A VISION-BASED PERCEPTUAL LEARNING SYSTEM FOR AUTONOMOUS MOBILE ROBOT

[1]KOMAL SHARMA, [2]MANISH MOTGHARE, [3]POONAM AGRAWAL, [4]MONIKA JAISWAL

[1]Priyadarshini Institute of Engineering and Technology, Nagpur, India
[2,3] Priyadarshini Indira Gandhi College of Engineering, Nagpur, India
[4] Tulsiramji Gaikwad Patil College of Engineering, Nagpur, India
E-mail: [1]komalmsharma@gmail.com, [2]motgharemanish@gmail.com, [3]apoonamagrawal@gmail.com, [4]monikajais11@gmail.com

Abstract- Autonomous robots are intelligent machines capable of performing tasks in the real world without explicit human control for extended periods of time. A high degree of autonomy is particularly desirable in fields where robots can replace human workers, such as state-of-the practice video surveillance system and space exploration. However, not having human's sophisticated sensing and control system, two broad open problems in autonomous robot systems are the perceptual discrepancy problem, that is, there is no guarantee that the robot sensing system can recognize or detect objects defined by a human designer, and the autonomous control problem, that is, how the robots can operate in unstructured environments without continuous human guidance. As a result, autonomous robot systems should have their own ways to acquire percepts and control by learning.In this paper, a computer vision system is used for visual percept acquisition and a working memory toolkit is used for robot autonomous control. Natural images contain statistical regularities which can set objects apart from each other and from random noise. For an object to be recognized in a given image, it is often necessary to segment the image into non-overlapping but meaningful regions whose union is the entire image. Therefore, a biologically based percept acquisition system is developed to build an efficient low-level abstraction of real-world data into percepts. Perception in animals is strongly related to the type of behavior they perform. To solve how the robots can learn to autonomously control their behavior based on percepts they've acquired, the computer vision system is integrated with a software package called the Working Memory Toolkit (WMtk) for decision making and learning.

Keywords- Robotics; perceptual learning; minimum spanning tree; Wmtk; Autonomous robot

I. INTRODUCTION

Robotics is a fascinating research topic, and for centuries people have been interested in building machines that mimic living beings. After four robot design generations, from the early days of the application of control theory to complex systems, followed by the sense-think-act cycle of the functional approach, and then to the Brooks' behavior-based robotics, and finally to the recent cognitive robotics, different branches, such as industrial robots, mobile robots, and humanoid robots, have come to a remarkable level of maturity as evidenced in the literature. However, to build and realize such artificial intelligent creatures, capable of moving and operating autonomously so as to replace trained individuals in a variety of applications, still remains as the ever-existing and continuous desire of mankind. As an absolute prerequisite for the design of such robots, human-like intelligence refers to the ability to self-organize, to learn and adapt to varying, complex, dynamic, and uncertain environments, and to interact with the environment and perform tasks efficiently and safely through using their on-board reasoning/problem solving capabilities. To imitate humans, any intelligent robot should be made up of elements such as sensing, perception, cognition, planning, control and actuation [1]

Perception is an awareness of things through the physical senses, especially, vision. Perceptual learning is the ability to construct compact representations of sensory events based on their statistical properties in the perceptual level as opposed to the behavior or cognitive level [2][3]. However, not having a human's sophisticated vision system, the robot perceives the world differently from the human, giving rise to the problem of perceptual discrepancy. Therefore, there is no guarantee that the robot vision system can recognize or detect objects defined by human designers as easily as well. As a result, to operate in realistic environments and to recognize objects in a given image, a robot must have the ability to form percepts on its own, either learned 2 from interacting with a supervisor or figured out by natual association among features of sensory information and, based on that, to segment the image into non-overlapping but meaningful regions whose union is the entire image [4]. In general, by meaningful, the regions in the segmented image should be homogeneous and should have well-defined boundaries, denoting reliable objects. With high spatial frequencies being filtered out, image segmentation reduces the amount of storage and makes behavior resistant to the loss of information. Our motivation is to develop an object recognition vision system based on a set of features associated with each object in the image to assist in the analysis of data that may be applied successfully across different application domains [5].

