Robot Navigation: A Developmental Approach

Laura M. Grabowski
Michigan State University
Department of Computer Science and Engineering
grabow38@msu.edu

ABSTRACT

Path integration, also known as dead reckoning, is a fundamental navigation strategy utilized by many diverse animal species. This strategy has been of interest for use in robot navigation for some time. Previous computational implementations have provided solutions either through a traditional engineering approach or through evolutionary computation. Although successful, these approaches result in systems that cannot adapt to changes in the environment in real time and are unable to autonomously correct for the error that is inherent in the path integration computation. The current work describes a simulated environment, modeled on Vickerstaff and Di Paolo [9], for developing navigation in an autonomous agent, and progress in establishing the developmental framework.

Keywords

Autonomous mental development, robot navigation, path integration, artificial intelligence.

1. INTRODUCTION

Robot navigation has traditionally centered on engineering the solutions to three fundamental questions: (1) Where am I? (2) Where are other places in relationship to me? (3) How can I get from one place to another? [4] In other words, traditional approaches to navigation for robots have focused on localization, mapping, and path planning. Although often very successful, such engineering approaches have the drawback of lacking flexibility and adaptability, and therefore cannot easily be generalized across different environments or different robots.

Animals are highly skilled navigators, and have evolved diverse spatial representation processes that reflect the varying cognitive demands of navigation tasks [2]. The designers of robot navigation systems have long been inspired and motivated by natural evolution's solutions to the challenges of navigation. The motivation behind the current research is to approach navigation from a developmental perspective (see e.g. [14-16, 18, 22, 26-27]). In animals, navigation itself consists of many different complex strategies, behaviors, and their corresponding representations. The focus of the current work is the strategy known as path integration. The overarching research question at the center of the current study, using path integration as the cognitive task, is how animals use environments to autonomously develop skills, and how that might be used to inform the design of robotic navigation systems.

2. OVERVIEW

2.1 Path Integration

Path integration, also known as dead reckoning, is the continual updating of position relative to a location, based on velocity, temporal, and acceleration information [6]. First postulated by Charles Darwin [1], path integration is fundamental and ubiquitous, operating in many diverse species, both invertebrate and vertebrate, including humans. Path integration functions automatically and constantly whenever the animal moves in continuous space [3, 7]. The path integration system works even in terrain in which landmark cues are absent, or otherwise unreliable. Any egocentric navigation system of this kind has two pitfalls: 1) it must run without interruption as long as the animal is moving, and 2) the system is inherently susceptible to cumulative error [12].

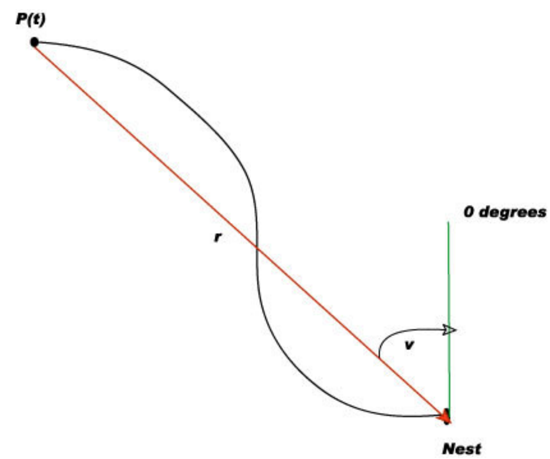

Figure 1. Path integration, or dead reckoning. The animal's position relative to its starting point (nest) at time t, P(t), is given by the vector (v, r).

2.2 The Developmental Approach

Founded in human cognitive development and supported by ongoing research in neuroscience, the developmental approach to artificial intelligence (AI) seeks to realize an automated learning mode that closely mirrors the way in which humans and animals learn: "automated animal-like learning" [13]. In this learning mode, interaction with the environment is continuous, dynamic, on-line, performed in real time, and input is rich and often multi-modal [13, 16-17].

The framework of autonomous mental development includes both intrinsic and developmental capabilities. In a natural brain or an embodied artificial system, the intrinsic developmental program--

encoded genetically in a natural system and written by the programmer in an artificial system--controls the development of cognitive capabilities through autonomous, real-time interactions with both the external world and its own internal world, through its sensors and effectors.

Unlike previous approaches, autonomous mental development (AMD) is not task-specific. Instead, AMD utilizes a task-independent paradigm, characterized by the following steps:

- Design a robot body according to the ecological conditions in which the robot will function.
- Design the developmental program.
- The robot begins executing the developmental program at "birth."
- Humans interact with the robot in real-time, to develop the robot's brain [15, 22].

