

(51) International Patent Classification:
G06F 19/00 (2011.01)(21) International Application Number:
PCT/US2014/039239(22) International Filing Date:
22 May 2014 (22.05.2014)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/826,387 22 May 2013 (22.05.2013) US(71) Applicants: **NEURALA, INC.** [US/US]; The Photonics Center, 8 St. Mary's Street, Suite 613, Boston, Massachusetts 02215 (US). **TRUSTEES OF BOSTON UNIVERSITY** [US/US]; One Silber Way, Boston, Massachusetts 02215 (US).(72) Inventors: **GORSHECHNIKOV, Anatoly**; 15 Jefferson Street, Newton, Massachusetts 02458 (US). **VERSACE, Massimiliano**; 846 East 3rd Street, Boston, Massachusetts 02127 (US). **BARNES, Tim**; 1305 E. Howell Street, Apt. 101, Seattle, Washington 98122 (US).(74) Agents: **TEJA, Joseph** et al.; Cooley LLP, 1299 Pennsylvania Ave., Suite 700, Washington, District of Columbia 20004 (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, 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, KZ, LA, LC, LK, LR, LS, LT, 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, 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, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: METHODS AND APPARATUS FOR EARLY SENSORY INTEGRATION AND ROBUST ACQUISITION OF REAL WORLD KNOWLEDGE




Figure 20

(57) Abstract: The systems and methods disclosed herein include a path integration system that calculates optic flow, infers angular velocity from the flow field, and incorporates this velocity estimate into heading calculations. The resulting system fuses heading estimates from accelerometers, 5 gyroscopes, engine torques, and optic flow to determine self-localization. The system also includes a motivational system that implements a reward drive, both positive and negative, into the system. In some implementations, the drives can include: a) a curiosity drive that encourages exploration of new areas, b) a resource drive that attracts the agent towards the re-charging base when the battery is low, and c) a mineral reward drive that attracts the agent 10 towards previously explored scientific targets.

**Methods and Apparatus for Early Sensory Integration and Robust Acquisition of Real
World Knowledge**

CROSS-REFERENCE TO RELATED PATENT APPLICATION

[0001] This application claims priority, under 35 U.S.C. § 119(e), from U.S. Application No. 61/826,387, filed May 22, 2013, and entitled "METHODS AND APPARATUS FOR PATH INTEGRATION AND NEURAL ACQUISITION OF LOCATION ESTIMATE," which application is hereby incorporated herein by reference in its entirety.

GOVERNMENT SUPPORT

[0002] This invention was made with government support under Contract No. NNX12CG32P awarded by NASA. The government has certain rights in the invention.

BACKGROUND

[0003] To date there has been little success in development of robotic intelligent behaviors without reliance on resource-intensive active sensors such as but not limited to laser range finders, sonars, etc. Furthermore, these sensors are often tailored and used exclusively for individual tasks. At the same time, robots do have other multi- purpose passive sensors, but the precision of knowledge based on these sensors leaves much to be desired. The lack of reliable cross-task sensory system is partially due to sensory noise and partially to an approach of independent, stove-piped processing of individual sensory streams.

SUMMARY

[0004] The system and methods described herein is based on a combination of an arbitrary number of sensory inputs correcting each other by compensating weaknesses of one sensor with the strengths of other sensors in the process of early fusion, then processing these sensory inputs through a redundant and robust neural-like system. An exemplary system may be used for a variety of applications, including but not limited to spatial navigation, visual object segmentation and recognition, or robotic attention.

[0005] In some implementations, the system may generate a noise-tolerant distributed data representation model via generating a set of data cells collectively representing a sensory

input data point in the range of measurement for a particular type of data, and may use the data to define a spatial or temporal resolution of the set of data cells (e.g., a data representation model) from the data range of the incoming sensory information and the number of cells in the set. A spatial or temporal scale for the representation may be determined using the range of the incoming sensory information, and may be used to generate a set of graph connections among the set of data cells with a first set of weights (in some implementations, these connections are designed to ensure continuity between representation of points in continuous space and some distance measure so that data cell representation reflects the distance between input data points).

[0006] It should be appreciated that all combinations of the foregoing concepts and additional concepts discussed in greater detail below (provided such concepts are not mutually inconsistent) are contemplated as being part of the inventive subject matter disclosed herein. In particular, all combinations of claimed subject matter appearing at the end of this disclosure are contemplated as being part of the inventive subject matter disclosed herein. It should also be appreciated that terminology explicitly employed herein that also may appear in any disclosure incorporated by reference should be accorded a meaning most consistent with the particular concepts disclosed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The skilled artisan will understand that the drawings primarily are for illustrative purposes and are not intended to limit the scope of the inventive subject matter described herein. The drawings are not necessarily to scale; in some instances, various aspects of the inventive subject matter disclosed herein may be shown exaggerated or enlarged in the drawings to facilitate an understanding of different features. In the drawings, like reference characters generally refer to like features (e.g., functionally similar and/or structurally similar elements).

[0008] Figure 1A shows experimental recordings from a visual cortex that illustrate continuity of sensory data and its representation in the brain.

[0009] Figure 1B shows an example tuning curve for the experimental data shown in Figure 1A.

[0010] Figure 1C shows an example of neuronal representation with narrow on-center (local excitation) and wide off-surround (global inhibition, this is the extreme case when the surround includes all other cells).

[0011] Figure 1D shows an example of a neuronal representation of an animal's heading.

[0012] Figure 2 is an illustration of beats of low frequency created by overlapping two higher frequencies. This can be used in space or time to encode intervals in wide ranges using only a limited set of base frequencies. Vernier scale and hyper-acuity is related to the inverse application of this principle.

[0013] Figure 3 illustrates the process of segmenting objects from their "background", by adaptively combining in a feedforward/feedback process various sources of potentially imprecise/noisy visual and non-visual information to estimate what is figure and what is ground, and create a form-fitting (or object-fitting) mask to enable segmentation of the object of interest.

[0014] Figure 4 depicts a biological path integration system. On the right, Figure 4 names the mammalian brain areas subserving path integration. On the left it describes core modeling work that was combined to create an example PInNAcLE implementation.

[0015] Figure 5 shows the results of PInNAcLE system. Individual motor, vestibular and optic flow (visual) based estimates are plotted without feedback correction described below and show consistent overshoot for motor and undershoot for optic flow. Combined estimate follows the true heading (2D and 3D headings) much closer.

[0016] Figure 6 shows an example flow chart for IMU on the fly calibration algorithm using flow field information.

[0017] Figure 7 shows obstacle map generated by the navigation model presented herein and how it affects path planning search algorithm.

[0018] Figure 8 shows assembly diagram of robotic sensors used in the example embodiment

[0019] Figure 9 shows software components of example robotic embodiment.

[0020] Figure 10 shows an experimental paradigm used to test the navigation model presented herein.

[0021] Figure 11 shows an example trajectory of a simulated robot in a virtual environment.

[0022] Figure 12 shows a map view of a trajectory from Figure 11 as it builds incrementally through time.

- [0023] Figure 13 plots a battery level during the exploration experiment.
- [0024] Figure 14 shows results of a model test in a robotic setup described herein.
- [0025] Figure 15 provides an overview of processing streams in a mammalian visual system.
- [0026] Figure 16 illustrates the result of OpenEye model recognizing rock types in a virtual environment
- [0027] Figure 17 depicts Virt-U interface between neural computation and 3D virtual engine.
- [0028] Figure 18 provides an example structure of Virt-U classes.
- [0029] Figure 19 shows a structure of a proxy layer designed to interface Virt-U with Cog Ex Machina.
- [0030] Figure 20 shows an exemplary navigation system architecture.
- [0031] Figure 21 shows a flow chart illustrating an exemplary embodiment of path navigation.

DETAILED DESCRIPTION

[0032] This disclosure is subdivided in several sections. The first section describes the basic principles of the biological sensory processing and claims how they can be applied to improve robotic and other engineering applications. The next section describes an example robotic navigational system that is based at least on the principle of the first section and is used to test the applicability of the presented approach. The disclosure presents an example of an artificial visual system that illustrates applicability of the principles in the first section to a completely different domain thus implying the generality of these principles. It also discloses a software layer that allows communication between neural models based on the principles of the first section and built in various software packages with robotic or virtual sensory motor scenarios that allow testing and further development of the above principles.

[0033] Biological Sensory Processing and Engineering Practices

[0034] Sensory information processing in an animal nervous system model differs from conventional engineering approaches in several respects: data representation, horizontal interactions between processing streams within and across modalities, and feedback from higher level results to early processing stages within and across processing streams. Neural-

like data representation, as well as neural horizontal and feedback interactions deliver advantages in engineering of artificial information processing systems, as detailed below.

[0035] *Data Representation*

[0036] In order to save resources, engineers try to minimize the representation of the sensory data. Whether it is luminance of a visual or radar pixel, acceleration, or sound pitch, conventional systems tend to represent it as a single floating point number with possible duplication or even triplication for the systems where robustness is important. This does not prevent the data from experiencing abrupt changes due to noise. Consequences of this include disregarding a temporal continuity of the physical world around us and (without additional processing) not establishing a relationship between nearby values, thus disregarding a spatial continuity of the world.

[0037] Figures 1A–1D illustrate a neural representation which uses ensembles of cells to represent every data point, which sacrifices the storage space for the sake of robustness and preserving the continuities of the world. Specifically, Figure 1A illustrates this principle for a cell population that codes the orientation of a bar in a visual field of an animal from experimental recordings. In visual cortex a set of cells 100 that represent orientations exist for each point in the visual space, and if a bar of a certain orientation 102 appears at this point not only the cell with the corresponding orientation becomes active 104 but so do cells with nearby orientations 106, the closer their preferred orientations are to the one presented, the higher is the cell activation (so called tuning curve of the cell). As a result, for each data point there is a distributed bell-shaped representation 108 (shown in Figure 1B) that is centered at the precise value of the preferred orientation and gradually decaying as the nearby cells tuning moves away from this value. This distributed representation is ubiquitous in the nervous system, and is used to represent headings in parasubiculum, places in the hippocampus, tones in auditory cortex, and many other sensory readings or processed data. Figure 1C illustrates how the neurons may be connected to achieve such a representation using local excitation 110 and global or long range inhibition 112. Figure 1D shows the activities of the resulting cells showing bell-shaped profile for a particular input (114).

[0038] This representation bears several advantages. First, it is more resistant to damage to the processing hardware. For example, the head direction cells that represent the animal's heading may have tuning curves about 90 degrees wide, which suggests that losing the heading signal completely involves losing at least a quarter of head direction cells. Even then, losing a quarter of the heading cells may affect only a quarter of possible headings.

[0039] The second advantage of the neuronal representation is preservation of the continuity of the world. The cells in the representation are not standalone memory elements, but a connected ensemble. Usually this connectivity follows some version of on-center off-surround (e.g., Mexican hat) profile, where the cells excite nearby (in terms of data values or tuning curves) cells and inhibit cells that are further away. This connectivity creates an attractor dynamics where it is easy to shift from the current data value to nearby values, but where it is much harder to shift to further away values. This reflects a natural continuity of data: if a person is facing East at one moment, it is virtually impossible for that person to face West in the next moment, and much more likely that the person's heading would still be very close to East the next time the person's heading was requested. From an engineering standpoint, this implementation may also work as a low-pass filter, which can be a first cascade in the sensory processing that serves to reduce or eliminate the high frequency jitter most sensors produce. Therefore, for the fair comparison of storage and processing requirements of state of the art engineering system and biological system we should compare the initial data representation of the latter with a combined data representation and low-pass filter of the former. Using attractor on-center off-surround dynamics may increase the robustness of the representation for continuous values.

[0040] Furthermore, biological systems often utilize multiple scales of on-center off-surround connectivity (corresponding to low-pass filter frequencies) through different lengths of on-center and off-surround connections, resulting in multiple scales of sensory representation. This allows the brain not only to process information at an appropriate scale (e.g., one does not need to know the distance between Boston and New York in feet in order to plan a trip) but also to dynamically switch between scales (e.g., gradually increase one's resolution as one approaches one's destination in New York). Due to a principle of beats phenomenon that occur when one combines two nearby frequencies 202 and 204 and produces a signal with a significantly lower frequency 206 (e.g., as shown in Figure 2), a limited set of base scales can allow for representing a wide range of data (e.g., up to least common multiple of the scales) with high precision (e.g., at least as high as the smallest scale) without a necessity to involve large numbers of cells in the representation. Furthermore, interactions between multiple scales can lead to phenomenon of hyper-acuity, when the information derived from the representation has higher resolution than the original sensory data (this is based on the same principle as Vernier scale used in engineering for

precise measurements). This principle can be generalized, and can be applicable as long as there is either spatial or temporal extent in the input signal.

[0041] To summarize, using neuronal-like data representation on every processing stage allows robustness of the system, built-in low pass filtering, increase in data precision beyond the precision of a sensor, preservation of spatial and temporal continuity of the signal, and easily scalable ability to handle wide ranges of data with high precision. Bell shaped and often Gaussian profiles of neuronal activity are also very convenient for combining signals from different inputs as discussed in the next section.

