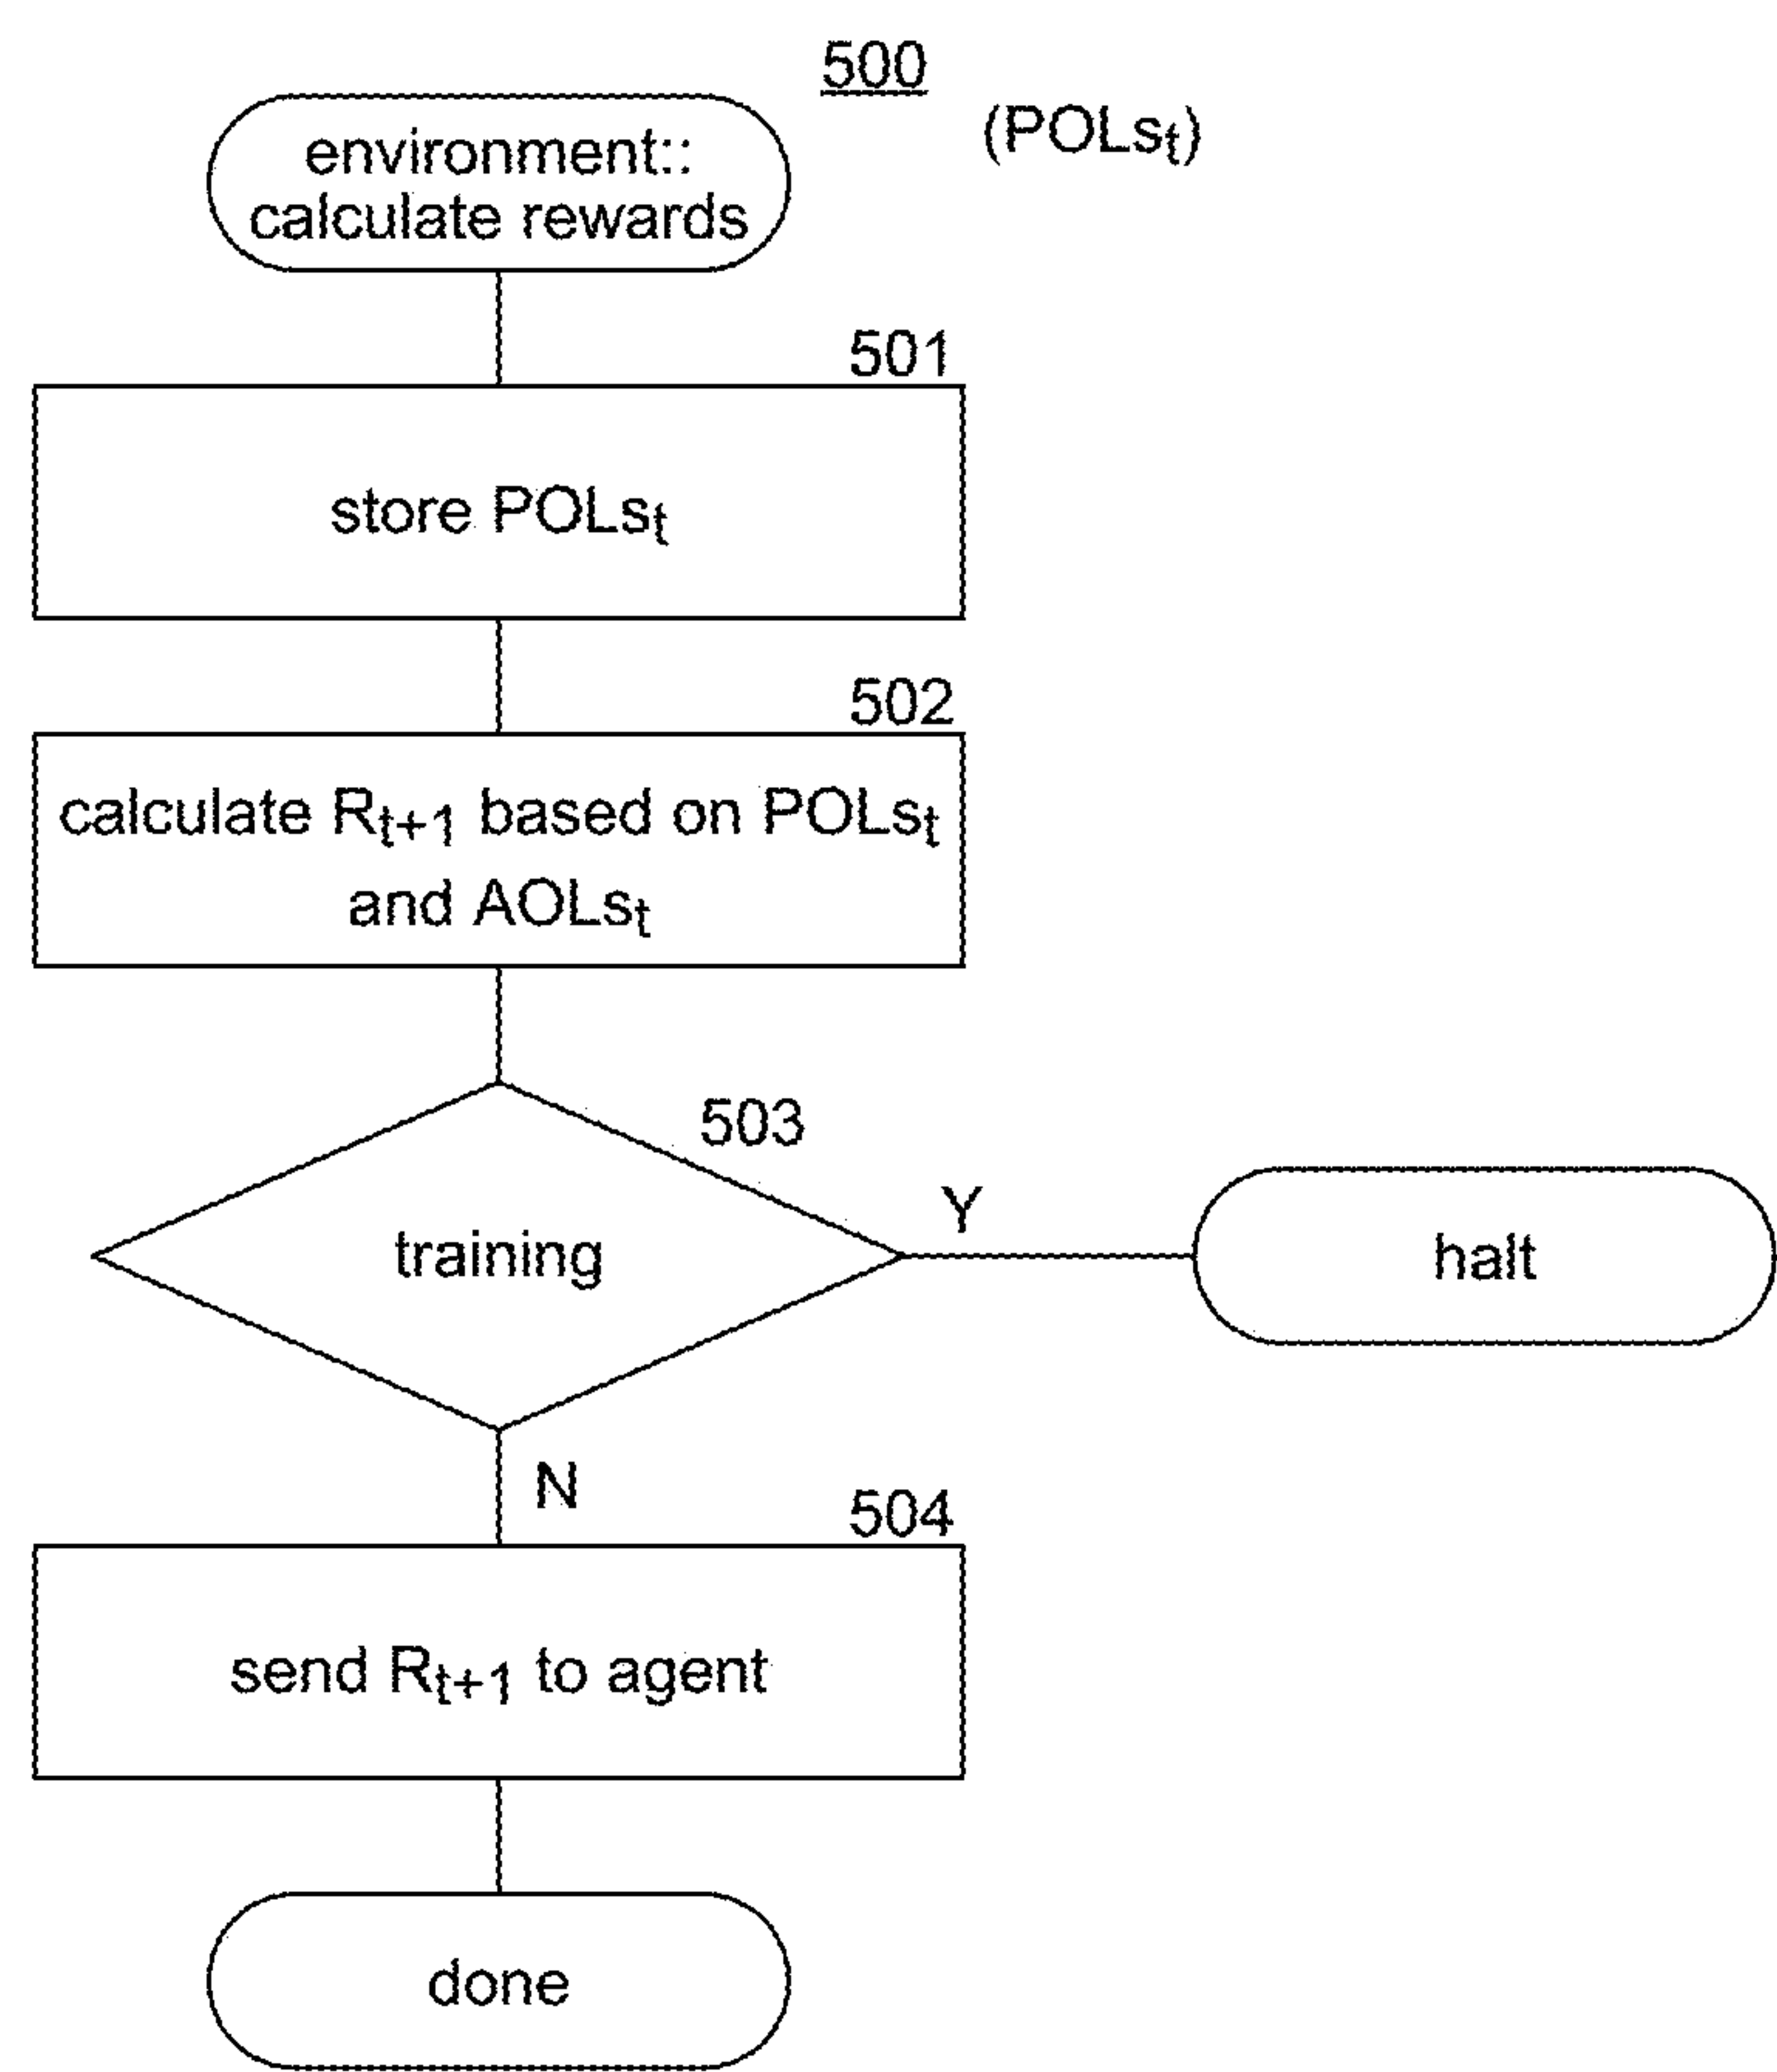


FIG. 5

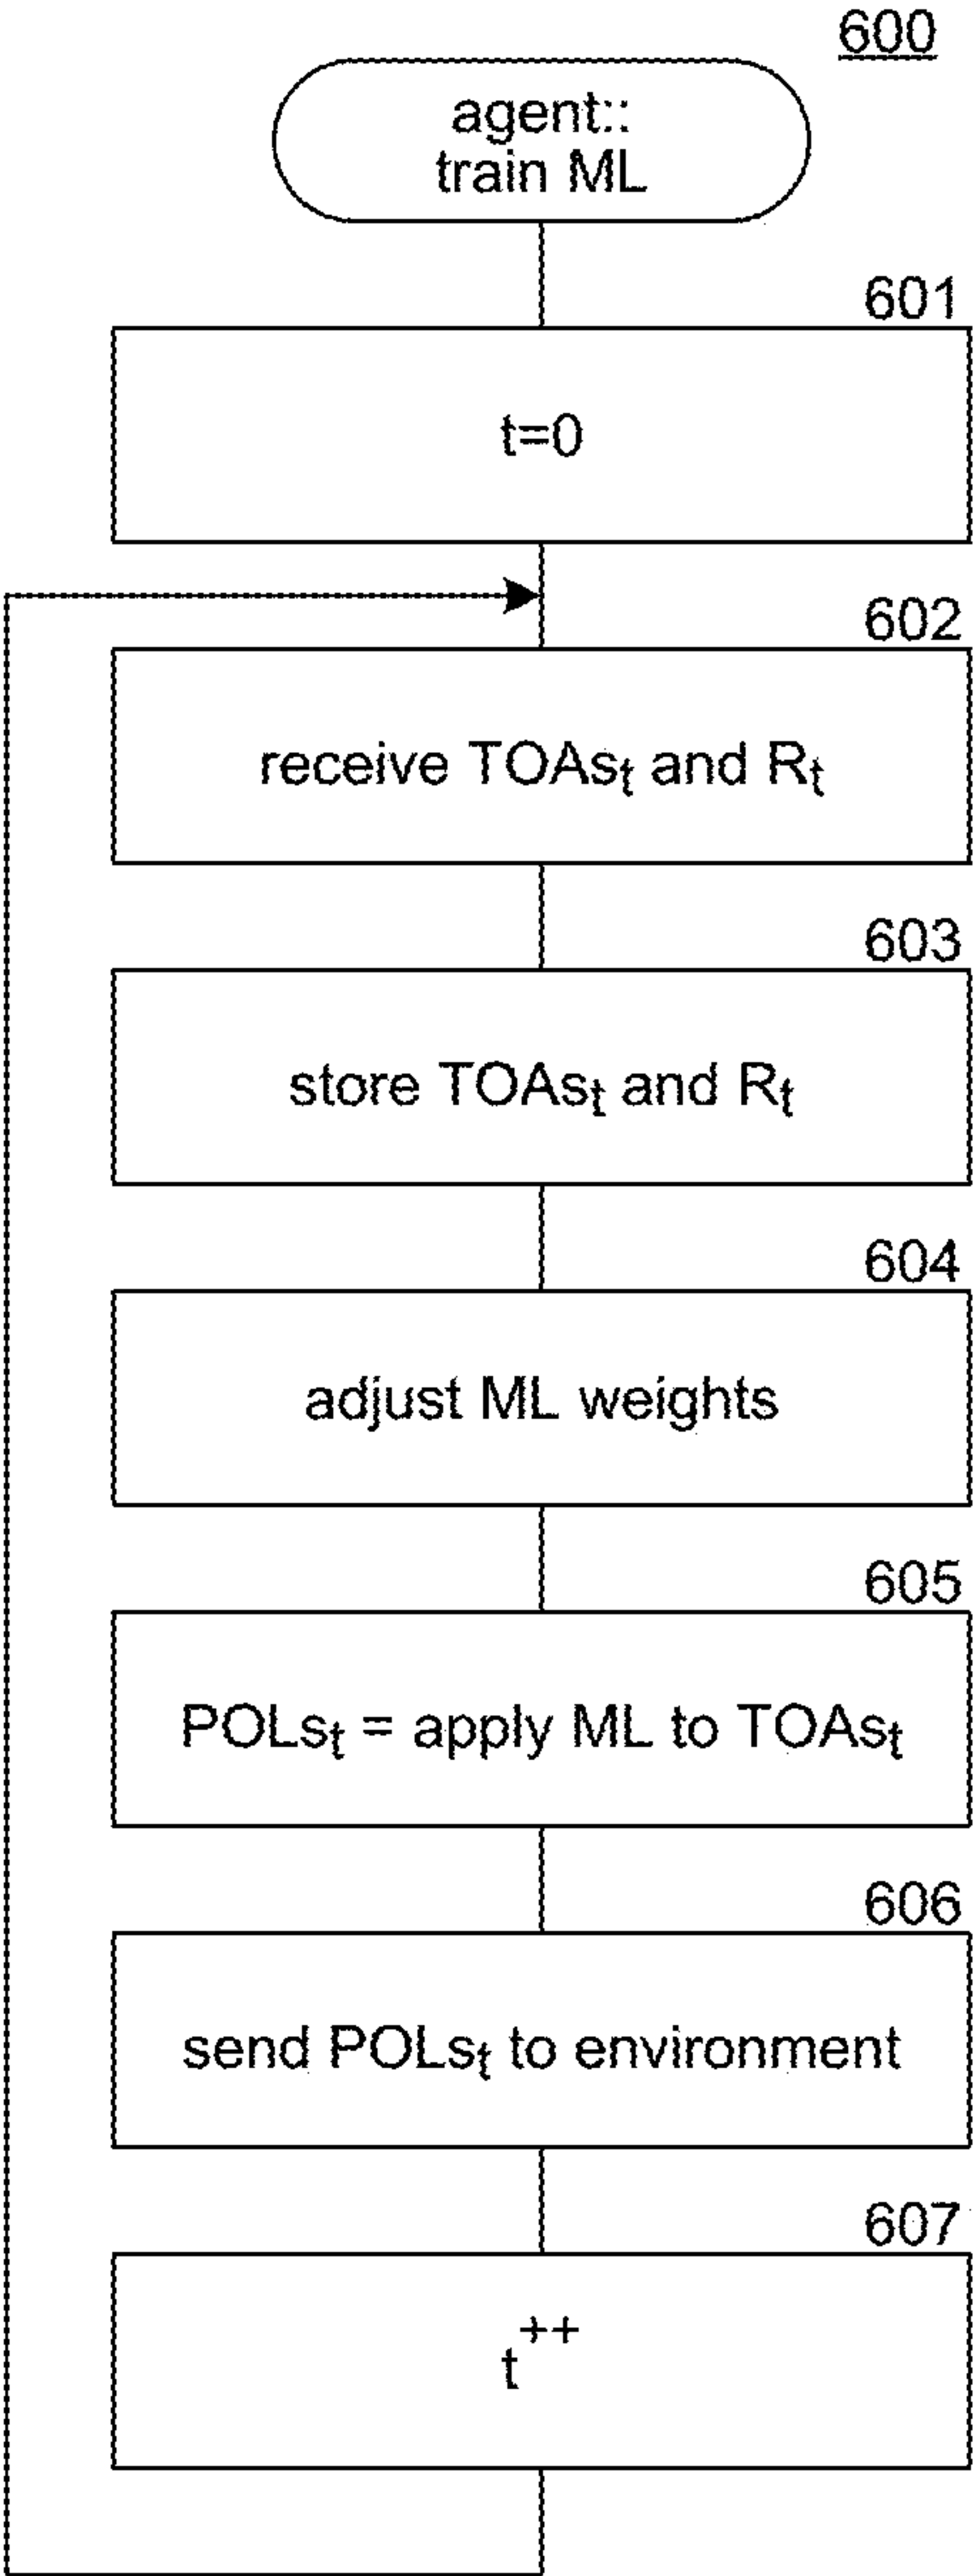


FIG. 6

AUTONOMOUS SENSE AND GUIDE MACHINE LEARNING SYSTEM

CROSS-REFERENCE TO RELATED APPLICATION(S)

[0001] This patent document is a divisional application of and claims priority to U.S. patent application Ser. No. 16/697,652, filed Nov. 27, 2019, which is incorporated herein by reference in its entirety.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH

[0002] The United States Government has rights in this invention pursuant to Contract No. DE-AC52-07NA27344 between the U.S. Department of Energy and Lawrence Livermore National Security, LLC, for the operation of Lawrence Livermore National Laboratory.

BACKGROUND

[0003] Unmanned aerial vehicles (“UAVs”), also referred to as unmanned aircraft systems (“UAS”) or drones, are employed for a wide variety of applications such as military, scientific, commercial, humanitarian, and recreational applications. UAVs can be controlled in various ways that range from autonomous control by a navigation system to remote control by an operator. The navigation (or guidance) systems that provide autonomous control may be on board the UAVs or at ground stations that receive data from the UAVs and transmit instructions to the UAVs. A navigation system may simply navigate the UAV to a destination location along a specified route (e.g., a straight line) defined by Global Positioning System (“GPS”) coordinates. More sophisticated navigation systems may interface with an onboard imaging system that collects images of the environment near the UAV. The navigation system may process the images to identify obstacles in the way of the UAV (e.g., buildings, mountains, and trees) and direct the UAV on a route to avoid the obstacles. When a UAV is under remote control of an operator, the UAV may have an onboard camera system that streams images of the environment to the operator. If the UAV does not have an onboard camera system, the operator needs to have a line of sight to the UAV. The operator may use various cockpit-type controls to guide the UAV.

