Agent-Augmented Co-Space: Toward Merging of Real World and Cyberspace

Ah-Hwee Tan and Yilin Kang

School of Computer Engineering
Nanyang Technological University, Nanyang Avenue, Singapore 639798
asahtan@ntu.edu.sg, kang0028@ntu.edu.sg

Abstract. Co-Space refers to interactive virtual environment modelled after the real world we are situated in. Through realistic 3D modelling and animation technologies, Co-Space simulates the real world in terms of look-and-feel of our physical surrounding. With the advancement in pervasive sensor network, Co-Space may also capture and mirror the happening in the physical world in real time. The development of Co-Space thus offers great opportunities for delivering innovative applications and services. Specifically, for enriching the experience of users in Co-Space, it is essential to incorporate knowledge facilities in the form of intelligent agents to enhance the interactivity and playability within. This paper will begin with a brief review of this emerging field of work related to agents in virtual worlds and integrated cognitive architectures. We then discuss the key requirement, issues and challenges in making Co-Space interactive and intelligent. Following the notion of embodied intelligence, we propose to develop cognitive agents, based on a family of self-organizing neural models, known as fusion Adaptive Resonance Theory (fusion ART). Our ultimate aim is to have such agents roaming freely in the landscape of Co-Space, developing an awareness of its surrounding and interacting with avatars of real human. As an illustration, a case study of our effort in building the Singapore Youth Olympic Village (YOV) Co-Space will be presented.

Keywords: Co-Space, Intelligent agents, virtual world.

1 Introduction

Virtual world has become a popular platform used in a variety of contexts, including education, business, and e-commerce [41]. Studies in South Korea have recently shown that users prefer virtual world to television [42]. Gartner even predicted that 80 percent of the Internet users will be actively participating in non-gaming virtual world by the end of 2011. To date, many popular virtual worlds exist, such as Second Life and Active Worlds, enabling users to create artificial content in virtual environment.

In our work, we are particularly interested in a special class of virtual world, called Co-Space, referring to interactive virtual environment modelled after a real physical world in terms of look-and-feel, functionalities and services. Through

B. Xie et al. (Eds.): ATC 2010, LNCS 6407, pp. 298–312, 2010.

realistic 3D modelling and animation technologies, Co-Space simulates the real world in terms of look-and-feel of our physical surrounding. With the advancement in pervasive sensor network, Co-Space may also capture and mirror the happening in the physical world in real time. Besides providing a much faster and easier access to information and services, the development of Co-Space has offered great opportunities for delivering innovative applications and services. Specifically, intelligent agents can be deployed in Co-Space enhancing its interactivity and playability.

Despite the appealing potential, deploying intelligent agents in a virtual world, just like in a real world, poses many challenges not addressed by traditional AI and machine learning algorithms. In particular, learning in virtual world is typically unsupervised, without an explicit teacher to guide the agent in learning. Furthermore, it requires an interplay of a myriad of learning paradigms. Due to such difficulties, most virtual worlds tend to constrain agents' actions to a very coarse level, dictated by hard coded rules [27,43,28].

In contrast to existing approaches, we take the view that a large part of an agent's intelligence is acquired through its interaction with the environment. This is in keeping with the view in modern cognitive science that cognition is a process deeply rooted in the body's interaction with the world [2]. Furthermore, we hypothesize that a cognitive autonomous system, with the appropriate architecture and the necessary adaptation mechanisms, is a sufficient self to learn and interact in a dynamic environment. Embodied cognition is also akin to the intensive study on reinforcement learning [31] in which an autonomous agent learns to adjust its behaviour according to evaluative feedback received from the environment.

Following the notion of embodied intelligence, we develop cognitive architectures, based on a family of self-organizing neural models, known as fusion Adaptive Resonance Theory (fusion ART) [35]. Fusion ART is a generalization of self-organizing neural models known as Adaptive Resonance Theory (ART) [5,8]. By extending the original ART model consisting of a single pattern field into a multi-channel architecture, fusion ART unifies a number of important neural models, developed over the past decades, including the original ART models, Adaptive Resonance Associative Map (ARAM) designed for supervised learning [32,39], and Fusion Architecture for Learning and Cognition (FALCON) [33,35], designed for reinforcement learning.

With the properties of self-adaptation, generalization, and fast yet stable real-time learning, fusion ART makes a suitable building block for designing learning agents in virtual world. By incorporating fusion ART, an agent will be able to learn from sensory and evaluative feedback signals received from the virtual environment largely without involving human supervision and intervention. In this way, the agent needs neither an explicit teacher nor a perfect model to learn from. Performing reinforcement learning in real time, it is able to adapt itself to the variations in the virtual environment and changes in the user behavior patterns.

In the next section, we shall review the related work on intelligent agents in virtual worlds and integrated cognitive architectures. We then discuss several issues and challenges in designing agents for an agent-augmented virtual environment.

We then present an embodied intelligence approach building upon the learning and adaptation capability of fusion ART and its extended architectures that integrate learning with higher level functions. Finally, we present a case study on a Co-Space known as Youth Olympic Village (YOV) Co-Space. The final section concludes and highlights possible future directions.