[0042] *Horizontal interactions between processing streams within and across modalities*

[0043] A conventional solution for processing multiple sensory streams is to stove-pipe several algorithms, each being finely tuned for a specific task, and combine their outputs for further analysis and guidance of actions. For example, visually estimated self motion is often used to correct odometry-based motion estimates, but it is usually done in discrete steps, rather than continuously, and by resetting the odometry data, rather than fusion of two sensory streams. Each individual algorithm is a result of a long effort to achieve maximal performance for a certain task; however these algorithms in aggregate rarely allow meaningful use of the results of intermediate processing stages. Furthermore, combining the algorithms involves reducing the interdependencies between processing streams to keep the system tractable, leading to little information about how the processing streams relate to each other.

[0044] An approach based on biological principles may involve improving processing, within each stream, using the early and intermediate results of the stream along with other streams. Such an approach may be used, for example, with image segmentation by the visual system. The visual system may be configured to allow a user to build an accurate representation of the world, including relative positions of all objects and how they move with respect to an entity (e.g., a robot). A conventional approach would be to select objects of interest, fine tune detection algorithms for these objects, and use outputs of these algorithms to feed motion detection. A human brain's approach is quite different; it extracts the information from the visual input in parallel, and uses different feature sets 300 of this information to help processing of other features 302 as early as it is useful and long before the objects are detected and identified 304 (e.g., see Figure 3). In the example implementation of an OpenEye system discussed below, determining the boundary of an object can be done using a change in luminance, color, texture, or other surface property, as well as depth

information that can come from binocular disparity or from accommodation of the eye, and discontinuities in the motion signals like optic flow. All of these cues can be combined on the early stages of processing.

[0045] Furthermore, at this stage the system does not need to know what the objects are, and how many are there; it may be completely agnostic towards identifying objects. As a result, the output of these interactions – a video stream segmented into areas that have similar properties and move together - can be used for a plurality of types of object detection and recognition, thus unifying a significant chunk of processing and making it scalable for any arbitrary number of objects. Additionally, the object-aware part of the system may become independent of the number of sensors used to produce visual segmentation, and can be reused, for example, on different robots with different sensory compositions. Since motion, depth, and different surface properties are processed at different resolution (e.g. motion signals may have coarser resolution than surface property signals), the system may also utilize all advantages that come with multiple scales of neuronal representations outlined in the previous section. Various modalities and even data extraction methods of sensory processing often provide incomplete and imprecise results. Early fusion of the outputs of these algorithms and modalities makes it possible to feed more complete and precise data to computationally expensive high-order processing algorithms.

[0046] Another advantage of early interactions between sensory processing streams may be a shortened response time. For example, for the human brain, interaural differences in sound arrival timings allow the brain to detect the source of sound in space. This detection happens early in the auditory stream in the inferior part of colliculus before the information reaches even the primary auditory cortex. The superior part of colliculus controls eye movements, and a direct link from the auditory stream to visual stream allows humans to move their eyes towards the source of sound long before they can recognize what this sound is. Similarly, a quick estimate of optic flow by a vehicle can be used to determine that there is an object on a collision course with this vehicle and trigger an avoiding maneuver before the nature of the colliding object is processed. The linkage between early sensory processing stages and fast motor responses may reduce response times for important inputs.

[0047] In some implementations (referring to Figure 21), a system may generate a set of data cells 2100, the set of data cells representing sensory input data received from a group of sensors. The sensor may use the set of data cells to both determine a spatial and/or temporal resolution 2102 of the set of data cells, and to determine a spatial and/or temporal scale 2104

for the set of data cells and the range of available data. In some implementations the resolution may refer to the size of the data set with respect to the range of data. The system may then define a set of weighted connections from the data cells 2106, to the combined representation 2116. Fill in the sensory data 2110. Using the apriori information about each sensor precision, the system may alter the significance of the weights from each set of data points in 2112, e.g., based on the precision of the source of the data represented by each set of weights, and/or the like. The system may weight the sensory input data from each of the sensors 2114 based on the altered sets of weights, and may combine the weighted sensory input data 2116. The weighted sensory data may be fed back into the system in order to update data representation in 2110 as in 2118. Iterate through the loop between 2110 and 2118 as long as new sensory data comes in.

[0048] Another illustration of across modality interactions is a Path Integration and Natural Acquisition of Location Estimate (PInNAcLE) system based on biological path integration illustrated in Figure 4. In some implementations, the PInNAcLE system transforms velocity data 400 recorded from the agent's sensors and actuators into the robot's world position. The basic element of the system is a ring attractor network. Here the inputs from three sensory modalities 402 are combined: motor outflow, which precedes the motion, but can be incorrect since the body might be restrained and the effectors might slip; IMU input, which is slightly delayed and has issues with detecting very slow motions (see Table 1); and optic flow, which is good for slow rotations but becomes increasingly inaccurate as the speed increases and cannot estimate linear motions without additional information.

Size\Normalize d input	0.0001	0.0002	0.0004	0.001	0.002	0.004	0.01	0.02	0.04
36								0.14	0.34
72							0.085	0.18	0.36
180		0.002	0.004	0.01	0.018	0.036	0.092	0.184	0.368
360	0.001	0.002	0.004	0.009	0.018	0.037	0.092	0.184	0.368
720	0.001	0.002	0.0035	0.009	0.0185	0.037	0.0925	0.1845	0.3685

Table 1 Calibration of a ring attractor from Figure 1. Left column shows the size of a network in cells, top row shows a normalized between 0 and 1 angular velocity input, values in the table show attractor rotation in degrees per second. Pink shaded cells show where the input was too small for the

attractor to respond, yellow shading shows where attractor performed at the edge of loosing precision due to small values of the input.

[0049] Each of these streams has advantages and disadvantages relative to others and a proper fusion of this information within the PInNAcLE system allows these streams to capitalize on the advantages of early sensor fusion and also to self-correct problems with different sensors. The resulting system can reliably track the vehicle position during extended periods of exploration.

[0050] PInNAcLE derives estimate of current heading from each of the three inputs and represents these estimates as shown in Figure 1B. The Gaussian shape of these signals simplifies their combination, as the sum of Gaussians that are close enough together (how to ensure this condition is discussed in the next section) is also a Gaussian. But simple sum may not suffice; in some implementations, the motor input may only be considered if it is backed up by either vestibular or optic flow, while the latter two may be weighted proportionally to the movement speed, with optic flow having more weight for slow speeds and vestibular having more weight for fast speeds. The resulting heading estimate 500 may take advantage of predictive power of motor outflow and high precision of vestibular input at high speeds and optic flow at low speeds (e.g., see Figure 5).

[0051] In some implementations, optic flow (e.g., the set of local motions perceived across the visual field at a particular time), provides a distributed and robust cue for navigation. Biological organisms use optic flow for navigation; for example, when human visual and vestibular senses are in conflict, the visual sense of self-motion can often times overpower the vestibular one. An optic flow system is included in the model (described below) in order to leverage this source of information as a complementary signal to vestibular and motor commands for path integration. The custom algorithm used to determine optic flow is a hybrid of two models.

[0052] Firstly, an initial optic flow field is calculated using monogenic filters, which have been useful for fast optic flow computation. These filters are computationally efficient (especially on a parallel GPU-enabled platform) and naturally illumination-invariant because they are based on a local image phase only. Because the signal is local, however, it is often noisy; this noise can be dampened by spatiotemporal smoothing.

[0053] Thus, a method of smoothing the monogenic signal was developed based on a motion integration algorithm. In some implementations the algorithm estimates agent velocity parameters from the optic flow field according to a pinhole camera model.

[0054] *Monogenic filters*

[0055] The optic flow field (e.g., a 2D vector field over an image) is first constructed by calculating a set of features over the image and then correlating those features over time. The optic flow system extracts local image phase information through a set of monogenic filters over a set of spatial scales s . Monogenic filters are based on Gabor wavelet filters, which balance some spatial and some frequency selectivity. They can be parameterized by their center spatial frequencies; the algorithm uses a particular set of center frequencies spaced geometrically according to

$$\mu_s = \mu_0 \cdot \gamma^{-s},$$

where $\mu_0 = 0.25$ is the highest-frequency filter scale, subsequent filters are scaled by a factor $\gamma = 1.3$, the number of spatial scales $n_s = 6$, and so the set of spatial scales $s \in \{0, \dots, 5\}$.

Monogenic filters are based on bandpass filters $G_{f_x f_y s}$. The general shape of one bandpass filter is given by a Gaussian function in polar coordinates, which produces a ring in image frequency space:

$$F_{f_x f_y}(\mu, \sigma) = \frac{1}{2\pi\sigma_r} \exp\left(\frac{(f_r - \mu)^2}{-2\sigma^2}\right),$$

where the radial component f_r is given by

$$f_r = \sqrt{f_x^2 + f_y^2}.$$

[0056] This ring can be re-parameterized exclusively by the spatial scale s , in which case the width of the ring scales with the center frequency μ_s :

$$G_{f_x f_y s} = F_{f_x f_y}(\mu_s, 0.85\mu_s).$$

[0057] This bandpass filter estimates the amplitude of local frequency at an image; the local angle and phase at that frequency is estimated with filters that select for horizontal or vertical frequency energy. These two real-valued filters can be combined into one complex filter $H_{f_x f_y}$ to save on computational resources:

$$H_{f_x f_y} = \frac{if_x - f_y}{f_r}.$$

[0058] To get the complete monogenic signal, the image I_{xyt} is transformed into ordinary frequency space with the Fourier transform $\mathcal{F}: (x, y) \rightarrow (f_x, f_y)$, multiplied by the set of bandpass filters, and then transformed back into the original image space.

$$\vec{a}_{xyt} = \begin{bmatrix} \mathcal{F}^{-1} \left[\mathcal{F}(I_{xyt}) G_{f_x f_y s} \right] \\ \Re(\mathcal{F}^{-1} \left[\mathcal{F}(I_{xyt}) G_{f_x f_y s} H_{f_x f_y} \right]) \\ \Im(\mathcal{F}^{-1} \left[\mathcal{F}(I_{xyt}) G_{f_x f_y s} H_{f_x f_y} \right]) \end{bmatrix} = \begin{bmatrix} a_{xyt}^{(1)} \\ a_{xyt}^{(2)} \\ a_{xyt}^{(3)} \end{bmatrix}.$$

[0059] Note that, since I_{xyt} and $G_{f_x f_y s}$ are both real-valued, and the resulting filtered signal is also real-valued ($\vec{a}_{xyt} \in \mathbb{R}^3$). Changing this monogenic signal into spherical coordinates can recover the three components of local image phase: amplitude, angle, and phase roll.

The difference in this monogenic signal over time produces an optic flow field. The monogenic filter signals are specifically compared between time frames for a change in local estimated phase, which is taken as an estimate of local image movement in the direction of the monogenic signal's estimated phase angle. The following equation, based on quaternion rotation, keeps the movement vector in the same monogenic space:

$$\vec{b}_{xyt} = \begin{bmatrix} \vec{a}_{xyt(t-1)} \cdot \vec{a}_{xyt} \\ a_{xyt(t-1)}^{(1)} a_{xyt}^{(2)} - a_{xyt}^{(1)} a_{xyt(t-1)}^{(2)} \\ a_{xyt(t-1)}^{(1)} a_{xyt}^{(3)} - a_{xyt}^{(1)} a_{xyt(t-1)}^{(3)} \end{bmatrix}.$$

Spatiotemporal smoothing

[0060] The monogenic optic flow signal is efficient and local, but this spatial locality can produce unwanted noise in the resulting optic flow field. This noise can be reduced through a process of spatiotemporal smoothing. The following method of smoothing the monogenic signal was loosely based on a model of biological motion integration. This motion integration process allows for some amount of local smoothing while retaining the computational efficiency of monogenic filters. More complicated smoothing algorithms may be useful for scenes with independently moving objects.

[0061] The main smoothing process, loosely identified with the primary visual cortex, V1, takes the input flow signal, averages that signal over space and time, and incorporates feedback that averages over spatial scales. The process of averaging over spatial scales

roughly corresponds to the functional role of a downstream motion-sensitive brain area, the middle temporal area (MT). This averaged value can be read by other parts of the model as the final optic flow signal.