In this way, the developmental robot will "grow" from infancy to adulthood, learning much as Turing envisioned the education of the "child" machine [8]. Utilizing this approach, robots should be able to learn any task, and the programmer does not need to know in advance what the task will be, nor have any domain knowledge about the end task. The human programmer does not code any task-specific information or actions; rather, these are learned through the operation of the developmental program [15, 22].

This is a major departure from traditional, symbolic approaches to AI. As noted earlier, within the task-specific paradigm, the internal representations of the system are predetermined by the human designers. The automatic generation of self-organizing internal representations is pivotal to AMD. Traditional, symbolic AI representations are world centered, i.e., the items of the representation correspond to a world concept. Each component of the representation has a predefined meaning in relation to the exterior world, using unique variables for each attribute of the object, necessitating a unique representation for every single physical object in the environment. From this, it is clear that this idea is neither scalable nor capable of coping with the richness and ambiguity of human environments. Instead of world centered representations, AMD utilizes distributed body centered representations. Instead of being an atomic, human-designed representation, a distributed body centered representation is built from the body's sensors and effectors, and the generated representation is distributed over different areas, much as representations may be distributed over different cortical areas in a living brain. In this way, a developmental robot generates representations that are more general than hand-designed world centered representations, in the sense that the representations are not unique for a single object. Correspondingly, generated actions are not unique given different sensory inputs, even of the same object [15]. The net result is flexibility, enabling the developmental robot to function in unknown and dynamic environments.

The theoretical grounding of the developmental approach draws strongly from the fields of neuroscience and psychology. From the neuroscience research, emphasis is placed particularly on the information on neural plasticity; in psychology, the emphasis is on developmental psychology and animal learning. Taken together, these areas provide the basis for development's focus on experience and learning.

2.3 The Simulation Platform

The BeaconWorld simulation is modeled in part after Vickerstaff and Di Paolo [9], and is programmed in MATLAB. The goal of the simulation is to provide a simple virtual environment for the development of path integration under the influence of environmental beacons. A beacon is simply some feature of the environment that is somehow salient; in a natural environment, this could be a distinctive location or object, such as the nest, a potential landmark, or a food source. Vickerstaff and Di Paolo approached this problem using genetic algorithms to evolve neural network models of path integration, emphasizing the reproduction of the behavior of Saharan desert ants, *Cataglyphis fortis* Their model successfully evolved a bicomponent model of path integration (see [5]) when the neurons were given properties that allowed multiplication by way of changes in synaptic strength [9].

2.3.1 The Environment

The environment of the current simulation is an unbounded 2-D plane, with the home (or nest) location always at the origin, (0, 0). Each trial presents the agent with a varying number of beacons. The number and locations of the beacons may be either generated randomly or given by the experimenter. The world state is maintained in a structure, and the world information is used to produce the visualization in the simulation. Figure 2 shows an example of a BeaconWorld world with an agent. The home location is at (0, 0), shown by the green star ('*'), and the beacon locations are shown by the blue squares. The agent appears as the red filled circle, with its global orientation indicated by the needle.

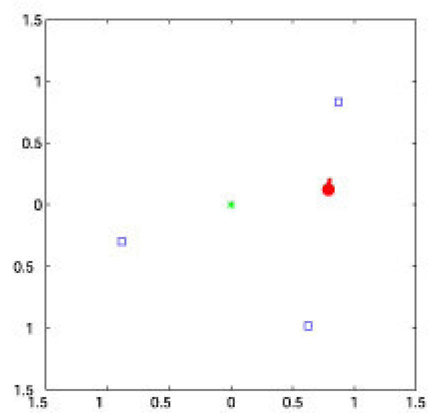

Figure 2. Example world with agent.

2.3.2 The Agent

The overall design of the agent's sensors is modeled on abstractions of vision, as a simplification of simulating synthetic visual input for the agent. Parts of the sensor design were based on Vickerstaff and Di Paolo [9], and other parts were added or modified, either as biologically motivated improvements over the model study [9] or to modify the sensory system to better accommodate a developmental approach.

The agent's state is maintained in a structure. The agent's body is modeled as bilaterally symmetric, with symmetrically paired sensors. The beacon and compass sensors have positive cosine-shaped activation functions, so that no sensor would give a negative value for its activation. The beacon power, home, speed, and food sensors are single-copy components. Table 1 gives a summary of the agent's sensor data fields and their contents.