Perceptual systems evolved to improve the performance of the motor system, however, not the other way around [6]. For example, given several sources of sensory stimuli present trying to attract attention in different directions, the agent may have

to choose between possible courses of action and, when one has been decided on, must prevent any irrelevant input from interfering with the behavior. Therefore, perception and action are tightly coupled and any basic behavior defines a category of action in response to a specific sensory input. In this process, learning plays a major part, "enabling animals to pass the motivational properties of innately reinforcing stimuli to accompanying or perceptibly related stimuli, making attention originally directed only to an innate reinforce direct to percepts associated with the reinforce, and allowing initially neutral objects to arouse a motivational state and links the motivation system to perceptual systems" [7].

For instance, to survive, animals must identify sources of nourishment and dangerous situations, and respond accordingly by approaching the food and fleeing the danger. It is in this sense that behavior learning is the process that forms association between perception and action, established by reinforcement in the form of reward and punishment. For higher animals such as humans, evidence suggests that the occurrence of a response depends on its predicted outcome, so that behaviors are planned on the basis of future benefits rather than present contingencies alone [8]. If the outcome of an action satisfies this expectation, an innate reward is issued, facilitating the 3 intended behavior. This kind of learning is referred to as reinforcement learning in the machine learning community and has been applied to many robotics tasks [9].

II. VISUAL PERCEPT ACQUISITION

Eighty percent of our perceived information about the external world reaches us by way of the eyes. As a result, vision is the primary sensory modality. Discussions of vision often consist of three distinct stages: sensation, perception, and recognition. As the initial stage of sensory transduction, sensation begins with photon receptors in the retina. Perception is the active process of selecting, organizing, interpreting the information detected by sensors, and transforming these raw sensory signals into distinct percepts. As certain categories in the mind, percepts are the brain's internal representation of specific external visual stimuli and are the results of complicated interactions between multiple visual areas.

The goal of perception is to allow animals to recognize things through the process of assembling sensory information into a useful and reliable representation of the world. Natural images contain statistical regularities which distinguish objects from each other and from random noise. For successful visual recognition, each object must have attributes that can be used to differentiate it from others and, based on that, to segment the whole image into meaningful objects. How does the brain represent a visual stimulus internally? Let's briefly look at the human visual system.