[0062] *Egomotion estimation*

[0063] Finally, the optic flow signal is used to estimate the camera's rotation, and thus the rotation velocity of the agent. This is done by projecting the optic flow field onto a template field that linearly scales with camera rotation speed. The template field only depends on the pixel position (x, y) and the camera parameters, namely, the focal length α and the physical pixel size β . Note, however, that the pixel indices range from negative to positive rather than starting at 0 or 1; for example, for a 640×480 image, $x \in \{-319.5, -318.5, \dots, 319.5\}$ and $y \in \{-239.5, -238.5, \dots, 239.5\}$. The template field is given by

$$\vec{R}_{xyt} = \frac{-1}{\alpha} \begin{bmatrix} \alpha^2 + (\beta x)^2 \\ \beta^2 xy \end{bmatrix},$$

and the projection operation can be written as

$$r_t = \frac{\sum_{xy} \vec{R}_{xyt}^T \cdot \vec{v}_{xyt}}{\sum_{xy} \|\vec{R}_{xyt}\|^2}.$$

[0064] If a spherical camera model such as the one initially proposed were used instead, then the template field $\vec{R}_{xyt}$ would be independent of pixel position. This simplification would allow for less memory storage but offers little to no speed advantage over the current image plane method. This contrasts with the calculation of linear velocities of optic flow, which depend on an estimate of depth at every pixel that changes over *time* as well as space, requiring more computational resources and communication. Unfortunately, modeling a spherical camera in the virtual environment poses particular challenges because OpenGL is heavily optimized for graphics in Cartesian coordinates. A spherical camera projection would require that all straight lines between triangle vertices to be deformed into curved lines on the camera sensor. The results determined by simulation shown below show that the pinhole camera model is sufficiently fast and accurate for egomotion estimation in the virtual environment.

[0065] Inertial measuring units (IMUs) have three sources of error with decreasing magnitude: Static Bias, Dynamic Bias and Random White Noise. Static bias is reduced by recalibrating the IMU between runs when the vehicle is stationary. Dynamic bias occurs

during the mission while the vehicle is underway and if not addressed can lead to drift — where the error from the previous position estimate further impacts the next position estimate. The correction of this drift can happen independently for rotational IMUs (gyroscopes) and translational IMUs (accelerometers) from the segregated optic flow or other flow field for example sonar depth field through time. Full flow field can be separated into depth-invariant component that represents rotation of the vehicle and depth-dependent component that represents its translation. Periods of time when there is no rotational flow can be used to recalibrate rotational IMUs, and periods with no translational flow can be used to calibrate translational IMUs. Effectively, this is a natural continuation of the weighted approach described in the previous paragraph: when the perceived motion is zero optic flow or sonar flow has its maximal effect on the position estimate.

[0066] Sensory fusion may weight the importance of incoming signals and modalities depending on the state of the system and the advantages of individual sensory streams. Some states shall be used to perform on-the-fly recalibrations of the sensors that can show significant drift during extended performance.

[0067] Figure 6 presents an example flow chart for this algorithm. System acquires feature field, for example an optic flow field on every iteration or frame 600. If there is no signal the field will be flat with all values equal 0, so the system proceeds to the next frame 602. If there is signal, then the system computes a velocity field by taking difference between this frame and the previous frame 604 and acceleration field by taking the difference between velocity field from the previous frame and the newly computed one 616. It may then determine whether the velocity field is flat 606. The velocity field can be flat only when there is no linear motion 614, so the system can override other inputs and ensure that the coordinate of the agent does not change its representation. After dividing the velocity field by distance 608 the system may determine if the field is flat 610; dividing may produce a flat field only if there is no angular motion 612, so the system can be sure that the heading of the agent does not change. Acceleration field 616 is used in a similar fashion to determine angular and linear accelerations and signal an opportunity to recalibrate IMUs when respective accelerations are 0 (e.g., the system checks to see if the acceleration field is flat 618, and may calibrate the accelerometer 626 if the field is flat; if the field is not flat, the system may divide the acceleration field by distance 620, and if the field has become flat 622, the system may calibrate the accelerometer 624.

[0068] Linear displacement is first estimated along three axes in the plane; the use of three axes instead of two provides redundancy and error detection (band cells 404 in Figure 4). Each axis has a set of corresponding ring attractors – all having a preferred direction aligned with the axis, but each having different spatial scales: the displacement covered by one full rotation of attractor around the ring. The neural system achieves its best localization results with spatial scales related to each other as different prime numbers (e.g. 7 to 11 to 17 for three scales). The merged heading estimate is filtered through a pair of Gaussian filters (one positive and one negative, centered 180 degrees apart from each other) to produce a cosine of an angle between the current heading and a preferred direction of a linear displacement ring attractor. This cosine, multiplied by the linear velocity of the agent (estimated by different sensory inputs), serves as an input to the corresponding set of ring attractors. The merged displacement estimates along axes are combined together into 2D displacement within the corresponding spatial scale (grid cells 406 in Figure 4). Finally, 2D maps of different spatial scales are merged together to produce a unique (within a map) current location estimate for the robot with respect to its starting location (place cells 408 in Figure 4).

[0069] Note that the final output of PInNAcLE system is an estimate of the agent's position and orientation in space, but instead of estimating these outputs from three different sensory inputs independently the information is combined on the stages of heading estimation and linear displacement estimation, and these combined more precise estimates are used downstream to produce final results. Similar to the visual processing example above, this downstream system also is independent of the sensory composition of early processing in PInNAcLE.

[0070] To summarize, unlike the vertical streams of processing used in conventional approaches, the biological approach suggests high level of horizontal interactions between streams and clear separation of processing stages in terms of defining intermediate outputs that can be used elsewhere in the system.

[0071] *Feedback from late stages to early stages within and across processing streams.*

[0072] Introduction of continuous delayed feedback in the system induces oscillatory regimes that usually complicate both mathematical analysis and performance tuning of the system. As a result, state of the art solutions for sensory processing tend to utilize corrective feedback mainly in the form of discrete resets. Furthermore, as a consequence of vertical stove-piping, the results of many sensory estimates are not trustworthy enough to use as a

corrective feedback. For example in loop-closure solutions visual input is used to recognize whether the agent completed the loop during navigation, and if so it is used to correct odometry-derived position estimate. Rarely if ever odometry is used to affect decision making on whether the loop was closed. As a result, the system makes errors when a view is identical in different parts of the environment (which happens often indoors), while a feedback from odometry could have eliminated a significant portion of these errors that just could not be loops even with low precision of odometers.

[0073] The brain on the other hand is an inherently oscillating system that shows us how oscillations can be utilized for synchronization, noise reduction, and dynamic control. Since brain puts oscillations to use rather than avoid them, the dynamic feedback is a key mechanism in the brain functionality and can be utilized for much more than just corrective input. Figure 5 shows how in PInNAcLE the fast motor input by itself tend to overshoot the motion, while optic flow based signal tends to undershoot, primarily because it takes longer to process and thus comes with a delay. Results of heading estimates from all sensors are merged together in a weighted fashion described in the previous section, allowing the system to take advantage of the best ranges of precision of individual sensors. The result of this merge is computed in the ring representing the final estimate and then provided back to individual sensors, correcting the individual errors and ensuring that all sensory estimates stay tuned to each other and do not diverge with time. In some cases, early corrective feedback to sensory readings reduce error in the system on early stages of processing.

[0074] Providing corrective feedback to these signals from a combined version on one hand introduces more inertia in the system (consider it another low-pass filter) and on the other hand allows keeping the Gaussians of all individual estimates within a short range of each other (as discussed in the previous section it is required for successful combining of these signals). Furthermore, the level of feedback can determine how much individual channel matters: increasing feedback to optic flow based estimate will reduce the lag of this estimate, but also make it less contributing input. Dynamic weights in the merging and feedback allow manipulation of the importance of each individual sensory stream for the final result during the operation of the system and depending on the conditions changing due to external events as well as the robot's own actions.

[0075] Since the fusion of different inputs happens in the system as early as possible, it reduces the error at early processing stages and prevents error propagation through later stages of processing. The resulting system a) produces positional estimates that are as good as

the best precision sensors it employs allow it make, b) can merge an arbitrary number of these sensory streams, and c) will continue to operate as long as at least one of these sensors is operational. Early fusion of sensory streams allows the rest of the system to be agnostic to the number and quality of sensors, which makes the system more robust and maintainable.

[0076] Correction of drift in the sensory information is not the only, and not even the main reason to use feedback in sensory processing. The main advantage the feedback can provide is the speed of processing. In the example implementation of active vision system OpenEye discussed below a visual object recognition cascade of the system may recognize a blob in a bottom left part of the visual field as a pickup truck. Motion detection sensory system may detect that it is moving to the right. The motor outflow system has full information about how the agent plans to move the camera. Given that all these processes are continuous, it automatically entails that in the next moment of time the OpenEye system knows precisely where in the visual field this truck shall be by combining feedback from object recognition with low level sensory/motor information.

[0077] This may result in the ability of OpenEye system to avoid further engagement of all the complicated cascades of the system, all the system has to do is to feed back to the early stages the expectation of what should be there and check how well it matches the input. If there is a match, then the expensive object recognition processing is not necessarily needed. If there is a mismatch, then the system engages full power to process why its prediction was wrong and how to correct it. This ability to disengage 90% of the computational capability and only concentrate on important parts of the input may outweigh the initial complexity the system. This advantage in terms of the cost of operating resources may manifest itself in wide spread applications as the energy efficiency becomes more and more important. Feedback across different modalities creates the predictive power in the intelligent system, so it can work in a predictor-corrector mode. This mode frees computational resources when correction is not needed.

[0078] The process described in the previous paragraph can be taken one step further by creating a memory map of a virtual visual field not limited by the current field of view. An example OpenEye implementation can place certain objects in the virtual field of view, turn 180 degrees, then turn back to the original view and expect to see the previously recognized objects there.

[0079] **Robotic Navigational Systems**

[0080] As an illustrative example of principles discussed in the previous section, this section provides a description of a robotic navigation system based on these principles.

[0081] Development of a biologically-inspired, robust, low-power, multi-component neural system able to perform self-localization and mapping will enable robots to autonomously navigate novel terrains without a global navigation satellite system (GNSS). By including the ability to learn about an environment as it explores, robotic agents will be able to autonomously negotiate novel terrains and send relevant, intelligently preprocessed information back to a human controller.

[0082] In some implementations, the system and methods described herein reduce development costs and time by allowing robots to be trained rather than explicitly programmed. The systems and method described herein can be implemented into robots, autonomous vehicles, and other moving agents for use in a plurality of applications, including: defense (e.g., Unmanned Aerial Systems (UAS) surveillance, patrol, rescue, demining); business (e.g., telepresence); home (e.g., cleaning, lawn mowing, toys); healthcare (e.g., remote diagnosis, assistive living); and agriculture (e.g., pot moving, autonomous seeding, crop assessment, wildlife conservation).

[0083] In some implementations, the navigation system and method described herein include neural systems mimicking the ability of biological organisms to navigate their environment as well as build usable and correctable representations of these environments through learning without a GNSS and active sensors.

[0084] The navigation system can navigate, map, and detect objects in land and aerial environments. Further capabilities of the system include visual learning and search, collision avoidance, a human-robot interface for robot training and deployment, and low-power hardware implementation of these algorithms.

[0085] According to one aspect of the disclosure, the system includes a path integration system that calculates optic flow, infers angular velocity from the flow field, and incorporates this velocity estimate into heading calculations. The resulting system fuses heading estimates from accelerometers, gyroscopes, engine torques, and optic flow to determine self-localization. The system also includes a motivational system that implements a reward drive, both positive and negative, into the system. In some implementations, the drives can include: a) a curiosity drive that encourages exploration of new areas, b) a resource drive that attracts the agent towards the recharging base when the battery is low, and c) a mineral reward drive that attracts the agent towards previously explored scientific targets.

[0086] In some embodiments, the system's biologically inspired navigational system includes a number of subsystems for an autonomous navigation. The subsystems can include: Path integration system for maintaining a sense of current position and direction relative to starting position and direction. Spatial representation incrementally built as the agent explores its environment. Path planning system that uses information from the spatial representation and path integration system to explore, return to base, revisit objects of interest, or avoid obstacles. Drive selection system to balance battery recharging, exploration, and return visits to explored locations of interest. Sensory system used to provide vestibular, motor, and optic flow information to drive and correct the path integration system.

[0087] The present disclosure also provides experimental results of the disclosed system in both a virtual environment and as implemented into an autonomous rover.

[0088] The navigation system can include an artificial brain model that constitutes integrated perceptual, decision-making, and motor systems. Using the navigation system, the agent can interact with its environment in an intelligent way based on its internally generated goals. The model can contain many interacting subsystems. For example, the subsystems can include a Path Integration and Neural ACquisition of Location Estimate (PInNAcLE) system. The PInNAcLE system can transform velocity data recorded from the agent's sensors and/or actuators into a world position of the agent. This world position can be used by the other modules described herein to learn positions of objects in the world. The PInNAcLE system is described in below. Another subsystem can be the MOTivation, Reward, and goal SElection (MoRSel). In some implementations, the MoRSel subsystem codes agent drives, learns maps where each drive can be satisfied through corresponding rewards, and/or selects a goal map location based on relative drive strengths. The MoRSel system is described below. Yet another subsystem can include the Spatial Planning and Allocentric Representation of Knowledge (SPARK) system. The SPARK subsystem can plan a path and provide a goal heading direction based on the current location, goal location, and/or the locations of learned obstacles. The SPARK system overview is given below.