Table 1. Agent Structure Sensor Data Fields

Field Name	Contents
loc	Current location of agent, in polar coordinates
thetainit	Initial orientation of agent, in radians
bsens.left bsens.right	Left/right beacon sensor output
alcomp.left alcomp.right	Left/right allothetic (global) compass output
idcomp.left idcomp.right	Left/right idiothetic (egocentric) compass output
bpower	Pre-attentive beacon power sensor output
bnovrwd	Beacon "novelty" reward
speed	Agent's speed (not yet implemented)
food	Whether or not the agent has found a food item (not yet implemented)
home	Whether or not the agent is at the home location (currently implemented in simplified form)

Each beacon and compass sensor respond maximally to a particular orientation of the agent, with each sensor pair having complementary values such that if one sensor responds maximally at $\theta_{Agent} = A$, the other responds maximally at $\theta_{Agent} = -A$. The beacon sensors activations, modeled on Vickerstaff and Di Paolo [9] are defined as

$left = [\cos(\theta_B - \pi/2)/2] + 0.05, \quad right = [\cos(\theta_B + \pi/2)/2] + 0.05,$

where θ_B is the current angle of the beacon to the agent's central body axis. The agent attends to only one beacon at a time, the beacon selected by the "pre-attentive" beacon power sensor, which uses both beacon distance and beacon size to select the beacon of highest "interest" to the agent. Closer beacons are of more interest than those that are farther away, but a larger beacon is more interesting than a smaller one when the distances to them are similar.

The allothetic (global) compass sensors activations, modeled on Vickerstaff and Di Paolo [9] are defined as

$left = [\cos(\theta_A - \pi/4)/2] + 0.05, \quad right = [\cos(\theta_A + \pi/4)/2] + 0.05,$

where θ_A is the agent's current global orientation in radians, in $-\pi$ to π. The allothetic compass sensors activations could be thought of as the response of visual or light sensors to a stationary sun at infinity in the east (i.e., over the positive x-axis).

The idiothetic (egocentric) compass sensors were added as a biologically-motivated enhancement over Vickerstaff and Di Paolo's [9] design, since animals generally have multiple, redundant compass systems that both interact and reinforce each other and function independently when conditions render one compass system unuseable [11]. The idiothetic compass was designed to model a vestibular sense of direction, independent of the visual system. This sense of direction can be defined arbitrarily in any environment, since its reference direction is set as the initial orientation of the agent at the beginning of an outbound journey, and is not dependent on maintaining global references. The idiothetic (egocentric) compass sensors activations are defined as

$left = [\cos(\theta_I - p_{left})/2] + 0.05, \quad right = [\cos(\theta_I + p_{right})/2] + 0.05,$

where θ_I = *(Current global orientation) - (Initial orientation),* and p_{left} and p_{right} are the left and right sensor's preferred directions, respectively, with p_{left} = *(Initial orientation)* + $\pi/4$, and p_{right} = *(Initial orientation)* - $\pi/4$.

The home sensor is implemented in a simplified form, with an activation value of 1 if the agent is at a distance of less than 0.01 from the home location, or 0 otherwise. The food and speed sensors are not implemented for this phase of the study.

2.3.3 Effector and Control Model

The agent models a two-wheeled robot, with control of the robot being a direct mapping from the wheel velocities to motion (the robot is assumed to have low inertia). The input to the two rotation motors varies according to the cognitive design of the agent. The agent's design included a forward motor, which was not implemented; the agent's motion was effectively modeled using only the two wheel motors.

2.3.4 Developmental Framework

Before any learning or development can take place, the agent needs to have a foundation for "intelligence," a framework for development. This framework includes giving the agent a sensory mapping, cognitive mapping, and a value system. To date, much of this framework has been established. In order to build and test the framework, the agent is being trained to visit the beacons in the environment. Although this task is not part of the path integration problem *per se*, the developmental paradigm will allow for generalization to the larger problems of developing path integration once the developmental framework is established.