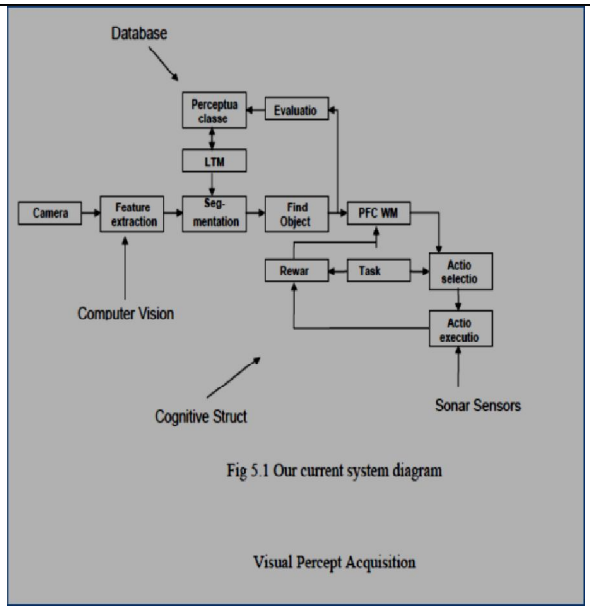


Figure: Visual Percept Acquisition

III. MINIMUM SPANNING TREE

The minimum spanning tree method is a graph analysis of arbitrary point sets of data. In a graph, two points can be connected by either a direct edge or a sequence of edges called a path. A loop in a graph is a closed path. A connected graph has one or more paths between any pair of points. A tree is a connected graph without closed loops. A spanning tree is a tree that contains every point in the data set. If a value is assigned to each edge in the tree, the tree is called a weighted tree. For example, the weights for each edge can be the distance between the two points. The weight of a tree is the total sum of edge weights in the tree. The minimum spanning tree (MST) is the spanning tree that has the minimal total weight among all possible spanning trees for the data set. The minimum spanning tree has the following property that can be used for clustering if the weight associated with each edge denotes the distance between the two points. That is, the weight associated with every edge in the minimum spanning tree will be the shortest distance between two subtrees that are connected by that edge. Therefore, removal of the longest edge will theoretically result in a two-cluster grouping. Removal of the next longest edge will result in a three-cluster grouping, and so on. These correspond to choosing breaks where maximum weights occur in the sorted edges. Two greedy algorithms are usually used to build the minimum spanning tree: Kruskal's algorithm and Prim's algorithm [33]. In this research, Prim's algorithm is used. The construction of the minimum spanning tree starts with some root node s and the tree T greedily grows from s outward. At each step, among all the edges between the nodes in the tree T and those not in the tree yet, the node and the edge associated with the smallest weight to the tree is added to the tree. This is an O(N2) problem. When the tree is built, after

sorting the edge in decreasing order, the edges can be cut to form clusters.

IV. WORKING MEMORY TOOLKIT (WMTK)

This working memory system learns to identify informational chunks worthy of retention using a model of the interactions between the brain's dopamine (DA) system and the PFC. The neuro-computational model for working memory is based on reinforcement learning, in which learning about stimuli or actions is solely based on the basis of rewards and punishments associated with them [66]. Unlike supervised learning, reinforcement learning is minimally supervised because what actions to take in particular situations is not told explicitly but must be worked out on the basis of the reinforcement given. Though the ability of animals to learn appropriate actions in response to particular stimuli on the basis of associated rewards or punishments is a focus of behavioral psychology, there exists evidence that the midbrain dopaminergic system is associated with reward learning. There are two broad classes of reinforcement learning.

In the first case, the reinforcement is delivered immediately after the action is taken. This makes learning relatively easy. In the second case, the reward or punishment is dependent on an entire sequence of actions and, therefore, is partially or wholly delayed until the sequence is completed. This makes learning more difficult since learning the appropriate action at each step in the sequence has to be based on future expectation of reward. However, this latter case is what we are most interested in. In the following, the temporal 32 difference (TD) learning algorithm, which is actually used in the WMtk, is given.

V. AUTONOMOUS NAVIGATION

The key word, autonomous, used in the title reflects upon the development of robotics systems that are more and more based on somewhat human-like attributes such as perception and cognition under the requirement that those systems operate in natural environments and perform goal-directed tasks successfully with minimal supervisory communication from a human. As we move deeper into the area of autonomous robotic systems, major areas of robotic research include the understanding of how humans robustly perform tasks such as vision, perception, cognition, thinking, reasoning, speech understanding, pattern recognition, decision-making, reasoning and control. In this paper, some basic cognitive aspects of path planning and navigation are introduced, followed by some approaches adopted in mobile robotic navigation research, which guide our design of intelligent sensory and control mechanisms. Finally, some landmark based scene recognition methods are explored path planning and navigation.