[0089] *Path INtegration and Neural ACquisition of Location Estimate (PInNAcLE)*

[0090] The navigation system's PInNAcLE subsystem follow the biological path integration system outlined in Figure 4 and can be based on joint information from a number of sources. The PInNAcLE system can include input from motor commands producing self-motion of the agent, input from vestibular or inertial sensors (e.g., accelerometers and gyroscopes), and optic or other sensory flow fields (e.g., from video cameras or sonar

sensors). The PInNAcLE system can maintain the agent's sense of location within the environment. The motor output commands can add the ability to anticipate movement before the agent performs the movement, and optic flow can add precision to the estimates of angular velocity especially important for low-velocity rotations as described in above. Detailed description of the optic flow computation used to generate result presented here is provided below.

[0091] In some implementations (e.g., Figure 20), a navigation system 2000 may be able to connect to a robot 2002 in order to provide guidance 2030 on where to travel within an environment. For example, the robot may have at least sensors A 2004 and B 2006, which each may correspond to a plurality of types of sensor data 2020a and 2020b. In some implementations, sensor types may include IMU sensors, motor outflow sensors, and optic flow sensors.

[0092] The robot may send the sensor data to the path integration module 2008, which may include a processing device 2008a (e.g., a processor and/or other computing device), ring attractor networks 2008b, and a map module 2008c. In some implementations there may be four ring attractor networks for IMU sensors, four ring attractor networks for optic flow sensors, and four ring attractor networks for motor outflow sensors. In some implementations there may be one ring attractor network per sensor for determining a heading, and one ring attractor network per sensor per spatial scale for linear displacement. The processing device may be able to process the sensor data from the sensors by weighting the sensor data (e.g., in relation to other sensor data) processed by the corresponding ring attractor network, and by combining the sensor data. The combined data may then be used to generate environment data (e.g. information about the robot's current position in its environment and/or other data related to the environment). The map module may use the environment data to generate a plurality of maps related to different data (e.g., different maps may be generated for different goals, different types of sensor data, and/or the like. In some implementations, the map module may be separate from the path integration system.

[0093] The path integration module may interact with a motivation module 2012, a goal selection module 2014, and a path planning module 2016 in order to generate those goals and/or to generate other data. For example, the path integration module may provide the map data 2022 to the motivation module so that the motivation module can determine reward/drive pairs 2024 to provide to the goal selection module for the robot. The reward/drive pairs may be determined by examining maps and determine potential reward

items in the environment and codes and/or values to relate to those reward items. The goal selection module then may weigh the reward/drive pairs against one another to determine at least one pair to prioritize. The prioritized pair may become a goal 2026 which may be sent to the path planning module. The path integration module may send the map data to the path planning module, in order to use the map data and the system's spatial understanding of the environment, and to use the goal criteria in relation to the environment, in order to generate a path 2028 from the current position of the robot to the robot's goal. The path planning module may forward this path to the robot in order to allow the robot to understand where to travel next 2030 and to path integration module as motor outflow signal to be integrated as expected path 2028.

[0094] Both path planning and path integration modules can utilize the same internal spatial representation system that follows the principles outlined in above and holds several specific 2D maps in it. A human readable map can be constructed by blending the goal selection maps activated by each drive (drives can be color coded to mark the corresponding reward type) and an obstacle map. The resolution of the map can depend on the spatial scales of grid cells in the PInNAcLE subsystem and can be adjusted according to the user's requirements.

[0095] *MOtivation, Reward, and goal SElection (MoRSel)*

[0096] The MoRSel (Motivation, Reward, and goal SElection) subsystem is an extensible internal motivation framework that can allow for the definition of an arbitrary variety of agent rewards. Each type of reward can be matched with a corresponding drive that, in some implementations, can be calculated as a function of the agent's internal state. These drives can be analog values that may be modulated or explicitly set by an end-user. When the agent navigates to a location that contains a combination of rewards, it can learn the global position of those rewards by increasing the value of synaptic weights between the corresponding drives and the place cell that codes for the agent's current location. Mapping reward values to drive weights may allow the agent, upon the growth a particular drive, to reconstruct an allocentric reward map in order to return to a reward location even after it has traveled away from that position.

[0097] The MoRSel subsystem can contain reward-drive pairs. In some implementations, the subsystem contains three reward-drive pairs, where each drive or reward signal can range from 0 to 1. The MoRSel subsystem can switch between navigation goals based on the value of drives and learned potential rewards. Exemplary reward-drive pairs can include: 1. A

resource drive that monitors the internal battery level and/or similar depletable resources. For example, as the battery drains, the drive to find a source of replenishment increases. 2. A mineral drive that is a drive to a target of interest. For example, the mineral drive can be satisfied when the agent is in proximity to a target rock that has a reward value. In some implementations, the reward value is designating as a scientific or other interest attached to it for example rocks on Martian surface that can potentially contain water. 3. A curiosity drive that can be a drive to explore an environment. In some implementations, there may be no explicit rewards that have been placed in the environment, but the curiosity drive is instead remembers all previously visited positions in its internal environment map and drives the agent to unexplored locations. This drive can decay once the agent's environment has been fully explored.

[0098] The navigation system can learn the positions of rewards and obstacles by navigating to those locations and making memories that record the presence or absence of motivationally (i.e., reward / deterrent) stimuli. These memories may be stored as synaptic weights between the drives described above, and global position coded by the Path Integration (PInNAcLE) system. These weights can be updated for example using simple associative learning rules, temporal difference reinforcement rules or Q-learning method.

[0099] In some implementations, the navigation system can treat obstacles and rewards differently (i.e., each has a separate map and/or a different learning rule). A separate map may be used because obstacles restrict path planning through the inhibition of the bidirectional search described below. Obstacle weights can decrease when the agent is located at that position, but the weight can increase when the agent is at a nearby location and simultaneously sensing an object in the direction that would lead to the weight's coded location. The process of transforming a local proximity or touch signal into a world position is described further below.

[0100] The goal selection system may balance the strengths of one or more drives with the agent's confidence of the availability of rewards at different environment locations in order to choose where to navigate to substantially satisfy its currently active drives. Before selecting a location as a goal, the system may mask locations that would either place the agent at the position of an obstacle or out of the range of positions that the agent can represent. . In some embodiments the presence of an obstacle attracts the mask value towards 0, and may keep it from being selectable as a goal location. In some embodiments, the mask value slowly recovers if the obstacle weight shrinks due to evidence that the location is

navigable. In order to keep the agent from leaving the area representable by the PInNAcLE system, the positional areas of the mask may be silenced so that they may not be selected as potential goal locations.

[0101] After the candidates for the goal are selected, a new goal position may be chosen through a competitive process, for example a winner-take-all function. Once this winner is selected it is held in memory either until a reset signal occurs or until the location is reached by the agent. If the reset signal happens before the goal is reached, the whole process of goal selection repeats and a new goal (potentially different from the old one) is selected.

[0102] **Spatial Planning and Allocentric Representation of Knowledge**

[0103] The Spatial Planning and Allocentric Representation of Knowledge (SPARK) system may translate agent senses and actions between egocentric and allocentric coordinates in order to control agent exploratory behavior. On one hand, SPARK transforms the agent's current allocentric position from the PInNAcLE system (above) and the goal location selected by MoRSel into a substantially optimal heading for quickly and successfully reaching the goal location and also driving the agent motors. On the other hand, the touch and proximity senses provide local egocentric information as to the proximal location of obstacles, which SPARK transforms into an allocentric spatial position to be learned as an obstacle position by MoRSel.

[0104] *Allocentric representation of knowledge*

[0105] The allocentric spatial representation in the system may model neural representations found in the mammalian brain. Place cells produced by the PInNAcLE system (above) may provide a 2D Cartesian coordinate of the agent on the surface, and head direction cells provide the heading. Both systems may use similar representations as a bump of activity in the neural field that is centered on a corresponding position or direction as discussed in above; the difference being that the positional map is 2D and the directional map is wrapped-around in 1D. The positional signal is may be used for goal and obstacle mappings, and the directional signal may be used in coordinate transformations between egocentric and allocentric representation for obstacle mapping and motor command calculation. Both systems may maintain their relative position and heading with respect to initial position and heading; anchoring these systems to absolute coordinates.

[0106] When the agent registers contact with an obstacle, SPARK translates this information into an allocentric world location of the obstacle that is learned as un-navigable when planning future routes through the environment. The obstacle's heading, compared to

the agent, is rotated according to the agent's internal heading in the PInNAcLE system to give an allocentric direction of the obstacle relative to the agent. The position of the agent, represented by an active place cell in the agent's PInNAcLE system, may be used to increase the obstacle weight for an adjacent place cell that represents the position where the obstacle lies according to the allocentric direction calculated above. In this way the proximal obstacle may be learned on an allocentric obstacle map.

[0107] A chain of neural populations may also do the inverse conversion of allocentric to egocentric coordinates in order to provide motor commands. An example system may convert the desired destination coordinate into the rotational velocity motor command, which then controls navigation for the agent. The position-to-direction map serves as the first stage of the transformation from desired location to motor command, and converts the allocentric position of the desired destination into an allocentric desired direction of movement by comparing the desired position with the current position of the animat. A desired direction ring compresses representation of the position-to-direction map into a one-dimensional allocentric desired direction vector. The direction-to-turn map then compares the allocentric desired direction with the allocentric current head direction and computes the egocentric turn force vector. To drive the vehicle, this signal can be interpreted as steering angle for the left or right turn.

[0108] *Spatial Planning*

[0109] SPARK's spatial planning algorithm may be based on a neuronal parallel breadthfirst asymmetric bidirectional graph search where each location in the path planning system corresponds to a node in a graph. Edges in the graph may initially exist between neighboring locations, but as the agent explores the environment and encounters obstacles the corresponding edges may become impassable through obstacle learning and may be excluded from the breadth-first probing. This may result in a system that can adapt to found configurations of obstacles and find the substantially shortest path around these obstacles.

[0110] In general, the SPARK algorithm may model propagation of waves of activity from the goal location (reverse activity spread). In an unconstrained environment these wavefronts may form expanding concentric circles around the goal location. Wave propagation may not be allowed if an obstacle is encountered, so the wavefront deviates from the circle to flow around the obstacle as shown in Figure 7 by depicting obstacle map 700, wave propagation 702, and a combined view of both populations 704. As a result of all deviations, the first wavefront segments that reaches the current location of the agent comes from the direction that corresponds to the substantially shortest path to the goal. In response

to this, a new wavefront is initiated at the current location (forward activity spread). The collision of this front with the second wave coming from the goal may then happen at the location closest to the current location of the agent and in the direction of the shortest path towards the goal. This is called the next desired destination. Then the agent moves to this desired location and the process is repeated. A reset signal may be generated when the graph search finishes, which allows a new goal to potentially be selected as described above.

[0111] *Robotic implementation*

[0112] The above described navigation system may be implemented into a robotic system, autonomous agent, or other moving body. This section describes one exemplary implementation of the navigation system into an example robotic system.

[0113] The navigation system can be implemented into any microprocessor. In this exemplary implementation, the navigation system was implemented with Lego Mindstorms NXTs outfitted with 7 sensors and 4 motors. A Lego NXT is a microcontroller that can be interfaced with many different types of sensors and motors, all which require no soldering or significant setup time. The NXTs communicate with a low-power, low-cost single-board computer development platform (Pandaboard) which controls communications with a neural application client. The Pandaboard runs an asynchronous server that listens for commands from the neural application client while reading data from the NXTs and camera. When data is received from an NXT, it is repackaged and sent over UDP to the neural application client. The Pandaboard utilizes OpenCV to process image data from a USB camera and preprocesses/broadcasts it to neural application over TCP through an onboard router connected to the board through a Cat5e cable. The neural application client is capable of sending movement commands and can also schedule tasks on the NXTs.

[0114] Figure 8 shows a hardware diagram and the protocols used to connect the devices. In some implementations, a Lego Mindstorm Robot 800 may have left-front wheels 802, right-front wheels 804, back-left wheels 806, and back-right wheels 808. The robot may have sensors such as an RFID scanner 810, an EOPD sensor 812, an accelerometer 814, a front ultrasonic 816, a left ultrasonic 818, a right ultrasonic 820, and a gyroscope analog sensor 822. The sensors on the robot may be configured to communicate with a Pandaboard 824 configured with LCP over USB connectivity. The board may also support a webcam 826 and/or other camera and/or image-capturing device, as well as a communication device 828 configured to allow communications between the robot and another electronic device (e.g., a computer 830).

[0115] Figure 9 illustrates the main objects within the software of the system and their respective hardware hosts and communication streams. In some implementations, a PandaBoard 924 may be configured with a MasterCommandThread 900 which may send NXT Commands 902a and 902b to a Lego Mindstorm board configured with an NXTCommandReceiver 904a and 904b which may receive instructions and/or other data from the PandaBoard. Each robot may also be configured to support action loops 906a and 906b, which may include actions such as PackageSensors 912 (which in some implementations may be configured to interact with sensors 914 on the NXTController 922), and/or PackageDistanceSensors 908 (which may be configured to interact with DistanceSensors 910 on the NXTController). The PandaBoard may also be configured with an USBtoTCPCamera module 916 which may allow the PandaBoard to receive data from a webcam 918 and/or a similar image-capturing device and convert it into a stream over TCP protocol. The NXTController may also periodically send control commands 920 to the PandaBoard when necessary.