2 Related Work

2.1 Learning Agents in Virtual Worlds

Though intelligent agents have been popularly used for improving the interactivity and playability of virtual worlds, most such agents are based on scripts or predefined rules. For example, in the Virtual Theater project, synthetic actors portray fictive characters and provide improvising behaviors. The agents are based on a scripted social-psychological model with the defined personality traits which rely on the values of moods and attitudes [27]. Agents in Metaverse, built using Active Worlds, are capable of performing the tasks typically associated with human beings, such as taking tickets for rides and acting as shopkeepers. However, these agents are basically reactive agents which work in a hard-coded manner. Virtual psychotherapist ELIZA [43], designed to take care of the 'patients', is also achieved with rule-based, adeptly modeled small talk. A conversational virtual agent Max has been developed as a guide to the HNF computer museum, where he interacts with visitors and provides them with information daily [28]. However, the design remains rule-based.

In view of the limitations of static agents, some researchers have adopted learning methods into service agents in virtual world. For example, Yoon et.al. present a Creature Kernel framework to build interactive synthetic characters in the project Sydney K9.0 [44]. Their agents can reflect the characters' past experience and allow individual personalization. But all the capabilities of the agents rely on past knowledge and couldn't adapt to user gradually during run time. The co-present agents in a virtual gallery [12] utilize a knowledge base containing general input response knowledge, augmented with knowledge modules for special domains. More recently, an embodied conversational agent that serves as a virtual tour guide in Second Life has been implemented by Jan [15]. It uses NPCEditor [21] to learn the best output for any input from a training set of linked questions and answers. Again, it learns from past experience but does not adapt over time according to the habits of a particular player or the changes in the environment.

2.2 Integrated Cognitive Models

On the other hand, research in intelligent autonomous systems has been a key focus in the fields of cognitive science and artificial intelligence in the past decades [9]. Below we review six cognitive architectures, namely Soar, ACT-R, ICARUS, BDI, the subsumption architecture, and CLARION, roughly classified according to their roots and emphases.

Soar [18,20], based on the physical symbolic hypothesis [23], is one of the earliest and most extensively developed AI architectures in the history. ACT-R [1] (also with a long history) and ICARUS [19] (a relatively recent model) are cognitive systems developed with the primary aim of producing artificial intelligence mimicking human cognition. While the three architectures share many features of classical artificial intelligence, including symbolic representation, production rule based inference, and means-end analysis for problem solving, ACT-R and ICARUS are notably different from Soar by their strong emphasis of producing a psychologically motivated cognitive model.

Belief-Desire-Intention (BDI) architecture [3,26] is a popularly used framework, incorporating beliefs, desires and intentions, for designing intelligent autonomous agents. Based on the studies of folk psychology and intentional systems, BDI has a special focus on intentions, representing an agent's commitments to carry out certain plans of actions [11]. Coined as the new artificial intelligence, the subsumption architecture [4] is notably different from the other cognitive architectures in its approach and design. The subsumption architecture is behaviour based and thus does not contain any problem solving or learning module. The idea of higher layers subsuming lower layers in the subsumption architecture has its root from neurobiology. CLARION [30,29] is a hybrid model integrating both symbolic and connectionist information processing. The design of CLARION is based on neural networks as well as cognitive psychology. As a result, it is similar to ACT-R as both models are based on a combination of artificial intelligence, cognitive psychology and some favour of neurobiology.

All cognitive architectures contain certain features which make them unique. However, no single cognitive architecture has provided a solution that is within the level of human cognition. Reinforcement Learning is a field that has received intensive research effort, but has not been incorporated into cognitive architectures in a major and principled way. Among these hybrid systems, temporal difference learning using gradient descent based function approximator has been most commonly used. However, the gradient descent methods learn by making small error correction steps iteratively. In addition, there is the issue of instability as learning of new patterns may erode the previously learned knowledge. Consequently, the resultant systems may not be able to learn and operate in real time.

3 Issues and Challenges

3.1 How to Create Autonomy and Self-Awareness?

Autonomy is the ability of an agent to act and make decisions on its own independently of the programmer or user. It is required for an agent-based autonomous entities to be self-awareness so as to increase the dynamics of the virtual environment. For example, in a Co-Space populated with agent-based avatars, the non-player characters should be autonomous in initiating actions and interacting with users proactively. Autonomy is also necessary to enable an agent to explore the virtual environment in the absence of continuous instructions by the user.

3.2 How to Enhance Interactivity?

In the domain of computer games, researchers studied how to improve the game's playability, of which interactivity has been identified as one of the most important factors. For example, a player who is constantly engaged will find the place more interesting comparing to one who has been left alone to explore in a lifeless land. As such, when agents are augmented as enemies or fellows in a virtual world, those agents should have frequent and natural interaction with the players to enhance playability.

3.3 How to Enable Situatedness?

Situatedness is also a consideration for intelligent agents in virtual world as they tend to be used in a dynamic, unpredictable and unreliable environment. As the environment changes rapidly, the agent cannot assume that the situation will remain stationary while it figures how to achieve a goal. The environment can also be unpredictable sometime due to the limitations of the agent in obtaining accurate and complete information about the environment or that the environment is being modified in ways beyond the agent's knowledge and reasoning capability. Finally, the environment can be unreliable in that the actions that an agent can perform may fail for reasons beyond an agent's control.