[0004] Navigation systems that process images to identify obstacles can be very expensive and can add significant weight to a UAV. These navigation systems include camera systems, image processing systems, obstacle detection systems, route planning systems, and so on. A significant portion of the expense is driven by the computational and storage resources that are needed to process the images. Each image may require tens of megabytes for storage. To process the images in real time, the navigation systems may require high-end processing systems and custom-designed processors. The expense and weight of such navigation systems make them impractical for use by UAVs except in high-end military applications. As a result, UAVs for commercial and recreational use are remotely controlled based on having a line of sight to the UAVs.

[0005] Some UAV systems employ an on-board sensor array rather than an imaging device to identify object locations of objects (e.g., to be avoided by the UAV or are destinations of the UAV) in a field-of-view of the sensor array. The sensor array includes transmitters that transmit

signals at intervals and receivers that collect return signals, which are transmitted signals reflected by objects. The sensor array may operate in monostatic mode or multistatic mode. In monostatic mode, the signal transmitted by a transmitter is received only by the receiver of that same transceiver. In multistatic mode, each transmitter transmits in sequence, but the return signal is collected by multiple receivers, generally all the receivers. When in multistatic mode, if the sensor array has M transmitters and N receivers the sensor array collects M x N return signals. Each return signal received by a receiver represents the reflection from each object in the field-of-view of the sensor array. For example, if there are three objects, the return signal received by a receiver will include three return pulses corresponding to the reflection of objects. The return pulses from objects that are closer to the sensor array will be received before those that are farther. The time-of-arrival of each return pulse represent the time between the transmitting of a signal and receiving of a return pulse. The set of M x N time-of-arrivals from each transmitter and receiver pair for each object is referred to as a “look.”

[0006] Once a look is collected at an interval, a UAV system typically performs a computation to determine the object location of each object given a look. A guidance system can then input the object locations and generate a guidance instruction for the UAV. Unfortunately, it can be computationally very expensive to determine the object locations. The computational expense is due in part to the difficulty in determining which return pulses correspond to which objects. Algorithms to determine object locations may employ triangulation and solve complex minimization problems to accurately determine the object locations. Unfortunately, UAVs typically do not have the computational power to perform that computation at an interval rate that allows effective guidance.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1 is a block diagram that illustrates verse architectures of SGML system in some embodiments.

[0008] FIG. 2 is a flow diagram that illustrates the processing of a train TOA/location ML component of the SGML system in some embodiments.

[0009] FIG. 3 is a block diagram that illustrates the components of the reinforcement learning component of the SGML system in some embodiments.

[0010] FIG. 4 is a flow diagram that illustrates the processing of a retrieve TOAs component of the environment system in some embodiments.

[0011] FIG. 5 is a flow diagram that illustrates the processing of a calculate rewards component of the environment system in some embodiments.

[0012] FIG. 6 illustrates the processing of a train ML component of the agent system in some embodiments.

DETAILED DESCRIPTION

[0013] Methods and systems for employing a machine learning system to assisting in guiding movement of a platform based on time-of-arrivals (“TOAs”) are provided. In some embodiments, a sense and guide machine learning (“SGML”) system trains a machine learning system to provide guidance information based on TOA information. The guidance information may be object locations or guidance instructions, and the TOA information may be TOAs or

object locations. The trained machine learning system may be deployed on a platform (e.g., UAV) with a sensor array of transmitters and receivers to determine object locations for use in guiding the platform. To train the machine learning system, the SGML system generates training data that includes TOA information labeled with guidance information. For example, the TOA information may be a look, and the labels may be a collection of object locations corresponding to the look. After the training data is generated, the SGML system performs the training of the machine learning system using the training data. The trained machine learning system can then be deployed (e.g., on a platform or remotely from a platform) to input TOAs and output guidance information. The guidance information can then be employed by a guidance system to guide the platform to avoid objects and/or to travel to an object. Although it may be computationally very expensive to generate the training data and train the machine learning system, the generating and training are performed off-line prior deployment for a platform. After training, the computations needed to determine guidance information are a very small fraction of that needed for the training and can be performed in real-time by an on-board computing system. Moreover, the on-board computing system can be relatively lightweight so that it can be installed on platforms where weight is a concern. The SGML system thus determines guidance information in a way that avoids the computational expense of prior systems that determined object locations, for example, by solving minimization problems.

[0014] In some embodiments, the SGML system trains a TOA/location machine learning system to determine object locations given a look and employs the trained TOA/location machine learning system to determine object locations. The trained TOA/location machine learning system may be deployed on a platform (e.g., UAV) with a sensor array of transmitters and receivers to determine object locations for use in guiding the platform. To train the TOA/location machine learning system, the SGML system generates training data that includes looks with each look being labeled with a collection of object locations. A collection of object locations may include one or more object locations. A look represents the expected TOAs of return signals given a collection of object locations of the label of the look. After the training data is generated, the SGML system performs the training of the TOA/location machine learning system using the training data. The trained TOA/location machine learning system can then be deployed (e.g., on a platform or remotely from a platform) to input observed looks and output an observed collection of object locations for each look. The observed collections of object locations can then be provided to a guidance system to generate guidance instructions to guide the platform to avoid objects and/or to travel to an object.

[0015] In some embodiments, the SGML system may employ a machine learning system to determine a guidance instruction based on TOAs. The SGML system trains a TOA/instruction machine learning system to generate guidance instructions given a look and employs the TOA/instruction machine learning system to guide movement of a platform. To train the TOA/instruction machine learning system, the SGML system generates training data that includes looks with each look being labeled with a guidance instruction. After the training data is generated, the SGML system performs the training of the TOA/guidance machine

learning system using the training data. The trained TOA/guidance machine learning system can then be deployed (e.g., on a platform or remotely from a platform) to input observed looks and output a guidance instruction for each look. A guidance system on-board a platform can use the guidance instruction to guide movement of the platform. Like the TOA/location machine learning system, the TOA/guidance machine learning system can be trained off-line and the trained TOA/guidance machine learning system can be deployed to a platform to generate guidance instructions in real-time.

[0016] In some embodiments, the SGML system may employ a machine learning system to determine a guidance instruction based on a collection of object locations. The SGML system trains a location/instruction machine learning system to generate guidance instructions given a collection of object locations (e.g., determined by a TOA/location machine learning system) and employs the location/instruction machine learning system to guide movement of a platform. To train the TOA/instruction machine learning system, the SGML system generates training data that includes collections of object locations with each collection being labeled with a guidance instruction. After the training data is generated, the SGML system performs the training of the location/guidance machine learning system using the training data. The trained location/guidance machine learning system can then be deployed (e.g., on a platform or remotely from a platform) to input observed collections of object locations and output a guidance instruction for each collection. A guidance system on-board a platform can use the guidance instruction to guide movement of the platform. Like the TOA/location machine learning system, the location/guidance machine learning system can be trained off-line and the trained location/guidance machine learning system can be deployed to a platform to generate guidance instructions in real-time.