[0116] The robot was equipped with three ultrasonic sensors that detect objects in proximity. These readings are then used by the navigation system to form an obstacle map to avoid collisions. These sensors were placed on the right, left, and front of the robot. An accelerometer and gyroscope were used to detect the robots acceleration and angular momentum for use with the PInNAcLE subsystem. RFID tags and an RFID scanner emulated the more complex chemical scanners used in Mars Rovers. The RFID tags were placed around particular rocks, with each emitter tuned to a particular frequency. When a tag was detected, the RFID signal simulated a chemical detection of various aspects of rock geology (e.g., a particular chemical signature). The model reward system was then triggered and the rock was placed in the robot's internal map of the environment. These tags are also used by the robot to simulate recognition of the recharging base.

[0117] The robot experiments were conducted in a ~2x2 meters indoor simulation of the Martian landscape. This indoor facility provided a scaled-down, enclosed area with rocks and small terrain variation (the terrain is a tarp sheet with small elevation discontinuities) to test vision and navigation aspects of the navigation system.

[0118] *Model Performance*

[0119] The robot was tasked to explore the experimental environment; tracking its past and present position through the intelligent integration of self-velocity parameters given by inertial, motor, and a visual flow sense, learning the position of task-relevant objects such as

a home base, scientific targets, and obstacles; and managing competing drives for exploration and management of internal resources such as its battery level. The structure and physics of the virtual Mars environment were created and simulated with the JMonkey game engine, passed to the model using Virt-U package, and the neuromorphic algorithm controlling the virtual rover was simulated separately after having been written in the Cog Ex Machina simulator or INDRA simulator (results not shown).

[0120] Figure 10 diagrams the specific phases of the main task. The virtual rover 1000 begins its exploration by leaving its base with full battery in order to begin exploring. (Figure 10, 1002). The rover explores unknown areas, maintaining its relative position and orientation with respect to the starting point and progressively building and updating its spatial representation of the environment mapping the obstacles 1004. This map is useful when the rover will, for instance, seek specific objects and locations in the environment or as a source for conversion into human readable format. After exploring for some time the rover's battery depletes 1006 to the point where it becomes motivated to return to the base 1008 where it began exploring by as direct of a route as possible (Figure 10, 1010). After recharging its battery the rover continues to explore from where it left off (Figure 10, 1012). Figure 11 show the actual results in the virtual environment, and Figure 12 shows an incremental map view of the same trajectory (e.g., 1200 corresponds to path 1002, 1202 corresponds to path 1010, and 1204 corresponds to path 1012).

[0121] Figure 5 compares the compass heading of the physical rover through the simulation to heading determined by the integration of velocity information from one of several senses. Note that the final head direction estimate is determined by the integration of a combined estimate from all three senses as outlined above.

[0122] Autonomous exploration may include the management of multiple, overlapping, and oftentimes conflicting drives that keep the agent relatively safe and efficiently completing its task(s). The simulation requires the rover to balance its exploratory drive with a need to keep its battery charged; the balancing of goals by the navigation above. Figure 13 plots the rover's battery level over the course of the simulation 1302. The battery is empty at the beginning of the simulation 1304 in order to motivate the rover to learn the position of the charging base. Gating the learning of a reward by the desire for that reward may be a biologically realistic method of conserving memory and computational resources while retaining survival-relevant memories. When the battery is low enough, the rover revisits its base in order to recharge, as shown by the shaded areas 1306 and 1308 in Figure 13. The

rover takes longer to find the base on the second visit at minute twenty-three, both because the rover has traveled further from the base and because accumulated path integration error forces the rover to search in the immediate vicinity around its memorized position of the base. In both cases, however, the rover is able to find the charging base, recharge, and continue its exploration; the second decision to recharge is also shown in Figure 12 as the part of the trajectory stretching from about 0 m latitude to about 100 m latitude in the middle graph.

[0123] In addition to testing the performance of the navigation system in the virtual environment, the model was also demonstrated to operate on a physical robot. The construction of the robot and environment is detailed above. Figure 14 shows a short simulation of the robot moving in a preset circular arc. While traveling in this arc the robot moves over radio emitters that symbolize the charging base and a scientific target. In this simulation the navigation system's path integration (PInNAcLE) subsystem is driven exclusively by the motor system's preset velocity, sent as individual tread torque levels. The robotic platform is able to map out its course in the environment 1402, learn the positions of the same types of task-relevant objects as in the virtual environment 1404, and is able to recharge a simulated battery when moving over the charging base 1406.

[0124] The robot was also turned in place by setting the angular velocity to the sine of the simulation time scaled so that it reverses its direction of rotation around the vertical axis every twenty seconds. This preliminary test shows that the camera on the physical robot is able to accurately determine the robot's angular velocity at slow rotation speeds. In some implementations, increasing the camera's frame rate can increase at which speeds the navigation system can detect rotations. In some implementations, detection of rotation can be improved by smoothing the motion of the camera by adding a suspension system to the robot or by gating optic flow's influence on path integration by the other senses' estimates of rotation speed as disclosed above.

[0125] **OpenEye: An artificial visual system based on neural principles**

[0126] The OpenEye model is an active vision system that autonomously make decisions about where to look, generates saccades to target locations, and learn online view-invariant representations of salient objects. OpenEye models the parallel visual processing streams of mammalian visual cortex, and integrates in a single framework several capabilities that are typically modeled separately, such as the influence of bottom-up saliency and top-down spatial evidence on saccade decision-making, view-based object recognition, stable

incremental learning, and space-variant approaches to vision modeling. In some implementations, the OpenEye and navigation system learns incrementally with a single presentation of the stimulus 1506, for example, it does not require a separate training and testing phase. OpenEye may also be able to operate with or without supervision.

[0127] In some implementations, OpenEye integrates bidirectional connections 1500 between dorsal (where) 1502 and ventral (what) visual processing streams 1504 (e.g., see Figure 15). OpenEye's where pathway may use simulated eye movements to navigate through a map of candidate areas established by bottom-up surface properties and top-down spatial evidence in order to explore distinct views of each object. By using a log-polar representation of the retinal visual field, OpenEye may sample the image "view" generated by each saccade in a space-variant manner such that an object in the center of the system's visual field (the fovea) is represented in high resolution while objects in the periphery are represented with nonlinearly decreasing resolution. This may save a processing power, as, for example, an image field of view can be processed in an order of magnitude smaller neural population using the log-polar compression.

[0128] Once the where pathway has chosen a point of foveation, it forms a surface-fitting attentional shroud around the object in order to isolate it from other surrounding objects. This is implemented using feedback principle outlined above. The attentional shroud may provide OpenEye with the capability to discover whether sequential views belong to the same or to different objects. The what pathway may be modeled as a series of hierarchical classifiers that cluster objects into view, object, and/or name categories. After the what pathway has identified an object, the what-to-where feedback may use learned prototypical object-feature maps to produce "smart saccades" in order to focus on spatial locations that yield greatest disambiguation of competing object identity.

[0129] The Martian Virtual Environment used to train OpenEye was created using the jMonkey 3 graphical engine and interfaced with the neural model using Virt_U package. OpenEye was trained to recognize two types of Martian rocks (type 1 1600 and type 2 1602) in Figure 16 with a supervising name signal for 250 cycles. The scene is then switched to a similar scenario, but this time, there was no supervision, and there was a third, novel type of rock that the system has not been trained on. The accuracy of classification in this testing scene is close to 100%. At first, the system is not able to recognize the novel rocks, and there is no name associated with the newly created object. This triggers a novelty signal in OpenEye, which can be returned to the user as an alert, or be utilized to capture snapshots of

the novel rock for closer examination. This capability allows the system to generate alerts and capture views 1604 not only of known science targets, but also novel targets 1606 of potential interest to the scientist. Eventually, once OpenEye has learned several views of these novel objects, it is able to classify them properly upon subsequent foveation, and the rock becomes part of the "object dictionary" of OpenEye, where a user can choose to attach a verbal label to the rock.

[0130] Virt-U: Apparatus and methods to connect off-the-shelf 3D engines or physical robots with neural simulation software packages.

[0131] The system disclosed herein includes a virtual universe (Virt-U) software component 1700 designed to simplify prototyping of neural models of robotic intelligence and test them in a virtual environment before implementing on physical robotic platform. With Virt-U a user can connect an existing neural modeling software package with another existing 3D simulation engine, for example a game engine.

[0132] In this scheme virtual environments 1702 supply rich sensory and proprioceptive signals attached to an agent, while the neural model 1704 (also referred to as the brain) receives and processes sensory data to produce motor control commands required to elicit the agent's corresponding behavior (e.g., see Figure 17).

[0133] The example system binds to a neural engine of choice that may allow for neural models to be executed in a substantially parallelized, distributed computation environment possibly on one or more interconnected computer systems with GPUs. In two example instantiations we connected Virt_U to Cog Ex Machina software from Hewlett Packard, and to INDRA engine developed by Neurala Inc. Virt-U then delivers sensory data to the neural model and behavioral commands (translated to physics and graphics) to the 3D engine. The same API that Virt-U exposes to neural engine for connectivity with the virtual environment is used to connect neural engine with the robotic platform.

[0134] *Virt-U Architectural Principles*

[0135] Virt-U architectural principles abstract the interface that connects a virtual world or robotic platform with its neural modeling environment. In order to segregate dependencies of the neural modeling environment from those of the virtual environment engine the Virt-U architecture includes of two major layers: the VE layer and the proxy layer. The VE layer abstracts the virtual environment engine and ads functionality that for example allows extracting sensory information that is not usually available from the game engine. The proxy layer delivers this sensory information to the neural model and extracts neural information

from the neural model in order to perform the behavior. The Virt-U can have two modes, particularly: 1. Simulation - when the behavior of the animat is controlled by the brain, and 2. No-brainer - no brain is attached, used for testing virtual environments.

[0136] Virt-U considers the world to be populated by virtual objects. These virtual objects can be animated (brained) and non-animated (brainless). Brained objects, called animats, can be controlled by neural model. An animat's brains receives sensory (e.g., visual) and/or proprioceptive signals and can then exhibit a behavior that follows the laws of physics to a predefined level. Animats can be either virtual animals (e.g., a rat) or vehicles (e.g., a rover), depending on how their motor functions are controlled, but a brain in both cases controls behavior. Brainless objects, called items, are entirely controlled by the virtual environment and obey its physics. Items include rocks, buildings, trees, etc. and can be considered a source for sensory information including rewards sensed by animats.

[0137] Virt-U considers an animat to be a collection of sensory organs and animat controls. Sensory organs may in turn include sub-sensory organs, which ultimately can be connected with virtual sensors dedicated to collecting sensory information. For example, an animat may have a sensory organ called "eyes" that comprises individual eyes, where each eye contains facets connected with virtual cameras. All sensory organs are responsible for maintaining and updating sensory information for the sub-organs they contain thus allowing for recursive sensory data update to be performed on an animat. Animat controls constitute outgoing interfaces that may be accessed by the brain to perform the predetermined behavior. Sensory organs and animat controls expose incoming and outgoing interfaces to be accessed from the Virt-U proxy layer in order to supply and deliver sensory and motor control information.

[0138] On the neural modeling side (brain side), proxy organ controllers perform access to the VE. These controller can retrieve and supply neural data by accessing a specific organ controllers' incoming and outgoing interfaces. Proxy controllers can abstract internal knowledge of sensory organs from the neural model by acting as hardware drivers within the neural modeling environment. In simulation mode, Virt-U can create the virtual world based on an XML description. This description is sufficient to define the world, the animats with all their sensory organs and controls, and all virtual objects with their associated rewards. This XML-based virtual world description ensures Virt-U portability between various neural modeling environments.

[0139] *Virt-U VE Layer for JMonkey Game Engine*

[0140] An example implementation of the Virt-U VE layer presented here complements the open source JMonkey game engine.

[0141] *Virt-U Application and Library*

[0142] The Virt-U library platform may be a general-purpose, Java-based game development environment that extends 3D graphics game development capabilities to the Java programming community. The Virt-U may include a system that maintains a variety of assets, for example, models, textures, sounds, whole scenes, rendering of 3D graphics, and the ability to simulate physical interactions.

[0143] The Virt-U library layer may isolate the development of a brain-controlled animat from the specifics and complexity related to 3D game programming. On the user level, a user may create and run a virtual world populated by an animat by writing Java code, instantiating the classes representing the world and the preexisting animat, and adding external cameras to the world as necessary. In some implementations, once the classes are instantiated, the VE is run by the JMonkey game engine and can be controlled via the JMonkey GUI.

[0144] In some implementations, Virt-U application is a platform application that builds custom virtual worlds at runtime, which extends the DebugWorld class and implements the WorldBuilder interface comprising a single construct method. Virt-U application is embodied in DebugApplication - an executable intended for testing custom virtual environments in the no-brainer mode and generation of world XML-based description to be used in simulation mode.

