

Fusion Adaptive Resonance Theory Networks Used As Episodic Memory for An Autonomous Robot

Francis Leconte, François Ferland, and François Michaud

IntRoLab, Interdisciplinary Institute for Technological Innovation (3IT)
Université de Sherbrooke, Sherbrooke, Canada

Abstract. Autonomous service robots must be able to learn from their experiences and adapt to situations encountered in dynamic environments. An episodic memory organizes experiences (e.g., location, specific objects, people, internal states) and can be used to foresee what will occur based on previously experienced situations. In this paper, we present an episodic memory system consisting of a cascade of two Adaptive Resonance Theory (ART) networks, one to categorize spatial events and the other to extract temporal episodes from the robot’s experiences. Artificial emotions are used to dynamically modulate learning and recall of ART networks based on how the robot is able to carry its task. Once an episode is recalled, future events can be predicted and used to influence the robot’s intentions. Validation is done using an autonomous robotic platform that has to deliver objects to people within an office area.

Keywords: Episodic memory, Adaptive resonance theory, Artificial emotions, Autonomous robots

1 Introduction

Autonomous service robots cohabiting with humans will have to achieve recurring tasks while adapting to the changing conditions of the world. According to Hawkins [6], predicting upcoming percepts and action consequences is key to intelligence. Collecting information about one’s experiences over time and their relationships within a spatio-temporal context is a role associated to an episodic memory (EM) [16]. External context information such as location, objects, persons and time [13] can be used, as for internal states such as emotions, behaviors and goals [11, 7]. Memory consolidation and recall can be accomplished by encoding and classifying events (e.g., by using a R-Tree [13]) and by using methods (e.g., probabilistic-based [3]) looking at contextual cues and history to determine if a memorized episode is relevant to the current situation. If an episode is recalled before the robot has completed its task, the memory can be used to anticipate upcoming percepts and actions for the task [8].

More bio-inspired approaches, like Adaptive Resonance Theory (ART) networks, have also been used to categorize patterns from contextual and state

data [1, 14, 18]. Wang et al. [18] use the concept of fusion ART, i.e., two ART networks in cascade [2, 18], to create an EM-ART model: one ART is used to encode spatial events and the other to extract temporal episodes from the experienced situations. Key parameters with this approach are the learning rates β and the vigilance parameters ρ . The learning rates set the influence a pattern has on weight changes, i.e., learning, and is associated to memory stability. Vigilance parameters are used as thresholds for the template matching process: high ρ produce a match when specific input patterns are presented, while lower ρ make more generic pattern matching, tolerant to noise and disparities between the learned pattern and the input pattern. In [19], validation of EM-ART was conducted using a first-person shooter game environment, looking for instance at the influences of ρ on how the episodic memory learns, demonstrating interesting performance of the EM-ART model. However, using EM-ART on an autonomous robot requires dealing with limited, noisy, imprecise and asynchronous perception processes, compared to having complete and continuous access to external context information and internal states of a virtual world. In addition, stability in the representation of events and episodes is required to make the EM-ART usable in the decision-making processes of a robot. Our solution to these issues is to dynamically set β and ρ associated to each events and episodes based on how the robot is able to carry its task, instead of keeping constant β and ρ associated to layers. Such evaluation is conducted using a simple model of artificial emotions. This paper present our EM-ART model, validated using IRL-1 robot platform programmed to deliver objects to people in an indoor environment.