[0017] In some embodiments, the SGML system may employ various techniques to generate the training data and various machine learning techniques to train a machine learning system. For example, to generate training data, the SGML system may generate collections of object locations and calculate, for each transmitter and receiver pair a TOA based on distances between the transmitter and an object and the receiver and the object and the speed of the transmitted signal (e.g., speed of light or sound). The SGML system may employ a simulator to simulate the TOAs factoring the transmission medium, object shape, object reflectivity, object size, and so on. Alternatively, or additionally, the SGML system may calculate TOAs based on transmitted and return signals generated during actual movement of a platform and objects. For example, the platform and objects may move within a training pen. At each interval, the TOAs may be calculated by a computing system on board the platform or by a remote computing system based on return signal information provided by the platform. In addition, the SGML may employ an on-board or remote sensor to determine the actual object locations at each interval. These sensors may be a LIDAR system, a multicamera system, acoustic system, and so on. The machine learning techniques may include supervised and unsupervised learning techniques. The supervised learning technique may include a neural network, random forest, decision trees, support vector machines, and so on. The unsupervised learning techniques may include reinforcement learning.

[0018] In some embodiments, the SGML system may train a machine learning system based on series of looks of sequential intervals with each series having a label. For example, the SGML system may generate training data that includes guidance information as a label for each sequential pair of looks or collections of object locations. Once trained, the SGML system may be used to input a sequential pair of looks to generate guidance information. When training is based on series of looks, the SGML system may be considered to factor in object direction, velocity, and acceleration when determining guidance.

[0019] In some embodiments, the SGML system may be part of an object sense and avoid (“OSA”) system on board a platform. The OSA system employ the machine learning system to determine guidance information that is provided to a flight controller system. The OSA system repeatedly uses the sensor array to collect sensor data of any objects in an object field (i.e., field of perception) around the platform. For example, the sensor array may transmit radar signals and receive the return signals that are reflected by the objects. Given guidance information, the flight controller controls the platform to move based on the guidance information.

[0020] The SGML system may support the guidance of a variety of autonomous vehicles (“AVs”) that are autonomously driven. The AVs may include UAVs, unmanned ground vehicles (“UGVs”), unmanned underwater vehicles (“UUVs”), and unmanned space vehicles (“USVs”). These vehicles are “unmanned” in the sense that a person does not control the guidance of the vehicle. For UUVs, the sensor array may be sonar-based. For USVs, the SGML system may be particularly useful in helping a satellite avoid collisions with space debris or other satellites. The sensor array for USVs may employ a larger field of perception that for a UAV to encompass a wider approach of objects. In addition, the SGML system may be used with a guidance system that factors in both the guidance information provided by the SGML system and estimated locations of known space objects determined from orbital parameters (e.g., Keplerian elements) of the space objects to help in determining whether return signals correspond to an object. Also, although the SGML system is described primarily based on a field of perception that is in front of the platform, the field of perception may surround the platform. In such a case, the platform may include multiple sensor arrays to sense the entire area around the platform. Such a surrounding field of perception may be useful, for example, to sense and avoid objects (e.g., space debris) that are traveling toward the platform from the rear.

[0021] In some embodiments, the SGML system may be employed in conjunction with an auxiliary object sense and avoidance system such as that described in PCT Publication No. WO 2018/190834 A1, entitled “Attract-Repel Path Planner System for Collision Avoidance,” filed on Apr. 12, 2017, which is hereby incorporated by reference. When objects are nearby a platform, such an auxiliary system may generate more precise guidance information based on a more accurate determination of object locations. When used in conjunction with such an auxiliary system, the SGML system may be employed while all objects are far, and the auxiliary system may be then employed while an object is not far.

[0022] The SGML system may provide guidance information so that a UAV takes evasive maneuvers to avoid an imminent collision with an object. For example, a predator

UAV may be attempting to intercept a prey UAV before it reaches its target by colliding with the prey UAV. In such a case, the SGML system of the prey UAV may receive new guidance information resulting in a change in travel direction that is more than 90° away from the target direction because the predator UAV moved from left of the prey UAV to immediately in front of the prey UAV. If the prey UAV was traveling in the target direction, the sudden and significant change in the travel direction by the prey UAV is effectively an evasive maneuver to avoid colliding with the predator UAV. The new guidance information may result in the prey UAV rapidly ascending or descending or even reversing direction. If the prey UAV has a surrounding field of perception, then the SGML system can provide guidance information to effect evasive maneuvers even if the predator UAV approaches from the rear.

[0023] The SGML system may also be used to control movement in robotic systems such as fixed-base robotic systems and free-moving robotic systems. A fixed-base system, such as those used on production lines, typically includes a robot manipulator, an end effector (e.g., tool or gripper), and a safety interlock system. The fixed-base robotic system may be taught its motion using a teach pendant, which is typically a handheld unit used to program the trajectory of the robot. The safety interlock system may include force sensors to detect a collision or a light curtain sensor to disable the robot manipulator when a person (i.e., unsuspected object) is near the workspace of the fixed-base robotic system. The SGML system allows a fixed-base robotic system to detect intrusions and alter the trajectory of the robot manipulator. As a result, worker safety can be improved, and throughput of a capital-intensive production line can be maintained by avoiding costly shutdowns. In addition, use of the SGML system can eliminate the need to teach with a teach pendant and the need for force sensors or a light curtain sensor. The SGML system can be used to provide guidance information for picking up and placing parts for a production line. The SGML system can be used to direct the end effector to a target location of a part within the cell of the fixed-base system for pickup while avoiding neighboring cells, for example, by controlling the roll, pitch, and yaw of the robot manipulator. The SGML system can then be used to direct the end effector to a target that is the desired location of a part on a production line. The output of the SGML system can be used to generate and issue commands to the robot manipulator to slow down, change orientation, or signal readiness to dock, interlock, mate, operate a tool, and so on. A sensor array may be located on the robot manipulator near the end effector or at a fixed location.

[0024] Free-moving robotic systems include servant robots and companion robots. The SGML system can be used to generate guidance information when a robot is moving parts or supplies within a hospital, production line, shipping facility, and so on. The payload of the free-moving robotic system may be a tool, part, supply, sensing system, interactive communicator (for a companion robot), and so on. The guidance information allows the robot to move while avoiding stationary and moving objects. A sensor array may be located on the front of a free-moving robotic system.

[0025] In some embodiments, the SGML system may be employed to design an architecture of a sensor array that satisfies accuracy, power consumption, weight, and so on

tradeoffs. Although a sensor array with more transmitters and receivers will in general result in more accuracy, a sensor array with a large number of transmitters and receivers may be too heavy or power-hungry for a platform. To evaluate architectures for a sensor array, the SGML system may be adapted to process TOAs resulting from actual object locations for that architecture. The SGML system may then be used to determine a collection of object locations. An accuracy metric may be calculated that indicates the similarity between the actual object locations and the collections of object locations. The accuracy metrics, power consumption, and weights of the architectures can be evaluated to identify an appropriate architecture. The architectures may include different numbers and types, locations, location curvatures, different beam patterns (e.g., main and side lobe patterns), and so on of transmitters and/or receivers. As an example, when a phased-array radar sensor is employed that includes a fixed number of transmitters and receivers, the SGML system may be employed to evaluate enabling of different numbers and/or positions of transmitters and receivers based on a power consumption and accuracy tradeoff.