[0145] As a command line argument, Virt-U application might take a name of a custom class that extends DebugWorld. An example of such a class is presented below.

[0146] Once the construct method yields, the debug application runs under control of the physics engine and can be interacted with via JMonkey GUI, which interprets keyboard and mouse actions as animat movement commands. The application also extends the Virt-U command set that allows switching between external cameras views and first-person views from the perspective of different animats. In some implementations, the Virt-U Builder can have one or more user customizable features, which can include:

[0147] a custom set of VE assets that determine how each individual object looks to the observer and allows for positioning of these objects in the scene;

[0148] different pre-created animats placed at specific locations in the VE scene; and

[0149] placement of virtual cameras characterized by a custom set of parameters to achieve a third-person view of certain portions of the scene.

[0150] *Running Virt-U Application in the No-brainer Mode*

[0151] In some implementations, the physics engine runs VE applications on its internal clock that synchronizes physical interactions with 3D graphic rendering. Users can control execution of the Virt-U application in no-brainer mode via the GUI described in Table 1.1 below. In other implementations Virt-U library exposes all the necessary API to an external application and provides it with all the necessary functionality to establish a similar GUI bypassing Virt-U application.

[0152] *Virt-U Implementation*

[0153] Implementation of Virt-U may comprise Java classes (e.g., see Figure 18) and interfaces that are extended or implemented to abstract the functionality to supply sensory information to the neural model as well as perform a behavior based on exposed motor control interfaces. The following sections discuss example Virt-U classes.

[0154] *VirtUApplication Class*

[0155] The VirtUApplication class 1800 is a focal point of communication between the game engine and Virt-U. The class implements JMonkey's SystemListener interface and thus provides an entry point for handling GUI events, rendering graphics, updating sensory information, and/or exercising controlled behavior. The VirtUApplication class initializes components of the JMonkey game engine, creates a rendering context, and instantiates the world. When JMonkey Engine calls the update method of the SystemListener interface in the VirtUApplication class, the latter may call the update method on the World class. In turn, the World class may call the update method on all objects populating the world.

[0156] VirtUApplication is extended by the DebugApplication class that is part of Virt-U application and runs Virt-U in no-brainer mode. In simulation mode, the VirtUApplication class may be extended by different application classes that expose the engine in different modes: synchronized to the neural simulation clock or running asynchronously. On the neural simulation proxy level these classes can be further extended or wrapped depending on the neural engine implementation. In one example of the proxy layer the CogVirtUApplication provides an entry point to control execution of Virt-U from Cog Ex Machina package written in Scala. In another implementation VirtUApplication is wrapped in Java Native Interface object to provide an entry point to the engine from INDRA package written in C++. Note that VirtUApplication is a singleton class and, in described implementations, only one instance of VirtUApplication can exist.

[0157] *World Class*

[0158] The World class 1802 serves as the main container for landscape, animats, items, and external cameras. The world appearance is determined by the scene model specified as an argument to the addScene method. The scene is encoded by the model's .j3o file that can be generated by Virt-U application or outside of Virt-U using third party software. In some implementations, by default, Virt-U adds an external chase camera positioned at a predetermined distance close to the animat. World is a singleton class. Animats may be added by the addAnimat method that takes a reference to the constructed animat and coordinates specification of an animat's initial location. External objects may be added to the world by calling addObject with an instance of the item class passed as an argument. External cameras may be added to the world instance by calling addCamera with arguments specifying location, orientation, and frustum of the camera.

[0159] The World class instance may be responsible for updating the state of animated and non-animated objects it contains. This may be performed by calling the update method on these objects. Similarly the update method is called on each non-animated objects. The World singleton object also calls the reset method for animated and no-animated objects it contains in a manner similar to that of update.

[0160] World class itself exposes the update and reset method called internally by the VirtUApplication object which exhibits SystemListener interface and thus is called by the engine on each step to update world's state and the state of its components.

[0161] *Animat Class*

[0162] The Animat class 1804 can serve as the base class for implementation of animats with a desired set of properties. Animat implementation details are described later in this report. In some implementations, another function of the Animat class is to serve as a container for sensory organs and animat control. The Animat class implements the update method from which it calls update for each sensory organ that the animat contains. Sensory organs may be added to the animat during its construction or initialization by calling an addController method that takes OrganController 1806 reference as an argument.

[0163] To control animat behavior, a control object may be created. The control object is animat-specific and depends on the type of outgoing interfaces exposed by the ultimate animat class. For example, walking animats such as a virtual animal expose the walking controls, allowing the caller to specify the direction and speed of walking or running. Vehicle animats expose vehicle controls, allowing the caller to specify acceleration and steering. Control objects are added via an addControl method typically called during animat

initialization or construction, passing a control object reference as an argument. All elements of an animat's behavioral interface may be accessed via its control object.

[0164] In some implementations, another function of the Animat class is to generate and hold a collision shape. The collision shape may be generated in the shape of a capsule surrounding the animat (depending on its dimensions) or as a more tightly fit shape following the animat shape. The collision shape serves as a proximity sensor as well as a chase camera boundary.

[0165] *Vehicle Classes*

[0166] The Vehicle class is derived from the animat class to encapsulate functionality related to vehicle control and vehicle composition. The specific vehicle class may be derived from the Vehicle class. The specific vehicle class may determine the appearance and navigational properties of the vehicle, sets up relevant sensory organs and motor controls, and determines its sensitivity to reward values as well as its collision properties. The Vehicle class may receive and interpret a user's input for control of the animat in no-brainer mode. The user's input may allow the user to accelerate/decelerate and steer the vehicle.

[0167] *Other Animat Classes*

[0168] The Virt-U may also include the CharacterAnimat and specific character classes that implement a walking character whose motion is based entirely on kinematics. A walking character may be controlled by standard GUI commands that allow the user manipulate the direction of motion (forward/backward, left/right turns, jumping).

[0169] *OrganController Class*

[0170] OrganController may serve as the base class for all sensory organs. It may define the update and reset methods 1812 that may be implemented by the derived classes. In some implementations, the purpose of classes that extend OrganController is either to supply sensory information specific to an animat's senses 1808a, 1810a or to contain other sensory organs 1808b, 1810b.

[0171] *OculoMotor Class*

[0172] This class handles visual inputs from single- or multi-camera sources as well as rotation of the cameras. Multi-camera sources are used here to simulate, with minimal distortion, large (>75 degrees) fields of view (FoV) that are attainable with certain real world cameras. In some embodiments, the multi-camera source works by combining inputs from several narrow FoV cameras (or facets) and deliver the result of this operation as a single

input data stream. The OculoMotor class acts as a container for Eye class objects, which in turn contain camera sensor facets.

[0173] *Vestibular Class*

[0174] This class may embody vestibular inputs that record self-motion of the agent. The vestibular sensors may provide the model with linear and angular velocities or accelerations of the agent as would be done if the respective accelerations were measured by accelerometers and gyroscopes of the real robotic device. Realistic noise models can be added to further improve the quality of simulated by Vestibular class real IMUs.

[0175] *Rewards and Collisions*

[0176] In some implementations, rewards and collisions may be accessed through their respective controllers to obtain reward and collision or touch sense sensory data. The update of touch and reward information may happen due to implementation of the CollisionListener interface by their respective animat classes. Collision may be detected via a collision body: an invisible and weightless shape surrounding the animat that serves as a proximity sensor. Collision or detection of an external object by proximity sensor can result in obtaining a reward value if a reward is attached to an object that the animate collides with or senses via proximity sensor. To sense a reward, the animat should have a corresponding reward sense. Sensing the reward information allows reinforcement learning to be implemented as part of the neural model controlling animat behavior. Lastly, a stationary object that extends the Item class can be associated with different rewards types holding a different reward value.

[0177] *TouchSense Class*

[0178] TouchSense abstracts touch/proximity sensors that provide a vector in polar coordinates centered on the agent's center of mass and pointing to the locations of collision (for touch sensor) or close-to-collision (proximity sensor). In some implementations, a difference between touch and proximity is the fitness of the collision surface for the agent. For example, for a touch sensor the collision surface substantially conforms to the shape of the body, whereas a shape surrounding the body with a radius corresponding to the sensor range is created for proximity sensors. In the touch sensor the length of the vector can relate to the strength of an impact, and in the proximity sensor the vector length can relate to the distance between the agent and the edge of an obstacle. In some implementations, multiple collisions can be detected and passed to the model simultaneously.

[0179] *RewardSense Class*

[0180] RewardSense abstracts an arbitrary type of reward in order to model positive or negative reinforcement. In the context of Mars exploration, RewardSense sensors can simulate results coming from chemical analysis of samples collected by the Rover. Any object in the VE can be assigned a reward type and value by the VE designer, and if the animat comes in contact with such an object during the simulation it will receive information about the associated reward type and its value.

[0181] *Other Sensory Organ Classes*

[0182] In some implementations, the Virt-U system can include additional sensors for testing and verification. These included a GPS-like sensor that provides the true coordinates of an agent at any point in time. Such a sensor could be used to verify and tune precision of the internal path integration system.

[0183] *AnimatVehicle Class*

[0184] The AnimatVehicle class provides entry points to control a vehicle that can use the JBullet ray-casting mechanism to mimic suspension of a physical vehicle. This class extends the JMonkey VehicleControl class and implements the CollisionListener interface. The latter constitutes a callback from the JMonkeyEngine that is called to handle collisions. The VehicleAnimat class can act asynchronously compared to the JMonkey game engine clock via the GUI interface, but actions can occur on game engine clock ticks. The update method implemented by the AnimatVehicle Class can execute queued control commands to the vehicle.

[0185] The collision method called back by JMonkeyEngine to handle collision events may allow the AnimatVehicle class to populate the touch or proximity sensor map. The touch map can be retrieved upon request from the TouchSense organ, but actual touch map generation can occur within the AnimatVehicle class. The reward may be updated if a collision body other than the animat has a reward type for which the animat subscribes by having an attached reward sense. Reward information (different from collision information) related to the nature, or quality, of the reward can also be filtered for environment.

[0186] *AnimatCharacter Class*

[0187] The AnimatCharacter class may provide entry points to control a walking character. The class extends JMonkey's PhysicalCollisionObject class that ultimately implements the CollisionListener interface. Actual movement control and collision handling may be implemented in the aggregated AnimatCharacterController object which implements the ActionListener interface. AnimatCharacterController controls movement based on linear

and angular velocities, and supports movement along an animat's own forward/backward and left/right axes as well as world up/down axis plus rotations of heading and elevation. Collisions may make the animat slide along the colliding surface by changing heading and/or elevation. AnimatCharacterController may also use the ghost object and convex sweep test to search for upcoming collisions. This may be combined with discrete collision detection to recover from collisions with objects. Lastly, AnimatCharacterController handles collision and reward information similarly to the manner described above for AnimatVehicle objects.

[0188] *Externally driven rendering context*

[0189] The physical clock that exists in JMonkey engine operates in units of time that corresponds to the physical world where objects have simulated masses, accelerations, torques, forces, and other physical properties in corresponding units. The graphical clock in the engine is synchronized to the physical clock, yet both clocks are designed to run asynchronously based on user input. In the case when neural simulation speed differs significantly between neural simulator and robotic embodiment this may not be appropriate, and the Virt-U system may synchronize the simulation clock with these two engine clocks so that inputs and outputs between the brain model and the animat's body operate on the same clock simulating the real time interactions. To achieve this synchronization, a custom rendering context was created that is driven by the simulation clock. On every tick of the simulation clock the world advances for a time interval that is equal to the length of the model processing cycle.

[0190] *Virt-U Proxy Layers for Cog Ex Machina and INDRA neural simulators.*

[0191] Virt-U proxy layer may expose Virt-U library interface to neural simulation software. Depending on the programming language used in the neural simulator, Virt-U proxy layer may be a collection of Java and/or Scala classes (example implementation for Cog Ex Machina) or a collection of C++ classes with internal Java Native Interface wrappers around Virt-U library calls (example implementation for C++ based INDRA neural simulator).

[0192] In some embodiments, the Virt-U proxy layer for Cog Ex Machina (e.g., see Figure 19) is designed for dynamic update and extraction of designated sensory (1904 to 1900, to 1906 to 1902) and motor control (1902 to 1910 to 1904) dynamic fields in the neural model (1908). Cog Ex Machina adopts a synchronous and entirely deterministic model of computation that may define the computational process as a sequence of steps. Updates occur

substantially simultaneously after execution of each step - this includes sensory information injection and motor control information extraction.

[0193] INDRA based proxy layer adopts an asynchronous model between virtual environment and neural model. This allows a user to test the performance of complicated neural models in real time and simplifies their transition to robotic embodiments. Virt-U proxy layer for INDRA may add time-averaging to the sensory readings for the cases when virtual environment queries sensors more often than the neural model can process the information.