2 EM-ART Modulated with Artificial Emotions

EM-ART is made of three layers: the Input Layer is used to represent the external context information and internal states on which to build the episodic memory; the Event Layer is made of nodes associated to events experienced; and the Episode Layer has nodes that represent the sequence of events making episodes as the robot accomplishes a particular task. Weights between the Input Layer and an Event node represent the pattern from the Input Layer associated to an event, while weights between the Event Layer and an Episode node are associated to the temporal order of events in an episode. As the robot accomplishes its intended task, the matching scheme of EM-ART is used to find similar events and episodes encoded in the memory, adapting weights to reflect variations in similar patterns or adding nodes with their associated weights to learn new events and episodes. Weight learning is influenced by the learning rates β , and the matching scheme by the vigilance parameters ρ . Simply by changing ρ , EM-ART can be used to recall specific events and episodes (e.g., the robot brought Paul a book from Peter in room 1002), or more generic situations (e.g., the robot brought someone an object from Peter in a room). In [18, 19], β and ρ are defined for layers, making learning and matching uniform across layers. In our EM-ART model, we exploit the influences of β and ρ by assigning them to each event and episode nodes according to how the robot is able to satisfy its intentions while accomplishing

the task, as monitored by the artificial emotion module. If a match between the current situation and a memorized episode is found, we also demonstrate how our EM-ART model can be used to predict upcoming event nodes simply by lowering their associated ρ and by ordering them using the memorized weights. Figure 1 illustrates our EM-ART model, described as follows.

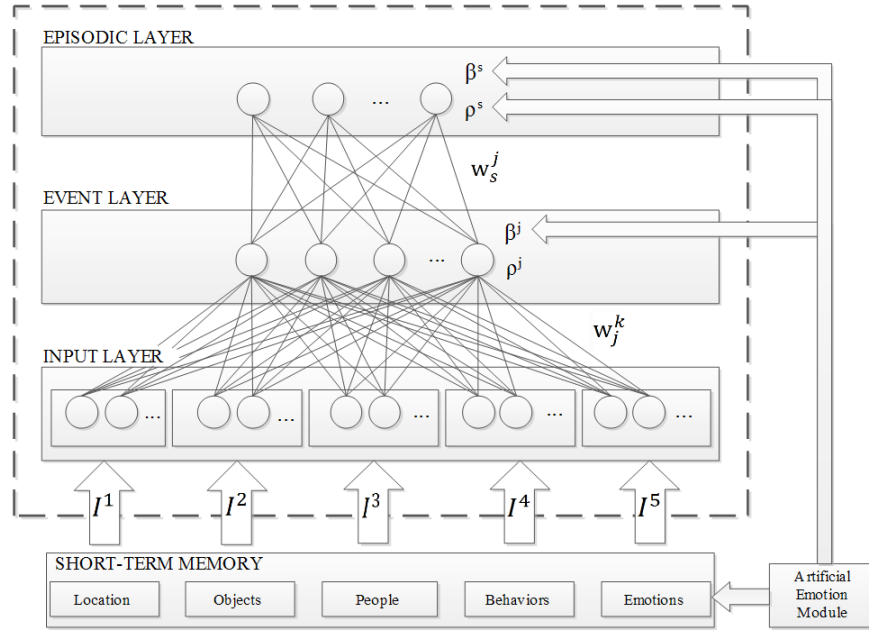


Fig. 1. EM-ART with an Artificial Emotion module

2.1 Input Layer

Let \mathbf{I}^k denote an input vector, with $I_i^k \in [0, 1]$ refers to input attribute i , for $i = 1, \dots, n$. \mathbf{I}^k is augmented with its complements $\bar{\mathbf{I}}^k$ such that $\bar{I}_i^k = 1 - I_i^k$ to define the activity vector \mathbf{x}^k of the Input Layer. Changes in the attributes of \mathbf{x}^k initiate the matching process with the Event Layer.

Input attributes are grouped into cn channels, and with IRL-1 we use five channels: location, objects recognized, people identified, IRL-1's exploited behaviors and its emotional state. A short-term memory buffer is used to synchronize percepts coming from different perceptual modules. For instance, the identity of the person interacting with IRL-1 and the object shown can be observed together even though they are derived using distinct and asynchronous perceptual processes. This allows the Input Layer to aggregate percepts related to more meaningful and significant changes in \mathbf{x}^k , which trigger the matching process in the Event Layer.