3.4 How to Learn and Function in Real-Time?

The ability to learn is another important issue. A good agent needs to acquire new knowledge and skills and improve its performance in carrying out a particular task over time. For example, if a user signals to an agent that it performs poorly on a task, the agent should be able to learn from this experience and avoid making the same mistake in the future. The capability of learning can also be an issue of playability, as an agent should not struggle to learn the rules of how to behave in an appropriate manner. Furthermore, this issue of learning becomes much more complex and challenging in a real-time dynamic multi-agent environment.

3.5 How to Learn about Users for Personalization?

Especially in the domain of education and gaming, one important aspect of agents that has received a great deal of attention is user modelling and personalization. For example, teaching in many virtual worlds is achieved with pedagogical agents as virtual teachers to monitor and provide learners with personalized guidance. To achieve this goal, personalization becomes an important issue in agent technology, i.e., adapting the teaching to the needs of various learners.

4 An Embodied Intelligence Approach

Over the past decades, a family of neural models known as Adaptive Resonance Theory (ART) [7,8] has been steadily developed. With well-founded computational principles, ART has been applied successfully to many pattern analysis, recognition, and prediction applications [10,22]. These successful applications are of particular interest because the basic ART principles have been derived from an analysis of human and animal perceptual and cognitive information processing, and have led to behavioral and neurobiological predictions that have received significant experimental support during the last decade [13,25].

In this paper, we show that Adaptive Resonance Theory lays the foundation of a unified model that encompasses a myriad of learning paradigms, traditionally viewed as distinct. The proposed model is a natural extension of the original ART models from a single pattern field to multiple pattern channels. Whereas the original ART models perform unsupervised learning of recognition nodes in response to incoming input patterns, the proposed neural architecture, known as fusion ART, learns multi-channel mappings simultaneously across multi-modal pattern channels in an online and incremental manner.

4.1 Fusion ART

Fusion ART employs a multi-channel architecture (Figure 1), comprising a category field F_2 connected to a fixed number of (K) pattern channels or input fields through bidirectional conditionable pathways. The model unifies a number of network designs, most notably Adaptive Resonance Theory (ART) [7,8], Adaptive Resonance Associative Map (ARAM) [32] and Fusion Architecture for Learning, COgnition, and Navigation (FALCON) [33], developed over the past decades for a wide range of functions and applications. The generic network dynamics of fusion ART, based on fuzzy ART operations [6], is summarized as follows.

Input vectors: Let $\mathbf{I}^{ck} = (I_1^{ck}, I_2^{ck}, \ldots, I_n^{ck})$ denote the input vector, where $I_i^{ck} \in [0,1]$ indicates the input i to channel ck. With complement coding, the input vector $\mathbf{I}^{ck}$ is augmented with a complement vector $\bar{\mathbf{I}}^{ck}$ such that $\bar{I}_i^{ck} = 1 - I_i^{ck}$.

Activity vectors: Let $\mathbf{x}^{ck}$ denote the F_1^{ck} activity vector for $k = 1, \ldots, K$. Let $\mathbf{y}$ denote the F_2 activity vector.

Weight vectors: Let $\mathbf{w}_j^{ck}$ denote the weight vector associated with the jth node in F_2 for learning the input patterns in F_1^{ck} for $k = 1, \ldots, K$. Initially, F_2 contains only one *uncommitted* node and its weight vectors contain all 1's.

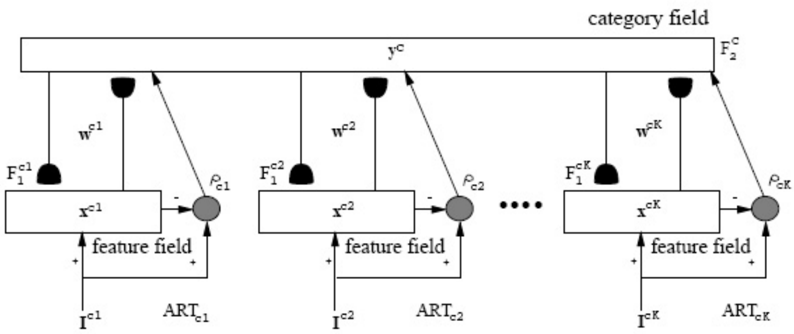

Fig. 1. The fusion ART architecture

Parameters: The fusion ART's dynamics is determined by choice parameters $\alpha^{ck} > 0$, learning rate parameters $\beta^{ck} \in [0, 1]$, contribution parameters $\gamma^{ck} \in [0, 1]$ and vigilance parameters $\rho^{ck} \in [0, 1]$ for $k = 1, \ldots, K$.

As a natural extension of ART, fusion ART responds to incoming patterns in a continuous manner. It is important to note that at any point in time, fusion ART does not require input to be present in all the pattern channels. For those channels not receiving input, the input vectors are initialized to all 1s. The fusion ART pattern processing cycle comprises five key stages, namely code activation, code competition, activity readout, template matching, and template learning, as described below.