[0194] *Manipulating Virt-U from Neural Application*

[0195] The neural application initializes a VE or World 1914 using a path to the XML description 1912 of the virtual world. This may create neural populations 1908 before the virtual world is instantiated and its sensory organs become accessible like in Cog Ex Machina proxy. This may also follow the opposite order and instantiate virtual world 1904 and sensory/motor buffers 1906 before any neural populations are created as in INDRA proxy. In both cases the subsequent construction within the proxy instantiates the missing parts and links them together. In some cases this linkage can be established by providing pointers to the neural data structures to the respective sensory/motor objects of the proxy 1900. In other cases it can be established by providing callbacks 1906 that execute sensory motor operations on the neural populations.

[0196] The mechanism by which the outside world connects to the neural model is called injection/extraction. Using injection, data can be placed into a neural population; extraction may allow data to be pulled from the neural population to perform a desired behavior.

[0197] *Proxy Controllers*

[0198] The proxy controllers 1900 serve as either input or output controllers between Virt-U proxy layer and Virt-U library. Actual data exchange occurs in the update method. Input controllers may acquire data from Virt-U sensory organs and inject them into virtual sensors. Output controllers may acquire data from virtual actuators and use animat control interfaces to perform the animat's desired behavior.

[0199] DebugWorld example class

```
public class TestMarsWorld extends DebugWorld implements WorldBuilder
{
    //...
    super();
```

```

    app.setDebugCollisionShapes(true);
}
@Override
public void construct()
{
    World.instance.addScene("Scenes/Mars.j3o");
    Rover animat = new Rover("Rover", "Models/RoverParts/RoverBody.j3o");
    World.instance.addAnimat(animat, new Vector3f(0.0f, 93f, 0f), 2.0f*((float)
    Math.PI)/3.0f, 0.0f);
    // add cameras
}

```

[0200] **GUI commands for the Virt-U Builder application**

User Action	System response
Key H	For animals: Set constant speed to move the active animat to the left. Speed remains constant until a conflicting command is issued. For vehicles: Apply steering to the left. Steering is only applied while the key is pressed.
Key K	For animals: Set constant speed to move the active animat to the right. Speed remains constant until a conflicting command is issued. For vehicles: Apply steering to the right. Steering is only applied while the key is pressed.
Key U	For animals: Set constant speed to move the active animat forward. Speed remains constant until a conflicting command is issued. For vehicles: Apply forward acceleration. Releasing the key leaves the speed constant (pending terrain slopes and obstacles).
Key J	For animals: Set constant speed to move the active animat backward. Speed remains constant until a conflicting command is issued. For vehicles: Apply brakes. Releasing the key releases the brake.
Key G	For animals: Set constant speed to move the active animat upwards. Speed remains constant until a conflicting command is issued. For vehicles: No effect.
Key B	For animals: Set constant speed to move the active animat downwards. Speed remains constant until a conflicting command is issued. For vehicles: No effect.
Spacebar Key	For animals: Set all speeds to zero. For vehicles: Jump in the air.
Key comma (,)	For animals: Set constant angular velocity of the active animat rotating it to the left. Velocity remains constant until a conflicting command is issued. For vehicle: No effect.
Key period (.)	For animals: Set constant angular velocity of the active animat rotating it to the right. Velocity remains constant until a conflicting command is issued. For vehicle: No effect.
Key F	For animals: Set constant angular velocity of the active animat rotating its front upward. Velocity remains constant until a conflicting command is issued. For vehicle: No effect.

User Action	System response
Key V	For animals: Set constant angular velocity of the active animat rotating its front downward. Velocity remains constant until a conflicting command is issued. For vehicle: no effect.
Key E	Switch to first person view from the eye of the active animat.
Key	Each press cycles through available eyes of the active animat, making the next eye the source of the first person view.
Keys 0-9	Switch the active external camera-to-camera. Camera 0 is a follow camera that always has the first animat in focus. Cameras 1-9 are optional cameras that are fixed in space. If the world has less than 9 fixed cameras the extra keys have no effect.
Mouse move	No effect if the camera 0 or first person view is active. Rotate the active camera if camera 1-9 is active.
Mouse button press + mouse move	If camera 0 is active, rotate the camera view.
Mouse wheel	Zoom-in/zoom-out with the camera.
Key C	Outputs the current position and direction of the active camera. This is useful for precise camera positioning.
Keys W,S,A,D	Move cameras 1-9 forward, backward, left, and right, respectively.
Keys Q,Z	Move cameras 1-9 upward and downward, respectively
Keys F1-F10	Set the active animat to animat 1-10. If there are fewer than 10 animats the extra keys have no effect.
Key TAB	Saves the current world configuration to an XML file.
Key P	Pauses and unpauses the rendering.
Escape Key	Exits the application.
Key M	Outputs memory usage of the application to the standard output.

Table I.1. GUI commands for the Virt-U Builder application

[0201] List of Virt-U sensors-actuators

- VirtualSensor
- VirtualVectorSensor
- VirtualColorSensor
- VirtualRetina
- VirtualColorRetina
- VirtualScalarFieldActuator
- Virtual0DSensor
- Virtual0DScalarFieldActuator

[0202] Proxy controllers implemented to support simulation during Phase I

- StaticCameraIC – injects color images from static cameras

- **EyesIC** – injects color input images from virtual vision sensory organs;
- **VestibularIC** – injects vestibular data in three 0D scalar fields;
- **RewardsIC** – injects reward values into an array of 0D sensors;
- **TouchIC** – injects touch information into a 2D scalar field representing a touch map in polar coordinates;
- **VehicleLinearSpeedIC** – injects vehicle linear velocity into a 0D scalar field;
- **VehicleAngularVelocityIC** – injects vehicle angular velocity into a 0D scalar field;
- **GlobalPositionIC** – injects GPS data into a 0D vector field;
- **EyesOC** – extracts eye movement information from the actuator field; and
- **DualMotorOC** – enforces vehicle angular and linear velocity based on values extracted from two 0D dynamic scalar fields.

[0203] **Example of Cog Virt-U application (animat wandering a virtual world)**

```
object Wandering extends CogBrowserApp("Wandering around", SimConfig.rate)
{
  val engine = new CogVirtUEngine("assets/Worlds/TestWorld", setupControllers)

  //The rest of the object is your network

  //create sensor and actuator fields
  val otolithX = new Virtual0DSensor("otolith X")
  val otolithZ = new Virtual0DSensor("otolith Z")
  val sscanal = new Virtual0DSensor("sscanal")
  val gps = new VirtualVectorSensor(3,"GPS")
  val somatoSensorSize = 100
  val touch = new VirtualSensor(somatoSensorSize, somatoSensorSize,"touch")
  val rewards = new Virtual0DSensor("comfort")
  val staticCameras = new StaticCameraFields
  val lv = new PresetLVField("LinearVelocity", 0.4f)
  val av = new PresetAVField("AngularVelocity", 30)
  val rmmcf = new RandomMicroMacroSaccadesField("EyeRotation", .01f*SimConfig.rate,
  .01f*SimConfig.rate, 30)
  val eye = new FovealField("Rat","Right Eye")

  //brain model .....

  //callback to initialize sensor/actuator controllers
  def setupControllers {
    val world = World.instance//engine.getWorld
    engine.registerController(new StaticCameraIC( world.cameras(),staticCameras))
    val animat = world.animats().get(0)
    engine.registerController(new EyesIC(animat,eye))
    engine.registerController(new VestibularIC(animat,otolithX,otolithZ,sscanal))
    engine.registerController(new GlobalPositionIC(animat,gps))
    engine.registerController(new RewardsIC(animat,rewards))
    engine.registerController(new TouchIC(animat,touch))
    engine.registerController(new EyesOC(animat,rmmcf.saccadeField))
  }
}
```

```
engine.registerController(new  
DualMotorOC(animat.asInstanceOf[CharacterAnimat],av.avField,lv.lvField, false))  
}  
}
```

[0204] Conclusion

[0205] While various inventive embodiments have been described and illustrated herein, those of ordinary skill in the art will readily envision a variety of other means and/or structures for performing the function and/or obtaining the results and/or one or more of the advantages described herein, and each of such variations and/or modifications is deemed to be within the scope of the inventive embodiments described herein. More generally, those skilled in the art will readily appreciate that all parameters, dimensions, materials, and configurations described herein are meant to be exemplary and that the actual parameters, dimensions, materials, and/or configurations will depend upon the specific application or applications for which the inventive teachings is/are used. Those skilled in the art will recognize, or be able to ascertain using no more than routine experimentation, many equivalents to the specific inventive embodiments described herein. It is, therefore, to be understood that the foregoing embodiments are presented by way of example only and that, within the scope of the appended claims and equivalents thereto, inventive embodiments may be practiced otherwise than as specifically described and claimed. Inventive embodiments of the present disclosure are directed to each individual feature, system, article, material, kit, and/or method described herein. In addition, any combination of two or more such features, systems, articles, materials, kits, and/or methods, if such features, systems, articles, materials, kits, and/or methods are not mutually inconsistent, is included within the inventive scope of the present disclosure.

[0206] The above-described embodiments can be implemented in any of numerous ways. For example, embodiments of designing and making the coupling structures and diffractive optical elements disclosed herein may be implemented using hardware, software or a combination thereof. When implemented in software, the software code can be executed on any suitable processor or collection of processors, whether provided in a single computer or distributed among multiple computers.

[0207] Further, it should be appreciated that a computer may be embodied in any of a number of forms, such as a rack-mounted computer, a desktop computer, a laptop computer, or a tablet computer. Additionally, a computer may be embedded in a device not generally

regarded as a computer but with suitable processing capabilities, including a Personal Digital Assistant (PDA), a smart phone or any other suitable portable or fixed electronic device.

[0208] Also, a computer may have one or more input and output devices. These devices can be used, among other things, to present a user interface. Examples of output devices that can be used to provide a user interface include printers or display screens for visual presentation of output and speakers or other sound generating devices for audible presentation of output. Examples of input devices that can be used for a user interface include keyboards, and pointing devices, such as mice, touch pads, and digitizing tablets. As another example, a computer may receive input information through speech recognition or in other audible format.

[0209] Such computers may be interconnected by one or more networks in any suitable form, including a local area network or a wide area network, such as an enterprise network, and intelligent network (IN) or the Internet. Such networks may be based on any suitable technology and may operate according to any suitable protocol and may include wireless networks, wired networks or fiber optic networks.

[0210] The various methods or processes (e.g., of designing and making the coupling structures and diffractive optical elements disclosed above) outlined herein may be coded as software that is executable on one or more processors that employ any one of a variety of operating systems or platforms. Additionally, such software may be written using any of a number of suitable programming languages and/or programming or scripting tools, and also may be compiled as executable machine language code or intermediate code that is executed on a framework or virtual machine.

[0211] In this respect, various inventive concepts may be embodied as a computer readable storage medium (or multiple computer readable storage media) (e.g., a computer memory, one or more floppy discs, compact discs, optical discs, magnetic tapes, flash memories, circuit configurations in Field Programmable Gate Arrays or other semiconductor devices, or other non-transitory medium or tangible computer storage medium) encoded with one or more programs that, when executed on one or more computers or other processors, perform methods that implement the various embodiments of the invention discussed above. The computer readable medium or media can be transportable, such that the program or programs stored thereon can be loaded onto one or more different computers or other processors to implement various aspects of the present invention as discussed above.

[0212] The terms “program” or “software” are used herein in a generic sense to refer to any type of computer code or set of computer-executable instructions that can be employed to program a computer or other processor to implement various aspects of embodiments as discussed above. Additionally, it should be appreciated that according to one aspect, one or more computer programs that when executed perform methods of the present invention need not reside on a single computer or processor, but may be distributed in a modular fashion amongst a number of different computers or processors to implement various aspects of the present invention.

[0213] Computer-executable instructions may be in many forms, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Typically the functionality of the program modules may be combined or distributed as desired in various embodiments.

[0214] Also, data structures may be stored in computer-readable media in any suitable form. For simplicity of illustration, data structures may be shown to have fields that are related through location in the data structure. Such relationships may likewise be achieved by assigning storage for the fields with locations in a computer-readable medium that convey relationship between the fields. However, any suitable mechanism may be used to establish a relationship between information in fields of a data structure, including through the use of pointers, tags or other mechanisms that establish relationship between data elements.

[0215] Also, various inventive concepts may be embodied as one or more methods, of which an example has been provided. The acts performed as part of the method may be ordered in any suitable way. Accordingly, embodiments may be constructed in which acts are performed in an order different than illustrated, which may include performing some acts simultaneously, even though shown as sequential acts in illustrative embodiments.

[0216] All definitions, as defined and used herein, should be understood to control over dictionary definitions, definitions in documents incorporated by reference, and/or ordinary meanings of the defined terms.

[0217] The indefinite articles “a” and “an,” as used herein in the specification and in the claims, unless clearly indicated to the contrary, should be understood to mean “at least one.”