[0026] FIG. 1 is a block diagram that illustrates various architectures of SGML system in some embodiments. Architecture **110** illustrates the overall architecture of the machine learning systems. The machine learning architectures input TOA information, such as TOAs or a collection of object locations associated with the look, and output guidance information, such as a collection of object locations or guidance instructions. Architecture **120** includes TOA/location machine learning system **121** and a non-machine learning system **122**. The TOA/location machine learning system inputs the TOAs of the look and outputs a collection of object locations. The non-machine learning system inputs a collection of object locations and outputs a guidance instruction. The non-machine learning system may be a conventional system used in flight control to determine guidance instructions based on object locations. Architecture **130** includes a TOA/location machine learning system **131** and location/instruction machine learning system **132**. The TOA/location machine learning system **131** operates in the same manner as TOA/location machine learning system **121**. The location/instruction machine learning system inputs a collection of locations and outputs a guidance instruction. The location/instruction machine learning system may be employed independently of a machine learning system that outputs collections of object locations such as a conventional flight system that inputs signals and outputs object locations. Architecture **140** includes a TOA/instruction machine learning system. The TOA/instruction machine learning system inputs TOAs and outputs a guidance instruction.

[0027] The computing devices and systems on which the SGML system may be implemented may include a central processing unit, input devices, output devices (e.g., display devices and speakers), storage devices (e.g., memory and disk drives), network interfaces, graphical processing units, tensor processing units, field programmable gate arrays (FPGAs), accelerometers, cellular radio link interfaces, global positioning system devices, and so on. The input devices may include sensors, keyboards, pointing devices, touch screens, gesture recognition devices (e.g., for air gestures), head and eye tracking devices, microphones for voice recognition, and so on. The computing devices may

include desktop computers, laptops, tablets, e-readers, personal digital assistants, smartphones, gaming devices, servers, and computer systems, such as massively parallel systems. The computing devices may access computer-readable media that include computer-readable storage media and data transmission media. The computer-readable storage media are tangible storage means that do not include a transitory, propagating signal. Examples of computer-readable storage media include memory such as primary memory, cache memory, and secondary memory (e.g., DVD) and include other storage means. The computer-readable storage media may have recorded upon or may be encoded with computer-executable instructions or logic that implements the SGML system. The data transmission media is used for transmitting data via transitory, propagating signals or carrier waves (e.g., electromagnetism) via a wired or wireless connection.

[0028] The SGML system may be described in the general context of computer-executable instructions, such as program modules and components, executed by one or more computers, processors, or other devices. Generally, program modules or components include routines, programs, objects, data structures, and so on that perform particular tasks or implement particular data types. Typically, the functionality of the program modules may be combined or distributed as desired in various embodiments. Aspects of the system may be implemented in hardware using, for example, an application-specific integrated circuit (“ASIC”).

[0029] FIG. 2 is a flow diagram that illustrates the processing of a train TOA/location ML component of the SGML system in some embodiments. A train TOA/location ML component **200** trains a machine learning algorithm using training data that includes TOAs of looks and corresponding collections of objects as training data. The training data may be generated based on the simulation of the TOAs given object locations or given actual TOAs and object locations of actual objects collected as a platform moves relative to the actual objects. In block **201**, the component selects the next collection of object locations. In decision block **202**, if all the collections have already been selected, then the component continues at block **207**, else the component continues at block **203**. In block **203**, the component selects the next pair of transmitters and receivers of the architecture. In decision block **204**, if all the pairs of transmitters and receivers for the architecture have already been selected, then the component continues at block **206**, else the component continues at block **205**. In block **205**, the component calculates the TOAs for the receiver of the selected pair based on return signals from each object location and then loops to block **203** to select the next pair of transmitters and receivers. In block **206**, the component adds the TOAs and a collection of objects locations to the training data and loops to block **201** to select the next collection of object locations. In block **207**, the component trains the machine learning algorithm using the training data and completes. The SGML system may train a TOA/instruction ML system in a similar manner except that in block **206**, TOAs and instructions are the training data. The SGML system trains a location/instruction machine learning system using training data that includes collections of object locations and corresponding guidance instructions. The object locations may be simulated object locations or actual object locations.

[0030] FIG. 3 is a block diagram that illustrates the components of the reinforcement learning component of the SGML system in some embodiments. The reinforcement learning component 300 includes an agent system 310 and an environment system 320. The agent system includes a train ML component 311, a weights store 312, and a TOAs/predicted object location store 313. The environment system includes a retrieve TOAs component 321, a calculate reward component 322, and TOAs/actual object location store 323, and reward store 324. The agent system inputs TOAs of looks and rewards provided by the environment system. The train ML component of the agent system then adjusts weights (stored in the weights store) associated with the machine learning algorithm based on the rewards and TOAs and predicted object locations. The agent system then applies the weights to the TOAs to generate predicted object locations, stores the TOAs and predicted object locations, and provides the predicted object locations to the environment system. Upon receiving predicted object locations, the retrieve TOAs component retrieves from the TOA/actual object location store the TOAs of a look and provides those TOAs as current TOAs to the agent. The calculate reward system inputs the predicted object locations and calculates a reward based on similarity between the predicted object locations and the actual object locations corresponding to the prior TOAs provided to the agent component. The rewards component may store information relating to the rewards for use in calculating subsequent rewards.

[0031] FIG. 4 is a flow diagram that illustrates the processing of a retrieve TOAs component of the environment system in some embodiments. The retrieve TOAs component 400 loops sending TOAs to the agent system and receiving predicted object locations from the agent system. In block 401, the component initializes an interval count. In block 402, the component selects TOAs for the interval. In block 403, the component sends the selected TOAs to the agent system. In block 404, the component waits to receive the predicted object locations for the current interval and provides the predicted object locations to the calculate rewards component. In block 405, the component invokes the rewards component passing an indication of the predicted object locations for the current interval and an indication of the interval. In block 406, the component increments the interval and loops to block 402 to selected TOAs for the now current interval.

[0032] FIG. 5 is a flow diagram that illustrates the processing of a calculate rewards component of the environment system in some embodiments. The calculate rewards component 500 is passed the predicted object locations for the current interval and an indication of the current interval and calculates a reward. In block 501, the component stores the predicted object location. In block 502, the component calculates a reward based on the predicted object locations and the actual object locations for the current time interval and, in some embodiments, earlier time intervals. In decision block 503, if the machine learning algorithm is trained, then the component completes the processing of the reinforcement learning system. In block 504, the component sends the reward for the next interval to the agent system and then completes.

[0033] FIG. 6 illustrates the processing of a train ML component of the agent system in some embodiments. The train ML component 600 receives TOAs, adjusts weights, and generates predicted object locations. In block 601, the

component initializes the interval. In block 602, the component receives TOAs and a reward for the current interval. In blocks 603, the component stores the TOAs and the reward. In block 604, the component adjusts the machine learning weights based on the reward. In block 605, the component applies the machine learning algorithm to the TOAs to generate predicted object locations. In block 606, the component sends the predicted object locations to the environment system. In block 607, the component increments the interval, loops to block 602, and waits to receive the TOAs and reward for the interval.