Code activation: Given the activity vectors $\mathbf{I}^{c1}, \ldots, \mathbf{I}^{cK}$, for each F_2 node j, the choice function T_j is computed as follows:

$$T_j = \sum_{k=1}^{K} \gamma^{ck} \frac{|\mathbf{I}^{ck} \wedge \mathbf{w}_j^{ck}|}{\alpha^{ck} + |\mathbf{w}_j^{ck}|}, \tag{1}$$

where the fuzzy AND operation $\wedge$ is defined by $(\mathbf{p} \wedge \mathbf{q})_i \equiv min(p_i, q_i)$, and the norm $|.|$ is defined by $|\mathbf{p}| \equiv \sum_i p_i$ for vectors $\mathbf{p}$ and $\mathbf{q}$.

Code competition: A code competition process follows under which the F_2 node with the highest choice function value is identified. The winner is indexed at J where

$$T_J = \max\{T_j : \text{for all } F_2 \text{ node } j\}. \tag{2}$$

When a category choice is made at node J, $y_J = 1$; and $y_j = 0$ for all $j \neq J$. This indicates a winner-take-all strategy.

Activity readout: The chosen F_2 node J performs a readout of its weight vectors to the input fields F_1^{ck} such that

$$\mathbf{x}^{ck} = \mathbf{I}^{ck} \wedge \mathbf{w}_J^{ck}. \tag{3}$$

Template matching: Before the activity readout is stabilized and node J can be used for learning, a template matching process checks that the weight templates of node J are sufficiently close to their respective input patterns. Specifically, resonance occurs if for each channel k, the *match function* m_J^{ck} of the chosen node J meets its vigilance criterion:

$$m_J^{ck} = \frac{|\mathbf{I}^{ck} \wedge \mathbf{w}_J^{ck}|}{|\mathbf{I}^{ck}|} \geq \rho^{ck}. \tag{4}$$

If any of the vigilance constraints is violated, mismatch reset occurs in which the value of the choice function T_J is set to 0 for the duration of the input presentation. Using a *match tracking* process, at the beginning of each input presentation, the vigilance parameter ρ^{ck} in each channel ck equals a baseline vigilance $\bar{\rho}^{ck}$. When a mismatch reset occurs, the ρ^{ck} of all pattern channels are increased simultaneously until one of them is slightly larger than its corresponding match function m_J^{ck}, causing a reset. The search process then selects another F_2 node J under the revised vigilance criterion until a resonance is achieved.

Template learning: Once a resonance occurs, for each channel ck, the weight vector $\mathbf{w}_J^{ck}$ is modified by the following learning rule:

$$\mathbf{w}_J^{ck(\text{new})} = (1 - \beta^{ck})\mathbf{w}_J^{ck(\text{old})} + \beta^{ck}(\mathbf{I}^{ck} \wedge \mathbf{w}_J^{ck(\text{old})}). \tag{5}$$

When an uncommitted node is selected for learning, it becomes *committed* and a new uncommitted node is added to the F_2 field. Fusion ART thus expands its network architecture dynamically in response to the input patterns.

4.2 Learning and Adaptation

The network dynamics described above can be used to support a myriad of learning operations. We show how fusion ART can be used for a variety of traditionally distinct learning tasks in the subsequent sections.

Learning by Similarity Matching. With a single pattern channel, the fusion ART architecture reduces to the original ART model. Using a selected vigilance value ρ, an ART model learns a set of recognition nodes in response to an incoming stream of input patterns in a continuous manner. Each recognition node in the F_2 field learns to encode a template pattern representing the key characteristics of a set of patterns. ART has been widely used in the context of unsupervised learning for discovering pattern groupings. Please refer to the selected ART literatures [7,8] for a review of ART's functionalities, interpretations, and applications.

Learning by Association. By synchronizing pattern coding across multiple pattern channels, fusion ART learns to encode associative mappings across distinct pattern spaces. A specific instance of fusion ART with two pattern channels is known as Adaptive Resonance Associative Map (ARAM), that learns multi-dimensional supervised mappings from one pattern space to another pattern space [32]. An ARAM system consists of an input field F_1^a, an output field F_1^b, and a category field F_2. Given a set of feature vectors presented at F_1^a with their corresponding class vectors presented at F_1^b, ARAM learns a predictive model (encoded by the recognition nodes in F_2) that associates combinations of key features to their respective classes.

Fuzzy ARAM, based on fuzzy ART operations, has been successfully applied to numerous machine learning tasks, including personal profiling [40], document classification [14], personalized content management [24,38], and DNA gene expression analysis [39]. In many benchmark experiments, ARAM has demonstrated predictive performance superior to those of many state-of-the-art machine learning systems.

Learning by Reinforcement. Reinforcement learning [31] is a paradigm wherein an autonomous system learns to adjust its behaviour based on reinforcement signals received from the environment. An instance of fusion ART, known as FALCON

(Fusion Architecture for Learning, COgnition, and Navigation), learns mappings simultaneously across multi-modal input patterns, involving states, actions, and rewards, in an online and incremental manner.