[0218] The phrase “and/or,” as used herein in the specification and in the claims, should be understood to mean “either or both” of the elements so conjoined, i.e., elements that are

conjunctively present in some cases and disjunctively present in other cases. Multiple elements listed with “and/or” should be construed in the same fashion, i.e., “one or more” of the elements so conjoined. Other elements may optionally be present other than the elements specifically identified by the “and/or” clause, whether related or unrelated to those elements specifically identified. Thus, as a non-limiting example, a reference to “A and/or B”, when used in conjunction with open-ended language such as “comprising” can refer, in one embodiment, to A only (optionally including elements other than B); in another embodiment, to B only (optionally including elements other than A); in yet another embodiment, to both A and B (optionally including other elements); etc.

[0219] As used herein in the specification and in the claims, “or” should be understood to have the same meaning as “and/or” as defined above. For example, when separating items in a list, “or” or “and/or” shall be interpreted as being inclusive, i.e., the inclusion of at least one, but also including more than one, of a number or list of elements, and, optionally, additional unlisted items. Only terms clearly indicated to the contrary, such as “only one of” or “exactly one of,” or, when used in the claims, “consisting of,” will refer to the inclusion of exactly one element of a number or list of elements. In general, the term “or” as used herein shall only be interpreted as indicating exclusive alternatives (i.e. “one or the other but not both”) when preceded by terms of exclusivity, such as “either,” “one of,” “only one of,” or “exactly one of.” “Consisting essentially of,” when used in the claims, shall have its ordinary meaning as used in the field of patent law.

[0220] As used herein in the specification and in the claims, the phrase “at least one,” in reference to a list of one or more elements, should be understood to mean at least one element selected from any one or more of the elements in the list of elements, but not necessarily including at least one of each and every element specifically listed within the list of elements and not excluding any combinations of elements in the list of elements. This definition also allows that elements may optionally be present other than the elements specifically identified within the list of elements to which the phrase “at least one” refers, whether related or unrelated to those elements specifically identified. Thus, as a non-limiting example, “at least one of A and B” (or, equivalently, “at least one of A or B,” or, equivalently “at least one of A and/or B”) can refer, in one embodiment, to at least one, optionally including more than one, A, with no B present (and optionally including elements other than B); in another embodiment, to at least one, optionally including more than one, B, with no A present (and

optionally including elements other than A); in yet another embodiment, to at least one, optionally including more than one, A, and at least one, optionally including more than one, B (and optionally including other elements); etc.

[0221] In the claims, as well as in the specification above, all transitional phrases such as “comprising,” “including,” “carrying,” “having,” “containing,” “involving,” “holding,” “composed of,” and the like are to be understood to be open-ended, i.e., to mean including but not limited to. Only the transitional phrases “consisting of” and “consisting essentially of” shall be closed or semi-closed transitional phrases, respectively, as set forth in the United States Patent Office Manual of Patent Examining Procedures, Section 2111.03.

CLAIMS

1. A method for aggregating data representing a moving body with respect to an environment, the method comprising:
 - (A) receiving, by a processor, a plurality of sets of data readings from at least one sensor, the plurality of sets of data readings representing independent measurements by the at least one sensor of the moving body with respect to the environment;
 - (B) mapping the plurality of sets of data readings received in (A) to respective first sets of data cells, each first set of data cells forming a corresponding first neural representation of the moving body with respect to the environment;
 - (C) applying respective weights to values in the respective first sets of data cells mapped in (B), the respective weights representing respective precisions of the independent measurements represented by the corresponding sets of data readings;
 - (D) combining the values in the first sets of data cells weighted in (C) to form a plurality of aggregated data readings; and
 - (E) mapping the plurality of aggregated data readings formed in (D) to a second set of data cells stored in the memory, the second set of data cells forming a second neural representation of the moving body with respect to the environment that is more precise than the first neural representations of the moving body with respect to the environment.
2. The method of claim 1, wherein (A) comprises receiving data readings representative of an orientation, motion, and/or position of the moving body with respect to the environment.
3. The method of claim 1, wherein (A) comprises processing data from a single sensor to form the first plurality of data readings and the second plurality of data readings by extracting different modalities of data reading from the data from the single sensor.
4. The method of claim 1, wherein (A) comprises receiving the first plurality of data readings from a first sensor and the second plurality of data readings from a second sensor.
5. The method of claim 1, wherein (A) comprises receiving the first plurality of data readings from at least one of a motor, a motor controller, an accelerometer, a gyroscope, an inertial measurement unit, an imaging sensor, a proximity sensor, a photosensor, a sonar, a lidar, a radar, a rangefinder, and a microphone.

6. The method of claim 1, wherein the first plurality of data readings represents at least one of a motor outflow, an optic flow, an inertial measurement unit measurement, a color measurement, a proximity measurement, a luminance measurement, a velocity measurement, an orientation measurement, an acceleration measurement, and a distance measurement.
7. The method of claim 1, wherein (A) comprises:
 - receiving a first set of data readings representing measurements by an inertial measurement unit (IMU) of the moving body;
 - receiving a second set of data readings representing commands issued to a motor of the moving body; and
 - receiving a third set of data readings representing an optic flow detected by an imaging sensor of the moving body.
8. The method of claim 7, wherein (C) comprises:
 - applying a first weight to the first set of data readings, a second weight to the second set of data readings, and a third weight to the third set of data readings;
 - increasing the first weight relative to the second weight and the third weight for data reading indicative of an increasing speed of the moving body;
 - increasing the third weight relative to the first weight and the second weight for data reading indicative of a decreasing speed of the moving body; and
 - varying the second weight from a value of zero for data readings indicative of no motion of the moving body to a value between the first weight and the third weight for data readings indicative of motion of the moving body.
9. The method of claim 8, further comprising:
 - (E) estimating the position and the orientation of the moving body with respect to the environment based on the second neural representation.
10. The method of claim 1, wherein (A) comprises:
 - receiving sensor data from an imaging sensor associated with the moving body;
 - deriving a first set of data readings from the sensor data, the first set of data readings representing at least one edge sensed by the imaging sensor;

deriving a second set of data readings from the sensor data, the second set of data readings representing an estimate of distance between the moving body and at least one object in the environment; and

deriving a third set of data readings from the sensor data, the third set of data readings representing motion of the moving body.

11. The method of claim 10, wherein (C) comprises:

applying a first weight to the first set of data readings, a second weight to the second set of data readings, and a third weight to the third set of data readings;

decreasing the first weight relative to the second weight and the third weight for data readings indicative of low contrast imagery sensed by the imaging sensor;

decreasing the second weight relative to the first weight and the third weight for data readings indicative of changes in depth below a predetermined threshold sensed by the imaging sensor; and

decreasing the third weight relative to the first weight and the second weight for data readings indicative of motion of the moving body.

12. The method of claim 1, wherein (B) comprises:

generating at least one first set of data cells representing respective data points in a measurement range of the at least one sensor, the respective data points representing a continuous portion of the environment;

defining a resolution of the first neural representation corresponding to the at least one first set of data cells based at least in part on a range of a corresponding set of data readings and on a number of data cells in the at least one first set of data cells;

determining a scale of the first neural representation corresponding to the at least one first set of data cells based at least in part on the range of the corresponding set of data readings; and

generating, with the processor, a set of graph connections among the at least one first set of data cells, the set of graph connections providing continuity between representations of the respective data points.

13. The method of claim 1, further comprising:

adjusting data readings received from the at least one sensor based at least in part on the plurality of aggregated data readings.

14. A system for navigating a robot through an environment, the system comprising:
an interface to receive first sensor data from a first sensor and second sensor data from a second sensor, the first sensor data and the second sensor data representing measurements of the environment with respect to the robot;
a memory, operably coupled to the interface, to store the first sensor data in a first neural representation of possible readings from the first sensor and to store the second sensor data in a second neural representation of possible readings from the second sensor; and
a path integration module, operably coupled to the memory, to:
(i) apply a first weight to the first sensor data in the first neural representation, the first weight representing a first measurement precision of the first sensor;
(ii) apply a second weight to the second sensor data in the second neural representation, the second weight representing a first measurement precision of the first sensor;
(iii) combine the first sensor data in the first neural representation with the second sensor data in the second neural representation so as to produce aggregated sensor data in an aggregated neural representation of the environment; and
(iv) estimate a position of the robot with respect to the environment based at least in part on the aggregated neural representation of the environment.
15. The system of claim 14, further comprising:
a map module, operably coupled to the path integration module, to generate a plurality of maps of the environment based on the position of the robot estimated by the path integration module.
16. The system of claim 15, further comprising:
a motivation and goal selection module, operably coupled to the path integration module, to implement a reward drive for the moving the robot and to generate a goal location for the robot based at least in part on the reward drive and the position of the robot estimated by the path integration module.
17. The system of claim 16, further comprising:
a path planning module, operably coupled to the path integration module, the map module, and the motivation and goal selection module, to generate a motor command for the robot based at least in part on the position of the robot estimated by the path integration

module, the plurality of maps of the environment generated by the map module, and the goal location generated by the motivation and goal selection module.

18. The system of claim 17, wherein the robot is configured to perform an action based on the motor command generated by path planning module.

19. The system of claim 16, wherein the reward drive comprises at least one of a curiosity drive, a consumable resource drive, and an external reward drive.

20. The system of claim 19, wherein the curiosity drive is configured to set the goal location for the robot in an unfamiliar area of the environment.

21. The system of claim 19, wherein the resource drive is configured to monitor a consumable resource of the robot and to instruct the robot to seek replenishment for the consumable resource when the consumable resource falls below a predetermined level.

22. The system of claim 19, wherein the external reward drive is configured to set the goal location in an area of the environment area where the robot had previously found an external reward.

23. The system of claim 19, wherein the goal location is based at least in part on a location of a reward object selected by the external reward drive, a location of a replenishment object selected by the resource drive, and a location selected by curiosity drive.

24. A method of providing an estimate of at least one of a position of a moving body and an orientation of the moving body, the method comprising:

(A) receiving a vestibular signal representative of the at least one of the orientation of the moving body and the position of the moving body;

(B) receiving a motor outflow signal representative of commanded motion of the moving body;

(C) receiving an optical flow signal representative of the at least of the orientation of the moving body and the position of the moving body with respect to an environment of the moving body;

(D) weighting the vestibular signal received in (A), the motor outflow signal received in (B), and the optical flow signal received in (C) based on a type of the motion of the moving body;

(E) combining the vestibular signal, motor outflow signal, and optical flow signal weighted in (D) to produce a localization signal; and

(F) generating, via at least one processor, the estimate of the at least one of the position of the moving body and the orientation of the moving body based on the localization signal produced in (E).

25. The method of claim 24, wherein (A) further comprises receiving the vestibular signal from at least one of a gyroscope and an accelerometer.

26. The method of claim 24, wherein (B) further comprises receiving the motor outflow signal from a controller operably coupled to a motor of the moving body.

27. The method of claim 24, wherein (C) further comprises:

(C1) receiving data from at least one of a camera, an infrared sensor, an ultrasonic sensor, a laser rangefinder, and a radio-frequency antenna; and

(C2) deriving the optical flow signal from the data received in (C1).

28. The method of claim 24, wherein (C) further comprises spatiotemporally filtering the optical flow signal.

29. The method of claim 24, wherein (D) further comprises at least one of:

(D1) determining rotation along at least one axis; and

(D2) determining translation along at least one axis.

30. The method of claim 24, wherein (D) further comprises selecting at least one weights based at least in part on at least one of a speed of the moving body, a velocity of the moving body, an acceleration of the moving body, and a trajectory of the moving body.

31. The method of claim 24, wherein the moving body comprises an autonomous vehicle.

32. A method for developing a robust noise-tolerant distributed data representation model of an environment based on data from at least one sensor, the method comprising:

(A) generating, with a processor, at least one set of data cells representing respective data points in a measurement range of the at least one sensor, the respective data points representing a continuous portion of the environment;

(B) defining, with the processor, a resolution of the neural representation corresponding to the at least one set of data cells based at least in part on a range of a corresponding set of data readings and on a number of data cells in the at least one set of data cells;

(C) determining, with the processor, a scale of the neural representation corresponding to the at least one set of data cells based at least in part on the range of the corresponding set of data readings; and

(D) generating, with the processor, a set of graph connections among the at least one set of data cells, the set of graph connections providing continuity between representations of the respective data points.

33. The method of claim 32, wherein the neural representation is first neural representation at a first scale and a first, and further comprising:

generating, with the processor, a second neural representation at a second scale and a second resolution according to steps (A) through (D);

combining, with the processor, outputs of the first neural representation with outputs of the second neural representation to produce a combined neural representation; and

estimating, with the processor, a position of a moving body associated with the at least one sensor based on the combined neural representation at a combined scale equal to least common multiple of the first scale and the second scale and a combined resolution equal to the smaller of the first resolution and the second resolution.




Figure 1C




Figure 1D





Figure 1B









Figure 2.

**Figure 3.**

**Figure 4.**

**Figure 5.**




Figure 6.




Figure 7.





Figure 8.

**Figure 9.**




Figure 10.




Figure 11.






Figure 12.

**Figure 13.**

**Figure 14.**








Figure 15.

1600



1604



1606

Figure 16.

1602




Figure 17.




Figure 18. 1800




Figure 19.




Figure 20




Figure 21