VI. A VISION-BASED AUTONOMOUS MOBILE ROBOT SYSTEM

The primary focus of our research is on computer vision-based perception and its interaction with an adaptive working memory system on a robot to facilitate developmental robotics. This research is conducted through implementing the methodologies outlined in II and III on a mobile robot platform and applying the system to mobile robot navigation. Since experienced-based learning is widely adopted by human beings, primates and many other animals, a major goal in robotics has been the development of robots that employ a natural process of learning through past experiences. There is biological evidence that adaptive working memory (WM) structures, existing in primate brains, are important to the learning and performing of tasks by focusing on the features most relevant to the current task and by providing the embodiment necessary for exploring the issues of task learning through accumulating rewards with regard to these features. However, an adaptive working memory cannot function in isolation. It must be part of a realistic system involving perception, actuators, reasoning, and short- and long-term memory structures. Toward this goal, a system is proposed in Fig 5.1 that is composed of working memory (cognitive structure), long term memory (database), and perceptual modules (computer vision). The center of this system is PFC WM. It communicates directly with other modules, receiving data from the perceptual system. The attentional focus of the robot is realized by the limited storage capacity of WM which has some means of deciding which information should enter the working memory and which should be removed, and holds the objects currently within the attention of the robot. However, without a human's sophisticated visual system, the robot perceives the world differently from a human. To solve this perceptual discrepancy problem, the first stage in the development of our vision-based 48 autonomous mobile robot system is to let it acquire percepts on its own by way of employing local color and texture features.

VII. COLOR HISTOGRAM

Database Cognitive Strict Feature Perceptuaclasse Find Evaluation PFC WM Task ActioselectioRewarActioexecutio extraction Object Segmentation Camera LTM Computer Vision Sonar Sensors 49 in the human retina: red cones, green cones and blue cones. The responses from the three different cones are compared to allow color discrimination. One of the most important functions of the visual system is its ability to recognize an object under a variety of different viewing conditions. For example, during a single day, the spectral content of daylight changes significantly. However, surfaces

and objects retain their perceived color despite the changes in the wavelength and energy composition of the light reflected from them. This phenomenon is called color constancy and is displayed in a wide range of species, including humans. Wavelength-selective cells whose responses seem to correlate with the human perception of colors have been reported in LGN, V1, V2 [95]. In V4 and beyond, it has been reported that there is no pre-specified wavelength composition that leads to a color and to that color alone and that V4 is concerned with automatic color constancy operations [96].

$$H=\begin{cases}(0+\dfrac{G-B}{MAX-MIN})\times 60 & if \quad R=MAX,\\ (2+\dfrac{B-R}{MAX-MIN})\times 60 & if \quad G=MAX,\\ (4+\dfrac{R-G}{MAX-MIN})\times 60 & if \quad B=MAX,\end{cases} \quad S=\frac{MAX-MIN}{MAX}, \quad V=MAX \qquad (5.1)$$

To follow this process, the images and videos taken from a camcorder are stored for each frame as a RGB (red, green, blue) color image. Next, the RGB color images are converted to the HSV (Hue Saturation Value) color space for use [97]. The Hue describes each color by a normalized number in the range from 0 to 1 starting at red and cycling through yellow, green, cyan, blue, magenta, and back to red. The Saturation describes the vibrancy of the color and represents the purity of a color such as the "redness" of red. The less saturation in a color, the more pale it looks (washed out). The Value describes the brightness of the color. For normalized RGB values in the ranges from 0 to 1, the conversion to HSV is done in the following manner: where MAX is the maximum value of (R, G, B), and MIN is the minimum. From the above formulas, it can be seen that, if MAX = MIN , H is undefined and S = 0, there is no hue and the 50 color lies along the central line of grays; if MAX = 0, V = 0 and S is undefined, the color is pure black and there is no hue, saturation and value.

Database tree

To build a memory network, in our approach, a tree-structured vector quantization method is used to generate a 3-way approximate nearest neighbor search tree. Given the obtained training set, a tree root node exists at the first level which includes all the indices to the database feature elements. Then an initial set of 3 representative patterns, called centers, are randomly selected from the elements of the root as its children cluster centers (or tree nodes) at the second level, and the whole set of the elements is clustered into three subsets by assigning each feature vector to its closest representative cluster center according to the mixed distance measure. At the third level, for each of the 3 clusters obtained at the second level, three feature vectors are selected randomly as its children cluster centers and elements are clustered correspondingly, resulting in 9 cluster tree nodes at the level. This procedure continues until either all the elements in a node belong to the same object class (a pure node) or the number of the elements in a node is below some limit, say, one hundred. Every feature vector in the leaf nodes has a percept landmark associated with it.