2.2 Event Layer

The matching scheme with the Event Layer consists of four steps:

1. **Activating an Event node.** Activation T of node j from the Event Layer is calculated using:

$$T_j = \sum_{k=1}^{cn} \frac{|\mathbf{x}^k \wedge \mathbf{w}_j^k|}{\alpha^k + |\mathbf{w}_j^k|} \quad (1)$$

where \mathbf{w}_j^k is the weight vector associated with the event j and input channel k , $\alpha^k > 0$ is the choice parameter, the fuzzy AND operation \wedge is defined by $(\mathbf{p} \wedge \mathbf{q})_i \equiv \min(p_i, q_i)$, and the norm $|\cdot|$ is defined by $|\mathbf{p}| \equiv \sum_i p_i$ for vectors \mathbf{p} and \mathbf{q} .

2. **Matching of \mathbf{x}^k and hypothesis J .** This step, known as resonance evaluation, examines if, for each channel k , \mathbf{x}^k matches the weights \mathbf{w}_j^k associated to the selected event node (identified as J), according to the following:

$$\frac{|\mathbf{x}^k \wedge \mathbf{w}_j^k|}{|\mathbf{x}^k|} \geq \rho^J \cdot \gamma^k \quad (2)$$

with $\rho_J \in [0, 1]$ being the vigilance parameter associated to the selected event node J and $\gamma^k \in [0, 1]$ being the relevance parameter associated to input channel k . γ^k make EM-ART sensitive to the precision of situational attributes, i.e. the resulting recognition threshold for an event is influenced by characteristics from a bottom layer using γ^k as opposed to vigilance parameter ρ which influences the recognition from a top layer. For instance, $\gamma^k = 0$ for the People channel generates an event regardless of the identity of the individual, while $\gamma^k = 1$ requires that a specific individual be identified to generate an event. Channels with zero relevance allow the system to keep specific information in memory without influencing pattern recognition, while providing useful information when an episode is recalled.

The evaluation starts by selecting node J with the highest T as the hypothesis. If any of the channel k fails to reach resonance with the event J , T_J is set to the next event J having the highest T_j until resonance occurs. If a resonant state is not reached, a new node is created as J .

3. **Learning.** Using J as the Event node, learning is performed according to:

$$\mathbf{w}_j^{k(new)} = (1 - \beta^J) \mathbf{w}_j^{k(old)} + \beta^J (\mathbf{x}^k \wedge \mathbf{w}_j^{k(old)}) \quad (3)$$

where $\beta^J \in [0, 1]$ is the learning rate parameter associated to the event J . $\beta^J = 1$ is used when a new node is created.

4. **Evaluating the activity vector $\mathbf{y} = y_1, \dots, y_m$ of the Event Layer.** For node J , $y_J = 1$, the activities of other nodes on the Event Layer decay linearly according to:

$$y_j^{(new)} = \max(0, y_j^{(old)}(1 - \tau)) \quad (4)$$

where τ is the decaying factor $\in [0, 1]$, which incidentally set the maximum number of event nodes that can be activated to derive an episode.

2.3 Episode Layer

The role of the Episode Layer is to recognize temporal patterns (or sequences of events) in the Event Layer and to predict upcoming events using the concept of temporal auto-association [5]. Whenever \mathbf{y} changes, the Episode Layer uses a matching scheme identical to the Event Layer, evaluating resonance with node S from the Episode Layer or creating a new node if no matches are found. Learning is done only when the task assigned to the robot is completed. Recognition of temporal patterns throughout a task must happen early enough to benefit from the prediction of these patterns before the end of the task, and therefore ρ^s are generally low by default, so that episode node can reach a resonant state more easily.