FALCON employs a three-channel architecture, comprising a category field F_2 and three pattern fields, namely a sensory field F_1^{c1} for representing current states, a motor field F_1^{c2} for representing actions, and a feedback field F_1^{c3} for representing reward values. A class of FALCON networks, known as TD-FALCON [34,37], incorporates Temporal Difference (TD) methods to estimate and learn value function $Q(s, a)$, that indicates the goodness to take a certain action a in a given state s.

The general sense-act-learn algorithm for TD-FALCON is summarized below. Given the current state s, the FALCON network is used to predict the value of performing each available action a in the action set $\mathcal{A}$ based on the corresponding state vector $\mathbf{S}$ and action vector $\mathbf{A}$. The value functions are then processed by an action selection strategy (also known as policy) to select an action. Upon receiving a feedback (if any) from the environment after performing the action, a TD formula is used to compute a new estimate of the Q-value for performing the chosen action in the current state. The new Q-value is then used as the teaching signal (represented as reward vector $\mathbf{R}$) for FALCON to learn the association of the current state and the chosen action to the estimated value.

4.3 Integrating Desire, Intention and Learning

To address the issue of autonomy and self-awareness, a hybrid architecture has been developed that integrates BDI components, including desire and intention, with a reinforcement learning system known as Temporal Difference - Fusion Architecture for Learning and Cognition (TD-FALCON). Following the Belief-Desire-Intention (BDI) framework, the proposed connectionist BDI-FALCON (cBDI-FALCON) architecture consists of three modules, namely the *desire* module, the *intention* module, and a *reactive* module, each of which is implemented as a fusion ART network. The three key modules and their relationships are exemplified in Figure 2. The detailed algorithms and processes are described in [36].

Reactive module: The low level reactive learning module is a TD-FALCON model that interacts with the environment through the sensory, motor, and feedback channels. Based on the goals defined in the desire module and the sensory inputs received from the environment, TD-FALCON performs reinforcement learning so as to acquire a set of action and value policies that enables the agent to achieve its goals.

Intention module: The intention module maintains the plan set and supports the key processes of plan learning, plan selection, plan execution and plan evaluation. Given a set of active goals and the current sensory inputs, the plan selection process identifies the most applicable plan to perform. During the plan execution, the action sequence of the adopted plan is extracted and performed through the motor channel of the reactive module. The execution of plans thus

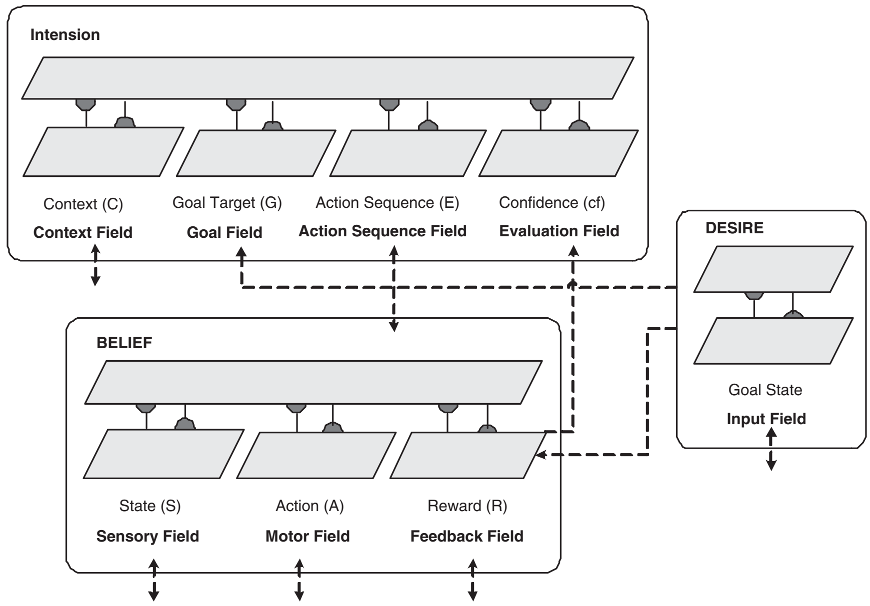

Fig. 2. A schematic architecture of cBDI-FALCON

enables an agent to perform a series of actions without the need of going through the typical sense-act-learn cycle for each action. This could potentially lead to saving in computation cost and enable the system to be more resilient in a challenging environment, wherein external signals may not be available all the time. Through a simple form of reinforcement learning, the plan evaluation process adjusts the confidence value of each adopted plan according to the outcome that it leads to. In contrast to other cognitive architectures, the intention module in the cBDI-FALCON is also modelled as a fusion ART neural network. Doing so enables plans to be learned and updated through reinforcement learning according to the outcomes of their use in a natural way.

Desire module: The desire module maintains an explicit representation of the agent's goals. Active goals are those that give direction to the agent's activities for it to achieve its objectives. By matching the defined goals with the corresponding current state attributes, the desire module computes how well the system has progressed towards the desired goals. The computed degrees of goal attainment in turn serve as reward signals to the feedback field of the reactive module and the evaluation field of the intention modules. Similar to the reactive and intention modules, the desire module in the architecture is also modelled as a one-channel fusion ART network. The overall design philosophy is thus aiming towards a unified framework by using a principled set of computational processes for supporting both intentional and reactive behaviour. In theory, all computations in the system can be operated in parallel on distributed neural networks, enabling the potential of speeding up.