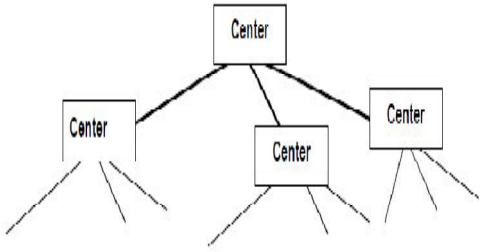


Fig 5.3 Illustration of database tree structure

Next, to search through the tree, given a new feature vector, its distances to the randomly chosen cluster centers at the second level are calculated and the winner is the center that the feature vector is nearest to. At the third level, the feature vector's distances with the three children centers of the winner at the second level are computed and the corresponding new winner is selected. This procedure continues until coming to a leaf node. If it is pure, then stop. Otherwise, do a nearest-neighbor search, and the winner is the training feature vector which gives minimum distance according to the chosen metric.

Integration of the WMtk into Robot Control

In use of the WMtk to build a robotic working memory system, a WorkingMemory object is created and configured to hold a list of candidate blob chunks that will be delivered by the perceptual modules. In order to evaluate the retention of a chunk or a combination of chunks, after an action is taken by the robot toward a goal, a prediction of task success in the form of a reward is passed to the WorkingMemory together with the candidate blob chunk list. Then the WorkingMemory object uses the TD learning algorithm to learn which chunk or chunk combination to retain and which to discard based on the reward. Next, the robot proceeds with another sensing, acting and reward delivering cycle. In the toolkit, the structure of chunks is nothing more than a pointer to a memory location and a string that will be defined by the user to denote the chunk type. As a result, there is no way that 57 the WMtk can automatically extract meaningful features from the candidate chunks. To solve real world problems, the WMtk allows the user to define four functions, one to map any chunk into a real-valued feature vector that the adaptive critic uses to assess the value of the chunk, one to map the robot's current situation into a real-valued state vector, one to return a scalar reward value associated with the current situation for the

WMtk to compute the TD error so as to drive learning in the adaptive critic, and finally one to release the memory occupied by the chunks when no longer needed.

Human-Robot Interaction Interface

As shown in Fig 5.4, a GUI interface was developed for human-robot interaction using GTKmm under Fedora Core 4 and 5 (Fedora Core is an RPM-based Linux distribution, developed by the community-supported Fedora Project and sponsored by Red Hat). GTKmm is a computer graphics library and allows us to create the interface which consists of a number of widgets such as those displaying windows, buttons, menus, etc. Derived from the GTK window class, the interface has four windows, one for the video stream which is received from a video camera, one for selecting a single frame from the stream, one for showing the processed image from that single frame, and one text based window to give human user the status of system execution. We also have a variety of buttons on the interface to control the system, such as loading a database, starting the video, capturing a specific image, processing the captured image, and moving the robot.