By default, every ρ^j are set high (0.95) to recognize specific contextual events, but are lowered to conduct a prediction if an episode is recognized before it is completed (to then be reset at the value before prediction occurred). If resonance occurs for episode node S at event node J , \mathbf{w}_S^j between the Event Layer and the episode S can be used to derive the relative order of events in the episode. The prediction \mathbf{y}_P of upcoming events can be calculated using the complement of \mathbf{w}_S^j and \mathbf{y} :

$$\mathbf{y}_P = \mathbf{w}_S^j \setminus \mathbf{y}, \quad w_S^j > w_S^J \text{ and } y_j > 0 \quad (5)$$

Anticipated events are subsequently reordered chronologically (in ascending order according to \mathbf{w}_S^j). To facilitate matching of these upcoming events, minor differences are tolerated by lowering ρ^j according to:

$$\rho^{j(new)} = \rho^{j(old)} \left(1 - C_\rho \left(1 - \frac{(p-1)}{\text{length}(\mathbf{y}_P)} \right) \right) \quad (6)$$

where p is the relative index of the event in the reordered sequence \mathbf{y}_P , and C_ρ is a constant that defines the maximum decrement for ρ^j . The next upcoming event ($p = 1$) has its vigilance parameter decreased the most. Lowering matching threshold of predicted patterns is a concept that is believed to be existing in the human brain [9]. Predicted events are more likely to appear in the current episode, so lowering ρ^j facilitates their activation and makes it possible to tolerate minor differences. To retrieve specific situational attributes related to \mathbf{y}_P , weights w_j^k can be read out one at a time following the sequential order of the anticipated events.

2.4 Artificial Emotion Module

The Artificial Emotion Module is used to adjust ρ^s and β^s to favor recall of the most relevant episode and to improve episode stability in memory. Two artificial emotions intensities $E_e \in [0, 1]$ are used: Joy (indicating that the robot behaves according to its intentions) and Anger (indicating that its intentions are not satisfied). The heuristic used is that when an episode is experienced with high emotion intensity, such episode needs to be stable in memory, meaning that it

should remain intact as future learning occurs. This is done by lowering the learning rates (β^s or β^j), which will limit weight changes:

$$\beta^{(new)} = \beta^{(old)} (1 - C_\beta \cdot (\max(E_e) - 0.5)), \quad \beta^{(new)} \in [\beta_{min}, \beta_{max}] \quad (7)$$

where C_β is a constant that defines the maximum decrement, and β_{min} and β_{max} limiting the range. A max E_e lower than 0.5 increases β , while a value above 0.5 decreases it.

Also, episodes with high emotional intensities must be recalled easily, meaning that ρ^s can be decreased according to:

$$\rho^{s(new)} = \rho^{s(old)} (1 - C_\rho \cdot (\max(E_e) - 0.5)), \quad \rho^{s(new)} \in [\rho_{min}^s, \rho_{max}^s] \quad (8)$$

where ρ_{min}^s and ρ_{max}^s limits the range.

Equations (7) and (8) are applied and saved when learning occur on the associated layer.

3 Experimental Setup

IRL-1 is a robotic platform composed of a humanoid torso on top of a mobile base [4]. IRL-1 uses a Kinect motion sensor for vision processing, a laser range finder for obstacle avoidance and simultaneous localization and mapping (implemented using [15]), and a 8-microphone array for speech interaction with people. IRL-1 detects a person by merging information from legs detection, voice direction and face detection, turning its head toward the person. People identification is implemented using a basic face recognition algorithm based on Principle Component Analysis on the detected face [17]. Objects recognition is done using 2D images from the Kinect using SIFT [10]. Two computers running Linux and Robot Operating System (ROS) [12] are used to implement IRL-1 control architecture.

For this experiment, IRL-1’s task is to deliver one of three objects O_1 , O_2 and O_3 to people in a different location, according to the following scenario:

- In room R_0 , a person P stops in front of IRL-1. IRL-1 then identifies and greets the person.
- Person P shows the object O_o to IRL-1. IRL-1 then recognizes the object and extends its left arm to grasp it.
- IRL-1 autonomously navigates to the other room R_1 , searching to deliver object O_o to somebody. When entering a room, IRL-1 asks if there is someone there to take object O_o . IRL-1 wanders around until a person D is located in area L inside room R_1 .
- IRL-1 extends its arm and delivers object O_o .