To combine planned and reactive capabilities, we develop two strategies, known as the follow-through and the re-evaluation strategies to coordinate the output produced by the intention and reactive modules. We have conducted extensive experiments to analyze the behaviours of the integrated system, in terms of plan utilization, efficiency, and the overall success rates. Our experimental results on a minefield navigation task show that the integrated neural architecture is able to combine intentional and reactive action execution, leading to improvement both in terms of task completion performance and efficiency.

4.4 Learning Personal Agents with User Modelling

Adaptive Player Modelling. To achieve real-time learning and personalization, we integrate learning personal agents with adaptive user modelling for service recommendation in Co-Space. Our personal agent is based on TD-FALCON [37] that employs a three-channel fusion ART and incorporates Temporal Difference (TD) methods to estimate and learn value functions of its recommendations. For player modelling, we adopt a two-channel fusion ART, that performs supervised learning through the pairing of the input patterns representing the recommendations and teaching signals representing the user's feedback received from the virtual environment. If an initial user profile is available, the model first initializes the player model by associating the attributes specified in the player's profile with positive reward signals. During play time, the player model learns the user's specific like's and dislike's by creating cognitive nodes associating the key attributes of the agent's recommendations to the user's feedbacks. Furthermore, the user's general interest could be inferred based on the frequency of the user's positive responses to recommendations given in a general interest category.

Integrating Player Model With Personal Agent. The overall recommendation agent, incorporating the personal agent, the player model, and a search

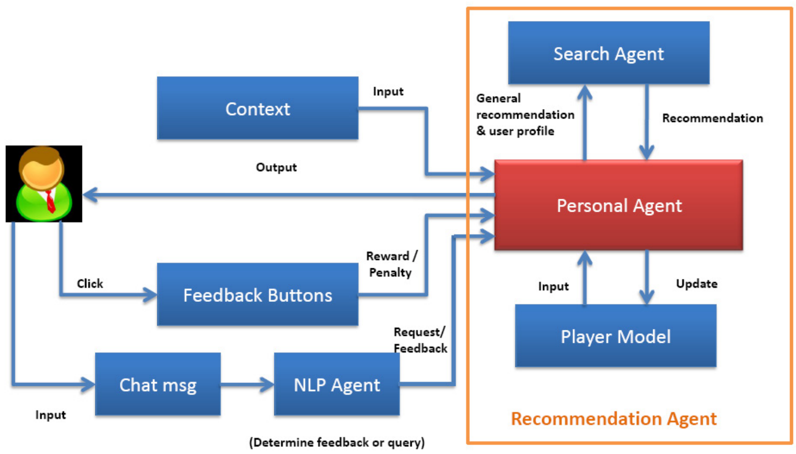

Fig. 3. Personal agent integrated with adaptive player model

agent, is shown in Figure 3. Specifically, the personal agent determines the appropriate class of services to recommend, such as accommodations, restaurants, YOV venues, shopping areas, and other general places of interest, according to the current situation and the user's context. Upon the user's feedback, the player model learns the player's specific preferences and updates the personal agent with the player's current general interests during the interplay. By incorporating the personal agent with the adaptive player model, the system would be more sensitive to the players' habits and eccentricity. Based on the recommendation output of the personal agent and the user preferences indicated by the player model, the search agent handles the retrieval of the requested information from the database.

5 The Youth Olympic Village Co-Space

The Youth Olympic Village (YOV) Co-Space is built to showcase the Youth Olympic Village (YOV) and the hosting country to visitors around the world in an interactive and playable manner. To achieve this objective, we are in the process of developing and populating human-like cognitive agents in the form of autonomous avatars that roam in the landscape of YOV Co-Space. The agents are designed to be aware of its surrounding and can interact with users through their human avatars. With the autonomous avatars befriending and providing personalized context-aware services to human avatars, we aim to make the content and services readily available to the users.

Figure 4 shows the architecture of the YOV Co-Space. As illustrated in this framework, the TD-FALCON based personal agent works in conjunction with the search agent in recommending functions and services to the users. Specifically,