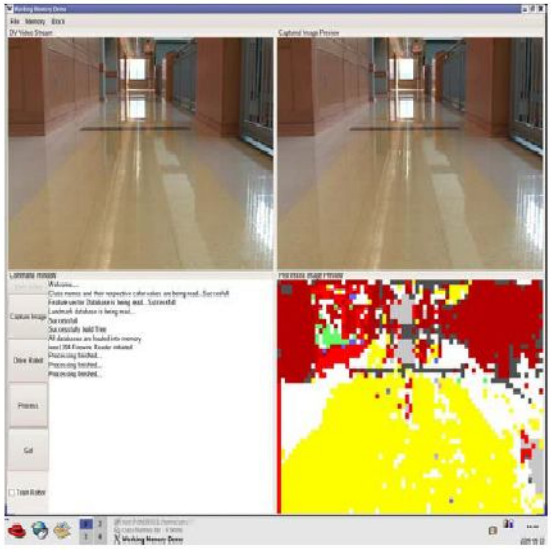


Fig 5.4 Vision-based human-robot interaction interface

Mobile Robot Platform

The mobile robot used in this work is a Pioneer 2 AT four-wheel-drive mobile robot [108], called Skeeter, in the Intelligent Robotics Lab, Vanderbilt University. Pioneer is a family of small and intelligent mobile robots, both two-wheel and four-wheel-drive, whose architecture was originally developed by Dr. Kurt Konolige of SRI International, Inc. and Standford University. These mobile robots contain all of the basic components for sensing and navigation in a realworld environment, including battery power, drive motors and wheels, position/speed encoders, and integrated sensors and accessories.To support accessories available for Pioneer as well as the user's own custom attachments, the controller has two RS232-standard communication ports and an expansion bus. In this project, a laptop and a workstation computer that both run the Fodera core 4 operating systems are used.

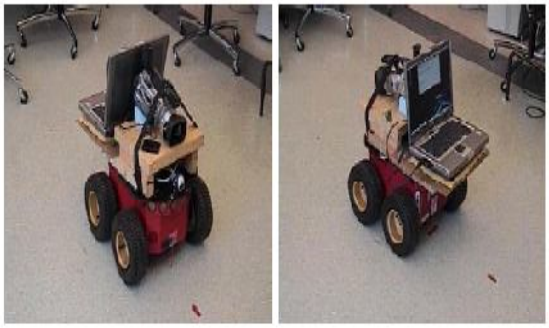


Fig 5.5 Mobile robot Skeeter with camera and laptop

VIII PROPOSED EXPERIMENTS

Experiment 1: Learn Percepts Autonomously

Hypothesis: simultaneous clustering of visual features in feature space and image space can be used for autonomously discovering and learning percepts. To deal with uncertain, ambiguous, contradictory and noisy data, the first proposed experiment in this research is to evaluate the performance of clustering with as little a priori knowledge input from a human teacher as possible. To find reasonable percepts, the fundamental idea is to cluster the obtained feature vectors in an unsupervised manner.

1. First the robot will use a video camera to take one picture at one location and extract feature vectors, then move 1 meters ahead and repeat the process until it moves a total of 20 meters, resulting in 66,740 feature vectors.
2. Upon the obtained 66,740, the robot will perform an unsupervised clustering algorithm to form clusters in the feature space. The number of clusters can be chosen reasonably larger than the "true" number of clusters in the images so that later a possible "merge" operation could be applied. When this grouping is done, the robot will assign a label to each cluster to form the longterm database.
3. Next, test images are grabbed at the same locations where the training images are obtained and new feature vectors are extracted. To the test images, do a simple nearest neighbor search with respect to the database to segment the test images, and then do the approximate nearest neighbor search with respect to the database tree implementation.
4. The segmentation of test images based on the clusters obtained should yield a set of connected regions in the image. A connected-component labeling algorithm is then applied to find these regions, which are called blobs in our research. These blobs can be used to produce working memory chunks.

5. Find all the image blobs (candidate working memory chunks) in the test images. Each blob not only represents one of the clusters in the feature space, but also represents a clustering of pixels in the image space (image space clustering).

Experiment 2: Learn the Meaning of Each Percept with Regard to Movement

Hypothesis: the robot can learn to associate each percept with regard to motion, the so-called sensory motor coordination. With percepts being identified in the previous experiment, the purpose of this experiment is to show the acquired models of objects could identify target objects with good reliability and to learn some properties of objects through the robot's own interaction with the world using the working memory toolkit. At the end of the first experiment, several meaningful blobs are obtained. Since the robot doesn't know anything about each blob, the task in this experiment is to find whether it represents an obstacle or it is open space. 1. Initially set all weights associated with each percept to be 0.5 and 2.5, respectively. 2. In the same place where training and test images are grabbed, ask the robot to take pictures, partition the pictures into blobs, and remember each blob together with their location. Each blob will be used as an input chunk to the working memory toolkit. 3. Ask the robot to move toward each object for 5 meters long. 4. If the robot moves 5 meters long without bumping onto that object, it will receive the maximum reward, 5, and gather evidence that the object is associated with an open space. 5. If the robot moves $x < 5$ meters long and then bumps onto that object, it will receive a less reward, x, and gather evidence that the object is associated with an obstacle.