Each occurrence of this scenario consists of an episode. Once the task completed, learning is triggered in the Episode Layer and IRL-1 is programmed to return to room R_0 to start again.

For the trials, parameters of our EM-ART are initialized as follows: $\rho^j = 0.95$, $\rho^s = 0.55$, $\rho_{min}^s = 0.45$, $\rho_{max}^s = 0.85$, $C_\rho = 0.20$, $\beta^j = \beta^s = 0.6$, $\beta_{min} = 0.1$, $\beta_{max} = 1$, $C_\beta = 0.25$, $\tau = 0.05$ and $\alpha^k = 0.01$. Joy and Anger are associated with the following control modules controlling IRL-1: *Teleoperation* (required when IRL-1 loses its position in the map), *Go To* (to navigate from one room to another) and *Wandering*. If these modules are activated (meaning that IRL-1 wants to satisfy the intended goal associated to these controllers) and exploited (meaning that IRL-1 is using these modules to control its actions) over time, then the intensity of Joy increases; otherwise, if they are activated but not exploited, Anger increases. For example, when IRL-1 activates *Go To*, Joy increases and Anger decreases as long as the module is exploited. If IRL-1 gets lost in its internal map, the *Go To* behavior is no longer exploited and therefore Joy decreases and Anger increases.

4 Experimental Results

To demonstrate the use of our EM-ART model, we conducted 10 trials for each of the following conditions, each trial initiated with an empty memory, to observe how it responds to different types of situations.

1. **Recall repeatability and prediction.** R_0 , R_1 , P , O_o , D and L remained identical throughout the trials, leading to only one episode. The EM-ART should therefore be able to recall the episode as soon as possible, allowing IRL-1 to predict where to go before having to wander in room R_1 , allowing IRL-1 to use L as a destination to go to. Successful recall of L occurred 8 times out of 9. Trial 1 led to the creation of an episode made of 15 event nodes. For trials 2 to 4, recall occurred relatively late in the episodes, i.e., while IRL-1 was wandering in R_1 . As the scenario was repeated and learned, recurring events stayed while sporadic events faded, and recall occurred as early as when IRL-1 was in R_0 , after having recognized the object O_o or the person P . In the one trial where recall was not observed, IRL-1 lost its position in its map: teleoperation was required, Anger were generated and 10 new event nodes were created, leading to a distinct episode.
2. **Recall repeatability and learning.** R_0 , R_1 , P , O_o and D remained identical throughout the trials, while L changed with each trial. The objective of this condition was to observe if the last L learned could be predicted as the destination when an episodic recall occurred. Successful recall of L happened 8 times out of 9. Each time, the destination predicted was the one from the previous trial, as expected, and IRL-1 started wandering from that point. When learning of the episode with the new destination L occurred, the weighted connections to the previous destination L were reduced. For the unsuccessful trial, a false detection of the object recognition module led to the creation of five new event nodes at the beginning of the episode, and consequently, created of a new episode.
3. **Semantic differences and creation of new episodes.** R_0 , R_1 , P and D remained identical throughout the trials, while O_o changed to be one of

three objects (O_1, O_2, O_3). This should lead to the creation of three episodes semantically different but with some similar events. Each object delivery was done at a specific location L for the object in question, to differentiate which episode was recalled and used to predict L . Figure 2 presents the total number of event nodes and episode nodes in memory after each trial, and with the object presented. As expected, each trial involving a new object O_o led to the creation of a new episode, for a total of three. The number of event nodes in the memory increased in the first five trials since percepts changed slightly over the episodes, but stabilized in the last five trials.