[0034] The following paragraphs describe various embodiments of aspects of the SGML system. An implementation of the SGML system may employ any combination of the embodiments. The processing described below may be performed by a computing device with a processor that executes computer-executable instructions stored on a computer-readable storage medium that implements the SGML system.

[0035] In some embodiments, a method performed by one or more computing systems for generating a machine learning system to generate guidance information based on locations of objects is provided. The method accesses training data that includes for a plurality of looks, training TOA information of the look and guidance information for the look. The guidance information is based on a training collection of object locations. The TOAs of a look represent, for each object location of a training collection of object locations, times between signals transmitted by transmitters and return signals received by receivers. The return signals represent signals reflected from an object at the object location. The method trains a machine learning system using the training data. The machine learning system inputs TOA information and outputs guidance information. In some embodiments, the training TOA information represents the TOAs of a look and the guidance information is a collection of object locations for each look. In some embodiments, the training TOA information represents a training collection of object locations corresponding to TOAs of a look and the guidance information is a guidance instruction. In some embodiments, the training TOA information represents TOAs of a look and the guidance information is a guidance instruction. In some embodiments, the object locations in a training collection of object locations represent locations relative to an array of transmitters and receivers. In some embodiments, a guidance system of a platform with transmitters and receivers employs the machine learning system to input TOA information as the platform travels and output guidance information. In some embodiments, the guidance system determines a guidance instruction for the platform based on the TOA information. In some embodiments, the platform is a component of a robot control system. In some embodiments, the platform is a satellite and the object locations are locations of objects in space. In some embodiments, the platform is an unmanned vehicle. In some embodiments, the training TOA information is calculated during movement of a platform with transmitters and receivers through a volume of objects and object locations of the training collections are identified by an actual object location sensor. In some embodiments, the actual object location sensor is mounted on the platform. In some embodiments, the actual object location sensor is external to the platform. In some embodiments, the method generates TOA information by simulating time-of-arrivals for training collections of

object locations. In some embodiments, the machine learning system is based on reinforcement learning. In some embodiments, the machine learning system is based on supervised learning.

[0036] In some embodiments, a method performed by one or more computing systems to guide movement of a platform. The method, for each of a plurality of intervals, receives TOA information derived from TOAs determined based on times between signals transmitted by transmitters and return signals received by receivers. A return signal is reflected from an observed object at an object location. The method determines guidance information by applying a machine learning system that inputs TOA information and outputs guidance information. The machine learning system having been trained using training data that includes TOA information and guidance information. In some embodiments, the TOA information is the TOAs of a look and the guidance information is a collection of object locations for each look. In some embodiments, the TOA information is a collection of object locations corresponding to TOAs of a look and the guidance information is a guidance instruction. In some embodiments, the TOA information is TOAs of a look and the guidance information is a guidance instruction. In some embodiments, the machine learning system includes a first machine learning system that inputs TOAs of a look and outputs a collection of object locations and a second machine learning system that inputs a collection of object locations and output a guidance instruction. In some embodiments, the platform is a component of a robot control system. In some embodiments, the platform is a satellite and the object locations are locations of objects in space. In some embodiments, the platform is an unmanned vehicle. In some embodiments, the method guides the platform based on the guidance instructions.

[0037] In some embodiments, a method performed by one or more computing system for evaluating an architecture of a sensor array for a platform is provided. The method accesses a plurality of architectures of sensor arrays. An architecture specifies number and positions of transmitters and receivers of the sensor array. For each of a plurality of architectures, the method generates an architecture metric. To generate an architecture metric, the method, for each of a plurality of TOAs of looks and an evaluation collection of object locations, applies a machine learning system that inputs the TOAs of the look and generates an estimated collection of object locations and generates a metric based on similarity of object locations of the evaluation collection and the estimated collection, where the machine learning system has been trained using training data that includes TOAs of looks and for each look an evaluation collection of object locations. The method then generates an architecture metric based on the metrics generated for the architecture. In some embodiments, the architecture further specifies curvature of a sensor array. In some embodiments, the architecture metric is further based on size or weight of a sensor array. In some embodiments, the sensor array is a phased sensor array. In some embodiments, an architecture further specifies a beam pattern of a transmitter.

[0038] In some embodiments, one or more computing systems to guide movement of a platform are provided. The one or more computing systems include one or more computer-readable storage mediums for storing computer-executable instructions for controlling the one or more computing systems and one or more processors for executing the

computer-executable instructions stored in the one or more computer-readable storage mediums. The instructions, for each of a plurality of intervals, receive TOA information derived from TOAs determined based on times between signals transmitted by transmitters and return signals received by receivers and generate guidance information by applying a machine learning system that inputs TOA information and outputs guidance information. The machine learning system is trained using training data that includes TOA information and guidance information. In some embodiments, the TOA information is the TOAs of a look and the guidance information is a collection of object locations for each look. In some embodiments, the TOA information is a collection of object locations corresponding to TOAs of a look and the guidance information is a guidance instruction. In some embodiments, the TOA information is TOAs of a look and the guidance information is a guidance instruction. In some embodiments, the machine learning system includes a first machine learning system that inputs TOAs of a look and outputs a collection of object locations and a second machine learning system that inputs a collection of object locations and output a guidance instruction. In some embodiments, the platform is a component of a robot control system. In some embodiments, the platform is a satellite and the object locations are locations of objects in space. In some embodiments, the platform is an unmanned vehicle. In some embodiments, the instructions further guide the platform based on the guidance information.

[0039] In some embodiments, one or more computing systems for evaluating an architecture of a sensor array for a platform are provided. The one or more computing systems include one or more computer-readable storage mediums for storing computer-executable instructions for controlling the one or more computing systems and one or more processors for executing the computer-executable instructions stored in the one or more computer-readable storage mediums. The instructions access an architecture that specifies number and positions of transmitters and receivers of the sensor array. For each of a plurality of TOAs of looks and an evaluation collection of object locations, the instructions apply a machine learning system that inputs the TOAs of the look and generates an estimated collection of object locations. The instructions also generate a metric based on similarity of object locations of the evaluation collection and the estimated collection. The instructions then generate an architecture metric based on the metrics generated for the architecture. In some embodiments, the machine learning system is trained using training data that includes TOAs of looks and for each look an evaluation collection of object locations. In some embodiments, the architecture further specifies curvature of a sensor array or a beam pattern of a transmitter. In some embodiments, the architecture metric is further based on size or weight of a sensor array. In some embodiments, the sensor array is a phased sensor array.