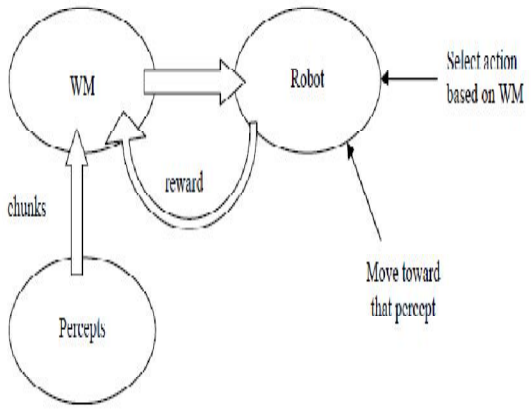


Fig 6.1 Percept learning

Experiment 3: Learn Perceptual Landmarks of a Location

Hypothesis: the robot can learn which percept or percepts are most useful for identifying a specific location even in the presence of distracting percepts. This experiment is to demonstrate that a robot can learn one or more percepts to locate objects and, based on that, build internal representations of objects in its environment without a priori knowledge of these representations from a human teacher. This basically is a classification problem that learns the correlation between the percept and the target location. For geometrical reasons, the landmarks are supposed to be ones that are close to the goal and large in size so as to best aid the pinpointing of a goal. Therefore, pictures are taken at the current and the goal locations.

CHAPTER VII
EXPERIMENT RESULTS AND CONCLUSION

Experiment 1 Results

To obtain a percept database, 20 pictures along the hallway were taken, using a SONY Digital Handycam DCR VX2000 digital video camera recorder, and 66,740 unlabelled feature vectors were extracted from these images. Based on the nearest neighbor distance, a minimum spanning tree is constructed from this data. Next the edges in the tree are sorted and cut in decreasing order until the number of feature vectors in the largest cluster is below a user-specified threshold obtained through error-and-trial. Then any cluster which has more than 100 feature vectors is believed to represent a percept. Finally, a nearest neighbor labeling algorithm is used to assign labels to the remaining points to build up the database. Each class was assigned a different color for displaying purposes. Eleven clusters were found using this method and corresponding percepts are denoted in the following table 7.1. It can be seen that each percept identified by a human has multiple percepts perceived by our robot vision system.

Table 7.1 Autonomously learned percepts

Denoting Symbol	Percept Name Identified by Human
Object1	BlackStripe
Object2	WoodPanel
Object3	LightReflection
Object4	WoodPanel
Object5	WhiteFloor
Object6	WoodPanel
Object7	WoodPanel
Object8	WoodPanel
Object9	WhiteFloor
Object10	YellowFloor
Object11	YellowFloor

CONCLUSION & FUTURE WORK

In this research, the basic issue of a vision-based autonomous mobile robot system has been addressed by integrating the working memory system into a vision-based mobile robot perceptual system. Particularly, an autonomous perceptual system has been addressed from the feature vector generation, similarity measure selection for natural association, to unsupervised clustering for database construction and

fast search implementation using tree structure. It seems that the mixed similarity metric performs better than the Euclidean metric and that fast segmentation with respect to the database search tree works reasonably well. The working memory toolkit, as a biologically inspired model of the PFC and midbrain dopamine system, has been incorporated into the robot control system and validated in the context of a real-world working memory task. Further, an application of the transverse patterning problem to landmark identification has been explored using the WMtk and a single perceptron. The specific scenario that we are studying assumes that there is a target location that the robot is able to recognize visually. Based on the experimental results, it can be concluded that conjunctive coding using the WMtk can do a better job than perceptron even with the appearance of a distracting percept. This is particularly useful for tasks that require visually-stable landmarks be identified and tracked so as to build a map of 92 the environment. However, questions exist with regard to how to find a reward mechanism so that the peak weights correspond to combinations of non-empty distinctive percepts. Finally, two non-conventional approaches for novel object detection are proposed. The online approach is fast but may be less reliable, the offline approach is more reliable but computationally expensive. This suggests that the latter method may be used as a so-called reflective process, i.e., run when the robot is not occupied with a task. It has been shown that both approaches can provide an accurate and robust solution to the novel object detection problem in applications. Therefore, the two novel object detection methods are complementary. Experimental results suggest that the proposed novelty detection methods provide a basis for continuous learning, which is important to supporting developmental robotics. Future work includes whether a more accurate cluster selecting criterion should be used,