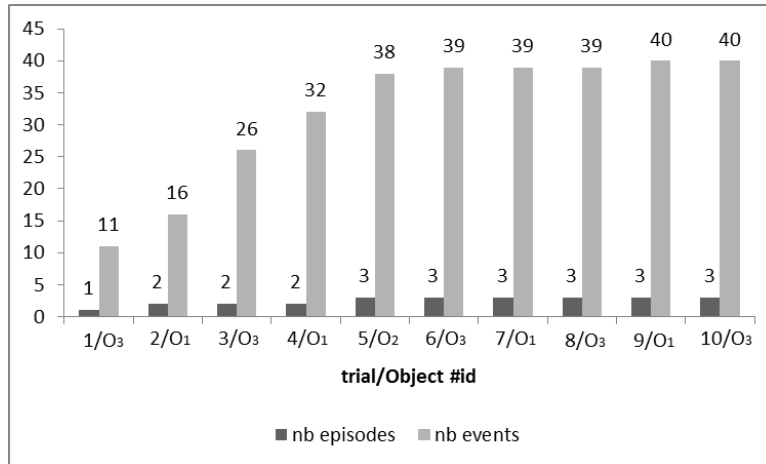


Fig. 2. Number of nodes of the Event Layer and Episode Layer as IRL-1 is being presented with three objects.

4. **Relevance of the input channel.** To test the influence of γ^k on an event node, R_0, R_1, P, D and L are kept identical over the trials, $\gamma^2 = 0$ for the Object channel, and a different object (O_1, O_2, O_3) was used between trials. This condition should lead to the same episode, making the objects carried by IRL-1 irrelevant for the episode, and as expected, after 10 trials, only one episode was learned. IRL-1 also recalled the episode when entering R_1 (once in the corridor between R_0 and R_1), and went directly to the delivery location L without wandering in the room R_1 , regardless of O_o .
5. **Episode with high emotional intensity.** R_0, R_1, P, O_o, D and L remained identical throughout the trials, but we forced IRL-1 to experience high emotional intensity during trial 1, by deliberately covering up the laser range sensor, making the *Go To* module unusable. This made Anger reached its maximum value as the episode was learned, leading to the decrease of β^s and ρ^s according to Eq. (7) and Eq. (8). This condition should lead to rapid episode recall, allowing IRL-1 to benefit from a prediction early on at the beginning of the task. Indeed, during trials 2 to 10, the episode learned in

trial 1 was recalled as soon as IRL-1 realized it was in R_0 : IRL-1 then decided to directly go to the delivery location L .

6. **Episode with no emotions.** R_0, R_1, P, O_o, D and L remained identical throughout the trials, but we set $E_e = 0$ for Joy and for Anger to illustrate the influences of emotions on recall. According to Eq. (7) and Eq. (8), ρ^s will increase over time, and the episode will not be recognized as easily. During trials 2 and 3, a successful episode recall was observed, allowing IRL-1 to predict the delivery location L . During trials 4 and 5, IRL-1 recognized the episode, but the prediction was not useful since L was already reached after having wandered for a while. Starting trial 6, episode recall did not happen before the end of the task because ρ_s was too high (0.85) to tolerate minor variations in the sequence of events, leading to the creation of new episodes in memory. After ten trials, the episodic memory contained three episodes rather than only one.

5 Conclusion and Future Work

The underlying objective of providing a robot with an episodic memory is to allow it to adapt its decision-making processes according to past experiences when operating in dynamic environments. This paper presents a variant of EM-ART in which the learning rate parameters and the vigilance parameters are associated to specific event and episode. Changing the learning rate influences weight adaptation to either learn quickly ($\beta = 1$) or preserve what was experienced in the past ($\beta = 0$), whether it is for an event node or an episode node. The vigilance parameters set what can be characterize as the granularity of the matching scheme: it can be identical ($\rho = 1$) or coarse ($\rho = 0.1$), in relation to input channels or to events. Keeping these parameters constant across layers consider that each episode has the same importance, which is unrealistic considering that the episode experienced may or may not be the results of appropriate actions according to the robot's intentions. Using a repeatable scenario involving people recognition, object recognition and location identification, we illustrate how adapting these parameters can lead to appropriate episode learning and recall, and how upcoming predicted events can be used to influence the behaviour of the robot. Results show that the robot successfully differentiates semantically dissimilar episodes and expands its memory to learn new situations online. To explore further all the potential of our EM-ART model, future work involves extensive testing with a higher number of trials experienced randomly, with different complex tasks in which repeatable sequences can be experienced, and observing how the EM evolves over time.