The sensory and cognitive mappings for the agent are performed by the Multi-layer In-place Learning Network (MILN) [20-21], using Lobe Component Analysis (LCA) [23]. In order for the agent to learn from its environment, it must have a sensory mapping--an effective internal representation of the environment, using a limited storage capacity. In the context of an artificial agent and machine learning, the continuous state space needs to be discretized, exploiting repeated or similar structure. Lobe Component Analysis (LCA) [23] was used for this discretization. The goal of the LCA algorithm is to learn an optimal representation of a set of input samples with a much smaller set of representation vectors (the lobe component). LCA is an "in-place" algorithm, since it operates incrementally, it does not store higher-order statistics (such as covariance matrix), and the network develops and learns as a side effect of competitive interactions, instead of using a separate developer to learn [21, 23]. LCA is used to transform the continuous space of the beacon sensors and the rotation motors into discrete neurons. These two types of neurons are then used to produce a discrete state space,

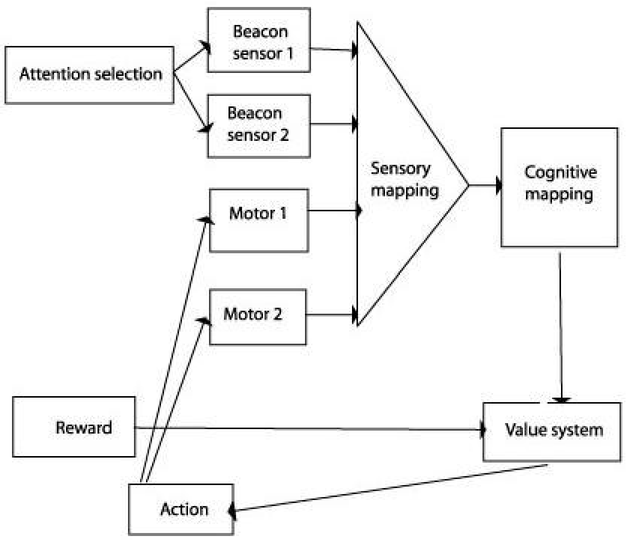

Figure 2. Developmental framework

with each state a combination of sensor (beacon sensor) and motor (rotation motor) neurons.. The state space is internal, i.e., it is in the agent's "brain," and is automatically generated from the agent's sensorimotor experience.

The cognitive mapping for the agent is accomplished by the Multi-layer In-place Learning Network (MILN). This mapping is essentially a Markov model, with one-step, or one sensor/effector frame, prediction. Here, prediction is framed as a regression problem. The MILN takes as input the discrete internal states as described above, using the top responding sensor and effector neurons as that state. The output is all possible next states; the agent learns through experience how to update the weights, or transition probabilities, to the next state. In this way, the agent's experience gives what states are most probable given the current state, leading to the prediction of what will happen next. The current value system is based on "novelty" rewards: the agent receives a reward when it reaches a beacon (through the beacon novelty reward sensor); after reaching a beacon, the agent becomes disinterested in that beacon for a period of time (i.e., the beacon is no longer novel), so the agent will not attend to that beacon. After the period of disinterest has passed, the beacon regains its original novelty, and the agent can again be interested by that previously visited beacon.

3. FUTURE WORK

The anticipated future directions of this study include several different components. The first element is the completion of the current phase, learning beacon homing using reinforcement learning. In this phase, the agent will learn how to use its sensory information and corresponding actions to earn rewards. During this process, the weights (transition probabilities) of the MILN will be trained using Q-learning [10].

The agent needs to have information about direction and distance for the path integration computation; the result of that computation is fundamentally an estimate of the bearing and distance to the goal location. The concepts of direction and distance will therefore need to be developed. Direction is potentially relatively straightforward, since the agent can make use of external directional references (i.e., the allothetic compass sensors). This will enable learning mappings from sensory experience. Distance is likely to prove very challenging, since there is no direct external reference. The agent will either have to develop an internal "distance sense," analogous to a sense of self-motion, or it will need to learn to map other sensory input, such as vision, to a sense of distance. Which of these directions will prove more advantageous has yet to be determined. The culmination of this research would be for the agent to integrate the direction and distance concepts in full-fledged path integration.

4. CONCLUSION

The useful nature of path integration is apparent by its near ubiquity in the animal world. It is a remarkably sophisticated behavior, yet it is performed effortlessly even by animals with relatively simple nervous systems. Path integration poses a great challenge to the developmental approach, and as with most challenges, offers tantalizing rewards. With its requirement for the higher-level senses of direction and distance, the path integration problem offers an opportunity for the developmental approach to break new ground in more complex concepts, perhaps moving one small step closer to the elusive goals of artificial intelligence.

5. REFERENCES

[1] Darwin, C. *On the origin of species*. Hill and Wang, NewYork, 1872/1979.

[2] Dyer, F.C. Spatial cognition: lessons from central-place foraging insects. in Balda, R.P., Pepperberg, I.M. and Kamil, A.C. eds. *Animal cognition in nature: the convergence of psychology and biology in laboratory and field*, Academic Press, San Diego, 1998, 119-154.