REFERENCES

[1] Fukuda, T., Michelini, R., Potkonjak, V., Tzafestas, S., Valavanis, K., &Vukobratovic, M. (2001). How far away is "artificial man"?. IEEE Robotics & Automation Magazine, 8(1): 66-73.

[2] Olshausen, B.A. & Field, D.J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. Nature, 381:607-609.

[3] Goldstone, R. (1998). Perceptual learning.Annu. Rev. Psyhchol, 49:585-612.

[4] Buhmann, J.M., Lange, T. &Ramacher, U. (2005). Image segmentation by networks of spiking neurons. Neural Computation, 17:1010-1031.

[5] Hunter, J., Human Motion Segmentation and Object Recognition using Fuzzy Rules, Proceedings of 14th Annual IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN 2005), Nashville, TN, August 13-15, 2005,pp 210- 216, 2005.

[6] Sperry, R.W. (1952). Neurology and the mind-brain problem.American Scientist, 40:291- 312.

[7] Milner, P. 1999. The Autonomous Brain: A Neural Theory of Attention and Learning. Mahwah, N.J.:L. Erlbaum Associates.

[8] Butler, A.B., and Hodos, W. (1996). Comparative Vertebrate Neuroanatomy: Evolution and Adaptation. John Wiley & Sons.

[9] Sutton, R.S. (1998). Reinforcement learning: an introduction. Cambridge, Mass. : MIT Press.

[10] Baldassarre, G. (2002) A modular neural-network model of the basal ganglia's role in learning and selecting motor behaviors. Cognitive Systems Research 1:5-13.

[11] Fuster, J.M. (1997). The Prefrontal Cortex.Lippincott-Raven Publisher, Philadelphia, New York.

[12] Schultz, W.S., Apicella, P., Romo, R. &Scarnati, E. (1995) Context-dependent activity in primate striatum reflecting past and future behavioral events. In Houk, C.J., Davis, L.J., and Beiser, G.D. (Eds.), Models of information processing in the basal ganglia. Cambridge, MA: The MIT Press.

[13] Hunter, W.S. (1913). The delayed reactions in animals and children.BehavMonogr, 2:1- 86.

[14] Baddeley, A. (2003). Working memory: looking back and looking forward. Nature Review Neuroscience, 4(10): 829-839. 119

[15] O'Reilly, R.C., Noelle, D.C., Braver, T.S., & Cohen, J.D. (2002). Prefrontal cortex and dynamic categorization tasks: Representational organization and neuromodulatory control. Cerebral Cortex, 12:246-257.

[16] Phillips J.L. & Noelle, D.C. (2005). A biologically inspired working memory framework for robots.Proc. of the 27th Annual Meeting of the Cognitive Science Society, Stresa, Italy, July.

[17] Wilkes, M., Tugcu, M., Hunter, J., & Noelle, D.C. (2005). Working memory and perception. Proc. of the IEEE Intl. Workshop on Robot and Human Interactive Communication (RO-MAN), Nashville, TN, August, 2005.

[18] Young, M.P., and Scannell, J.W., 1993. Analysis and modeling of the mammalian cerebral cortex.In Experimental and Theoretical Advances in Biological Pattern Formation, ed. H.G. Othmer, RK. Main and J.D. Murray, pp. 369-384. New York: Plenum Press.

★★★