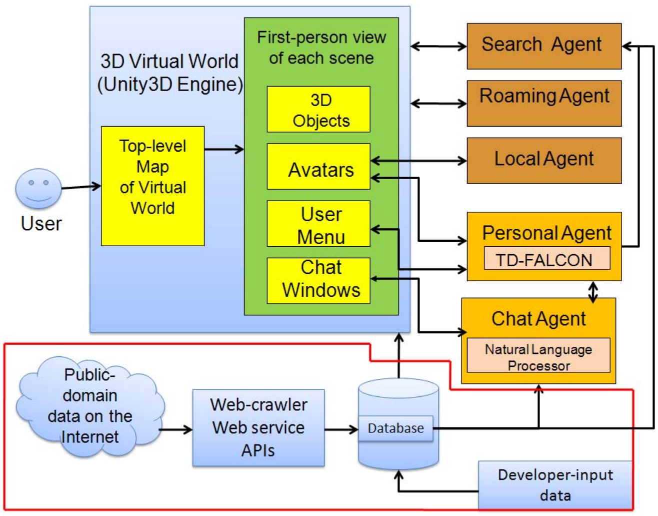

Fig. 4. The architecture of Co-Space

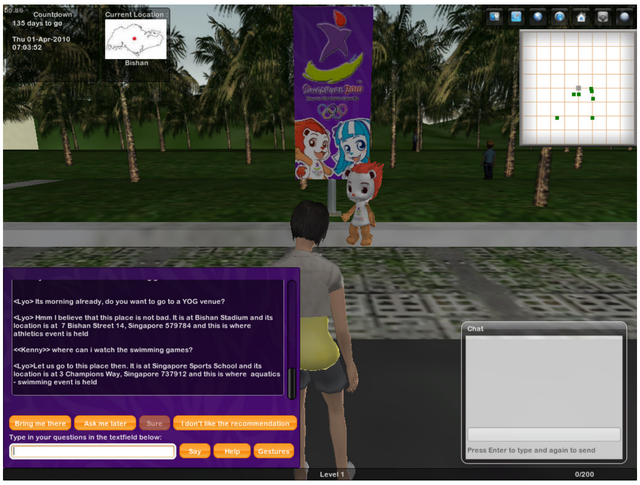

Fig. 5. A snapshot of TD-FALCON based personal agent in Co-Space

the personal agent determines the appropriate type of services to recommend whereas the search agent retrieves the specific services based on the environment situations as well as the users' context parameters. Figure 5 provides a snapshot of the virtual world, showing the personal agent Lyo serving the user. For a more detailed description of the YOV Co-Space and the experimental study, please refer to [16,17].

6 Conclusion

This paper has presented our ongoing work in developing and deploying intelligent agents for Co-Space based on a generalized neural model, known as fusion ART. Using a universal coding mechanism, the proposed model unifies a myriad of traditionally distinct learning paradigms, including unsupervised learning, supervised learning, and reinforcement learning. In fact, ART-style learning and matching mechanism seems to be operative in many levels of the cerebral cortex of the brain. The proposed framework may thus serve as a foundation model for developing high level cognitive information processing capabilities for intelligent agents, including awareness, reasoning, explaining, and surprise handling.

Acknowledgments. The reported work is supported by the Singapore National Research Foundation Interactive Digital Media R&D Program, under research Grant NRF2008IDM-IDM004-037.

References

1. Anderson, J.R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., Qin, Y.: An intergrated theory of the mind. Psychological Review 111, 1036–1060 (2004)
2. Anderson, M.L.: Embodied cognition: A field guide. Artificial Intelligence 149, 91–130 (2003)

3. Bratman, M.E., Israel, D.J., Pollack, M.E.: Plans and resource-bounded practical reasoning. Computational Intelligence 4(4), 349–355 (1988)
4. Brooks, R.A.: Cambrian Intelligence: The Early History of the New AI. MIT Press, Boston (1999)
5. Carpenter, G.A., Grossberg, S.: A massively parallel architecture for a self-organizing neural pattern recognition machine. Computer Vision, Graphics, and Image Processing 37, 54–115 (1987)
6. 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, 759–771 (1991)
7. Carpenter, G.A., Grossberg, S. (eds.): Pattern Recognition by Self-Organizing Neural Networks. MIT Press, Cambridge (1991)
8. Carpenter, G.A., Grossberg, S.: Adaptive Rresonance Theory. In: Arbib, M.A. (ed.) The Handbook of Brain Theory and Neural Networks, pp. 87–90. MIT Press, Cambridge (2003)
9. Chong, H.-Q., Tan, A.-H., Ng, G.-W.: Integrated cognitive architectures: A survey. Artificial Intelligence Review 28(2), 103–130 (2007)
10. Duda, R.O., Hart, P.E., Stock, D.G. (eds.): Pattern Classification, 2nd edn. John Wiley, Chichester (2001) Section 10.11.2
11. Georgeff, M.P., Ingrand, F.F.: Decision-making in an embedded reasoning system. In: Proceedings of IJCAI, pp. 972–978 (1989)
12. Gerhard, M., Moore, D.J., Hobbs, D.J.: Embodiment and copresence in collaborative interfaces. Int. J. Hum.-Comput. Stud. 64(4), 453–480 (2004)
13. Grossberg, S.: How does the cerebral cortex work? development, learning, attention, and 3d vision by laminar circuits of visual cortex. Behavioral and Cognitive Neuroscience Reviews 2, 47–76 (2003)
14. He, J., Tan, A.-H., Tan, C.-L.: On machine learning methods for chinese documents classification. Applied Intelligence 18(3), 311–322 (2003)
15. Jan, D., Roque, A., Leuski, A., Morie, J., Traum, D.: A virtual tour guide for virtual worlds. In: Ruttkay, Z., Kipp, M., Nijholt, A., Vilhjálmsson, H.H. (eds.) IVA 2009. LNCS, vol. 5773, pp. 372–378. Springer, Heidelberg (2009)
16. Kang, Y., Tan, A.-H.: Learning personal agents with adaptive player modeling in virtual worlds. In: Proceedings, IEEE/WIC/ACM International Conference on Intelligent Agent Technology, Toronto, Canada (2010)
17. Kang, Y., Tan, A.-H.: Self-organizing agents for reinforcement learning in virtual worlds. In: Proceedings of IJCNN, Barcelona, Spain, pp. 3641–3648 (2010)
18. Laird, J.E., Rosenbloom, P.S., Newell, A.: Chunking in Soar: The anatomy of a general learning mechanism. Machine Learning 1, 11–46 (1986)
19. Langley, P., Choi, D.: A unified cognitive architecture for physical agents. In: Proceedings of AAAI, pp. 1469–1474 (2006)
20. Lehman, J.F., Laird, J.E., Rosenbloom, P.S.: A gentle introduction to soar, an architcture for human cognition (update 2006)
21. Leuski, A., Traum, D.: A statistical approach for text processing in virtual humans. In: 26th Army Science Conference (2008)
22. Levine, D.S.: Introduction to Neural and Cognitive Modeling, ch. 6. Lawrence Erlbaum Associates, New Jersey (2000)
23. Newell, A.: Unified Theories of Cognition. Harvard University Press, Cambridge (1990)
24. Ong, H.-L., Tan, A.-H., Ng, J., Pan, H., Li, Q.-X.: FOCI: Flexible organizer for competitive intelligence. In: Proceedings of CIKM, Atlanta, pp. 523–525 (2001)