[3] Etienne, A.S. and Jeffery, K.J. Path integration in mammals. *Hippocampus*, 14. 180-192, 2004.

[4] Franz, M.O. and Mallot, H.A. Biomimetic robot navigation. *Robotics and Autonomous Systems*, *30*. 133-153, 2000.

[5] Mittelstaedt, H. Analytical cybernetics of spider navigation. in Barth, F.G. ed. *Neurobiology of arachnids*, Springer, Berlin, 1985, 298-316, 1985.

[6] Roche, R.A.P., Mangaong, M.A., Commins, S. and O'Mara, S.M. Hippocampal contributions to neurocognitive mapping in humans: a new model. *Hippocampus*, *15*. 622-641, 2005.

[7] Shettleworth, S.J. *Cognition, evolution, and behavior*. Oxford University Press, New York, 1998.

[8] Turing, A.M. Computing machinery and intelligence. *Mind*, *49*. 433-460, 1950.

[9] Vickerstaff, R. and Di Paolo, E.A. Evolving neural models of path integration. *Journal of Experimental Biology*, *208*. 3349-3366, 2005.

[10] Watkins, C.J.C.H. and Dayan, P. Q-Learning. *Machine Learning*, *8*. 279-292, 1992.

[11] Wehner, R. The ant's celestial compass system: spectral and polarization channels. in Lehrer, M. ed. *Orientation and communication in arthropods*, Birkhauser Verlag, Basel, Switzerland, 1997, 145-185, 1997.

[12] Wehner, R., Michel, B. and Antonsen, P. Visual navigation in insects: coupling of egocentric and geocentric information. *Journal of Experimental Biology*, *199*. 129-140, 1996.

[13] Weng, J., The developmental approach to intelligent robots. in *1998 AAAI Spring Symposium Series, Integrating Robot Research: Taking the Next Leap*, (Stanford University, 1998).

[14] Weng, J. Developmental robots: theory and experiments. *International Journal of Humanoid Robotics*, *1* (2). 199-236, 2004.

[15] Weng, J., A theory for mentally developing robots. in *2nd International Conference on Development and Learning*, (Cambridge, MA, 2002), IEEE Computer Society.

[16] Weng, J. Visual communication and information processing. in *Learning in computer vision and beyond: development*, Marcel Dekker Publisher, New York, 1999.

[17] Weng, J., Evans, C.H., Hwang, W.S. and Lee, Y.-B., The developmental approach to artificial intelligence: concepts, developmental algorithms and experimental results. in *NSF Design and Manufacturing Grantees Conference*, (Long Beach, CA, 1999).

[18] Weng, J., Hwang, W.S., Zhang, Y. and Evans, C.H., Developmental robots: theory, method, and experimental results. in *2nd International Symposium on Humanoid Robots*, (Tokyo, 1999), 57-64, 1999.

[19] Weng, J., Hwang, W.S., Zhang, Y., Yang, C. and Smith, R.J., Developmental humanoids: humanoids that develop skills automatically. in *First IEEE Conference on Humanoid Robots*, (2000), MIT Press.

[20] Weng, J., Lu, H., Luwang, T. and Xue, X., In-Place Learning for Positional and Scale Invariance. in *IEEE World Congress on Computational Intelligence , International Joint Conference on Neural Networks*, (Vancouver, BC, Canada, 2006).

[21] Weng, J. and Luciw, M.D., Optimal In-Place Self-Organization for Cortical Development: Limited Cells, Sparse Coding and Cortical Topography. in *5th International Conference on Development and Learning*, (Bloomgton, IN, 2006).

[22] Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M. and Thelen, E. Autonomous mental development by robots and animals. *Science*, *291*. 599-600, 2001.

[23] Weng, J. and Zhang, N., Optimal In-Place Learning and the Lobe Component Analysis. in *IEEE World Congress on Computational Intelligence , International Joint Conference on Neural Network*, (Vancouver, BC, Canada, 2006).

[24] Weng, J. and Zhang, Y., Developmental robots --- a new paradigm. in *Second International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, (Edinburgh, 2002), 163-174.

[25] Zeng, S., Cherba, D.M. and Weng, J., Dav developmental humanoid: the control architecture and body. in *IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, (2003), 974-980.

[26] Zeng, S. and Weng, J., Obstacle avoidance through incremental learning with attention selection. in *IEEE International Conference on Robotics and Automation*, (2004), 115-121.

[27] Zeng, S. and Weng, J. Online-learning and attention-based approach to obstacle avoidance using a range finder. *Journal of Intelligent and Robotic Systems*, *43* (2), 2005.