[0040] In some embodiments, one or more computing systems for generating a machine learning system to generate guidance information based on locations of objects are provided. The one or more computing systems include one or more computer-readable storage mediums for storing computer-executable instructions for controlling the one or more computing systems and one or more processors for executing the computer-executable instructions stored in the one or more computer-readable storage mediums. The

instructions access training data that includes, for each of a plurality of looks, training TOA information of the look and guidance information for the look. The instructions train a machine learning system using the training data wherein the machine learning system inputs TOA information and outputs guidance information. In some embodiments, the guidance information is based on a training collection of object locations. The TOAs of a look represent, for each object location of a training collection of object locations, times between signals transmitted by transmitters and return signals received by receivers. The return signals represent signals reflected from an object at the object location. In some embodiments, the training TOA information is generated during movement of a platform with transmitters and receivers through a volume of objects and object locations of the training collections are identified by an actual object location sensor. In some embodiments, the TOA information represents the TOAs of a look and the guidance information is a collection of object locations for each look. In some embodiments, the training TOA information represents a training collection of object locations corresponding to TOAs of a look and the guidance information is a guidance instruction. In some embodiments, the training TOA information represents TOAs of a look and the guidance information is a guidance instruction. In some embodiments,

[0041] The following paragraphs describe various embodiments of aspects of the SGML system. An implementation of the SGML system may employ any combination of the embodiments. The processing of the methods described below may be performed by a computing device with a processor that executes computer-executable instructions stored on a computer-readable storage medium that implements the SGML system.

1. A method performed by one or more computing systems to guide movement of a platform, the method comprising, for each of a plurality of intervals,

receiving time-of-arrival (“TOA”) information derived from TOAs determined based on times between signals transmitted by transmitters and return signals received by receivers wherein a return signal is reflected from an observed object at an object location; and

determining guidance information by applying a machine learning system that inputs TOA information and outputs guidance information, the machine learning system being trained using training data that includes TOA information and guidance information.

2. The method of claim 1 wherein the TOA information is the TOAs of a look and the guidance information is a collection of object locations for each look.

3. The method of claim 1 wherein the TOA information is a collection of object locations corresponding to TOAs of a look and the guidance information is a guidance instruction.

4. The method of claim 1 wherein the TOA information is TOAs of a look and the guidance information is a guidance instruction.

5. The method of claim 1 wherein the machine learning system includes a first machine learning system that inputs TOAs of a look and outputs a collection of object locations and a second machine learning system that inputs a collection of object locations and output a guidance instruction.

6. The method of claim 1 wherein the platform is a component of a robot control system.

7. The method of claim 1 wherein the platform is a satellite and the object locations are locations of objects in space.

8. The method of claim 1 wherein the platform is an unmanned vehicle.

9. The method of claim 1 further comprising guiding the platform based on the guidance information.

10. A method performed by one or more computing systems for evaluating an architecture of a sensor array for a platform, the method comprising:

accessing a plurality of architectures of sensor arrays, an architecture specifying number and positions of transmitters and receivers of the sensor array; and

for each of plurality of architectures,

for each of a plurality of time-of-arrivals (“TOAs”) of looks and an evaluation collection of object locations,

applying a machine learning system that inputs the TOAs of the look and generates an estimated collection of object locations, the machine learning system trained using training data that includes TOAs of looks and for each look an evaluation collection of object locations; and

generating a metric based on similarity of object locations of the evaluation collection and the estimated collection; and

generating an architecture metric based on the metrics generated for the architecture.

11. The method of claim 10 wherein the architecture further specifies curvature of a sensor array.

12. The method of claim 10 wherein the architecture metric is further based on size or weight of a sensor array.

13. The method of claim 10 wherein the sensor array is a phased sensor array.

14. The method of claim 10 wherein an architecture further specifies a beam pattern of a transmitter.

15. One or more computing systems to guide movement of a platform, the one or more computing systems comprising:

one or more computer-readable storage mediums for storing computer-executable instructions for controlling the one or more computing systems to, for each of a plurality of intervals:

receive time-of-arrival (“TOA”) information derived from TOAs determined based on times between signals transmitted by transmitters and return signals received by receivers; and

generate guidance information by applying a machine learning system that inputs TOA information and outputs guidance information, the machine learning system being trained using training data that includes TOA information and guidance information; and

one or more processors for executing the computer-executable instructions stored in the one or more computer-readable storage mediums.

16. The one or more computing systems of claim 15 wherein the TOA information is the TOAs of a look and the guidance information is a collection of object locations for each look.

17. The one or more computing systems of claim 15 wherein the TOA information is a collection of object locations corresponding to TOAs of a look and the guidance information is a guidance instruction.

18. The one or more computing systems of claim **15** wherein the TOA information is TOAs of a look and the guidance information is a guidance instruction.

19. The one or more computing systems of claim **15** wherein the machine learning system includes a first machine learning system that inputs TOAs of a look and outputs a collection of object locations and a second machine learning system that inputs a collection of object locations and output a guidance instruction.

20. The one or more computing systems of claim **15** wherein the platform is a component of a robot control system.

21. The one or more computing systems of claim **15** wherein the platform is a satellite and the object locations are locations of objects in space.

22. The one or more computing systems of claim **15** wherein the platform is an unmanned vehicle.

23. The one or more computing systems of claim **15** wherein the instructions further guide the platform based on the guidance information.

24. One or more computing systems for evaluating an architecture of a sensor array for a platform, the one or more computing system comprising:

one or more computer-readable storage mediums for storing computer-executable instructions for controlling the one or more computing systems to:

access an architecture that specifies number and positions of transmitters and receivers of the sensor array; and

for each of a plurality of time-of-arrivals (“TOAs”) of looks and an evaluation collection of object locations,

apply a machine learning system that inputs the TOAs of the look and generates an estimated collection of object locations; and

generate a metric based on similarity of object locations of the evaluation collection and the estimated collection; and

generate an architecture metric based on the metrics generated for the architecture; and

one or more processors for executing the computer-executable instructions stored in the one or more computer-readable storage mediums.

25. The one or more computing systems of claim **24** wherein the machine learning system is trained using training data that includes TOAs of looks and for each look an evaluation collection of object locations.

26. The one or more computing systems of claim **24** wherein the architecture further specifies curvature of a sensor array or a beam pattern of a transmitter.

27. The one or more computing systems of claim **24** wherein the architecture metric is further based on size or weight of a sensor array.

28. The one or more computing systems of claim **24** wherein the sensor array is a phased sensor array.

29. One or more computing systems for generating a machine learning system to generate guidance information based on locations of objects, the one or more computing systems comprising:

one or more computer-readable storage mediums for storing computer-executable instructions for controlling the one or more computing systems to:

access training data that includes, for each of a plurality of looks, training time-of-arrival (“TOA”) information of the look and guidance information for the look; and

train a machine learning system using the training data wherein the machine learning system inputs TOA information and outputs guidance information; and

one or more processors for executing the computer-executable instructions stored in the one or more computer-readable storage mediums.

30. The one or more computing systems of claim **29** wherein the guidance information is based on a training collection of object locations and wherein the TOAs of a look represent, for each object location of a training collection of object locations, times between signals transmitted by transmitters and return signals received by receivers wherein the return signals represent signals reflected from an object at the object location.

31. The one or more computing systems of claim **30** wherein the training TOA information is generated during movement of a platform with transmitters and receivers through a volume of objects and object locations of the training collections are identified by an actual object location sensor.

32. The one or more computing systems of claim **29** wherein the TOA information represents the TOAs of a look and the guidance information is a collection of object locations for each look.

33. The one or more computing systems of claim **29** wherein the training TOA information represents a training collection of object locations corresponding to TOAs of a look and the guidance information is a guidance instruction.

34. The one or more computing systems of claim **29** wherein the training TOA information represents TOAs of a look and the guidance information is a guidance instruction.

* * * * *