ACKNOWLEDGMENT

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

1. Carpenter, G.A., Grossberg, S.: A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer vision, graphics, and image processing* 37(1), 54–115 (1987)
2. Carpenter, G.A., Grossberg, S., Rosen, D.B.: Fuzzy ART: Fast stable learning and categorization of analog patterns by an adaptive resonance system. *Neural Networks* 4(6), 759–71 (1991)
3. Dodd, W., Gutierrez, R.: The role of episodic memory and emotion in a cognitive robot. In: *Proceedings of IEEE International Workshop on Robot and Human Interactive Communication*. pp. 692–7 (2005)
4. Ferland, F., Létourneau, D., Aumont, A., Frémy, J., Legault, M.A., Lauria, M., Michaud, F.: Natural interaction design of a humanoid robot. *Journal of Human-Robot Interaction* 1(2), 118–34 (2012)
5. Haikonen, P.O.: *Consciousness and robot sentience*, vol. 2. World Scientific Publishing Company (2012)
6. Hawkins, J.: *On intelligence*. Macmillan (2004)
7. Komatsu, T., Takeno, J.: A conscious robot that expects emotions. In: *Proceedings of IEEE International Conference on Industrial Technology*. pp. 15–20 (2011)
8. Kuppuswamy, N.S., Cho, S.H., Kim, J.H.: A cognitive control architecture for an artificial creature using episodic memory. In: *Proceedings of SICE-ICASE International Joint Conference*. pp. 3104–10 (2006)
9. Kurzweil, R.: *How to Create a Mind: The Secret of Human Thought Revealed* (2012)
10. Lowe, D.G.: Object recognition from local scale-invariant features. In: *Proceedings of IEEE International Conference on Computer Vision*. vol. 2, pp. 1150–7 (1999)
11. Nuxoll, A.M., Laird, J.E.: Enhancing intelligent agents with episodic memory. *Cognitive Systems Research* pp. 34–48 (2011)
12. Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., Wheeler, R., Ng, A.Y.: ROS: An open-source robot operating system. In: *ICRA Workshop on Open Source Software* (2009)
13. Stachowicz, D., Kruijff, G.: Episodic-like memory for cognitive robots. *IEEE Transactions on Autonomous Mental Development* 4(1), 1–16 (2011)
14. Taylor, S.E., Vineyard, C.M., Healy, M.J., Caudell, T.P., Cohen, N.J., Watson, P., Verzi, S.J., Morrow, J.D., Bernard, M.L., Eichenbaum, H.: Memory in silico: Building a neuromimetic episodic cognitive model. In: *Proceedings of World Congress on Computer Science and Information Engineering*. vol. 5, pp. 733–7 (2009)
15. Thrun, S., Fox, D., Burgard, W., Dellaert, F.: Robust Monte Carlo localization for mobile robots. *Artificial Intelligence* 128(1), 99–141 (2001)
16. Tulving, E.: Precis of elements of episodic memory. *Behavioral and Brain Sciences* 7(2), 223–68 (1984)
17. Viola, P., Jones, M.J.: Robust real-time face detection. *International Journal of Computer Vision* 57(2), 137–54 (2004)
18. Wang, W., Subagdja, B., Tan, A.H., Starzyk, J.A.: A self-organizing approach to episodic memory modeling. In: *Proceedings of International Joint Conference on Neural Networks*. pp. 1–8 (2010)
19. Wang, W., Subagdja, B., Tan, A.H., Starzyk, J.A.: Neural modeling of episodic memory: Encoding, retrieval, and forgetting. *IEEE Transactions on Neural Networks and Learning Systems* 23(10), 1574–86 (2012)