25. Raizada, R., Grossberg, S.: Towards a theory of the laminar architecture of cerebral cortes: Computational clues from the visual system. Cerebral Cortex 13, 200–213 (2003)
26. Rao, A.S., Georgeff, M.P.: Modeling rational agents within a BDI architecture. In: Proceedings of Second International Conference on Principles of Knowledge Representation and Reasoning, pp. 473–484 (1991)
27. Rousseau, D., Roth, B.H.: A social-psychological model for synthetic actors. Stanford Knowledge Systems Laboratory Report KSL-97-07 (1997)
28. Kopp, S., Gesellensetter, L., Krämer, N.C., Wachsmuth, I.: A conversational agent as museum guide - design and evaluation of a real-world application. In: Panayiotopoulos, T., Gratch, J., Aylett, R.S., Ballin, D., Olivier, P., Rist, T. (eds.) IVA 2005. LNCS (LNAI), vol. 3661, pp. 329–343. Springer, Heidelberg (2005)
29. Sun, R., Merrill, E., Peterson, T.: From implicit skills to explicit knowledge: a bottom-up model of skill learning. Cognitive Science 25(2), 203–244 (2001)
30. Sun, R., Peterson, T.: Learning in reactive sequential decision tasks: The CLARION model. In: Proceedings of IEEE ICNN, pp. 1073–1078 (1996)
31. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. MIT Press, Cambridge (1998)
32. Tan, A.-H.: Adaptive Resonance Associative Map. Neural Networks 8(3), 437–446 (1995)
33. Tan, A.-H.: FALCON: A fusion architecture for learning, cognition, and navigation. In: Proceedings of IJCNN, Budapest, Hungary, pp. 3297–3302 (2004)
34. Tan, A.-H.: Self-organizing neural architecture for reinforcement learning. In: Wang, J., Yi, Z., Żurada, J.M., Lu, B.-L., Yin, H. (eds.) ISNN 2006. LNCS, vol. 3971, pp. 470–475. Springer, Heidelberg (2006)
35. Tan, A.-H.: Direct code access in self-organizing neural architectures for reinforcement learning. In: Proceedings of IJCAI, Hyderabad, India, pp. 1071–1076 (2007)
36. Tan, A.-H., Feng, Y.-H., Ong, Y.-S.: A self-organizing neural architecture integrating desire, intention and reinforcement learning. Neurocomputing 73(7-9), 1465–1477 (2010)
37. Tan, A.-H., Lu, N., Xiao, D.: Integrating temporal difference methods and self-organizing neural networks for reinforcement learning with delayed evaluative feedback. IEEE Transactions on Neural Networks 9(2), 230–244 (2008)
38. Tan, A.-H., Ong, H.-L., Pan, H., Ng, J., Li, Q.-X.: Towards personalized web intelligence. Knowledge and Information Systems 6(5), 595–616 (2004)
39. Tan, A.-H., Pan, H.: Predictive neural networks for gene expression data analysis. Neural Networks 18(3), 297–306 (2005)
40. Tan, A.-H., Soon, H.-S.: Predictive adaptive resonance theory and knowledge discovery in database. In: Terano, T., Chen, A.L.P. (eds.) PAKDD 2000. LNCS, vol. 1805, pp. 173–176. Springer, Heidelberg (2000)
41. Thomas, D., Brown, J.S.: Why virtual worlds can matter. International Journal of Media and Learning 1(1), 37–49 (2009)
42. Weinstein, J., Myers, J.: Same principles apply to virtual world expansion as to china and other new markets. Media Village 11 (2006)
43. Weizenbaum, J.: ELIZA: a computer program for the study of natural language communication between men and machines. Communications of the ACM 9 (1996)
44. Yoon, S., Burke, R.C., Blumberg, B.M., Schneider, G.E.: Interactive training for synthetic characters. In: Proceedings of AAAI, pp. 249–254 (2000)