

Self-organizing Cognitive Models for Virtual Agents

Yilin Kang and Ah-Hwee Tan

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

Abstract. Three key requirements of realistic characters or agents in virtual world can be identified as autonomy, interactivity, and personification. Working towards these challenges, this paper proposes a brain inspired agent architecture that integrates goal-directed autonomy, natural language interaction and human-like personification. Based on self-organizing neural models, the agent architecture maintains explicit mental representation of desires, intention, personalities, self-awareness, situation awareness and user awareness. Autonomous behaviors are generated via evaluating the current situation with active goals and learning the most appropriate social or goal-directed rule from the available knowledge, in accordance with the personality of each individual agent. We have built and deployed realistic agents in an interactive 3D virtual environment. Through an empirical user study, the results show that the agents are able to exhibit realistic human-like behavior, in terms of actions and interaction with the users, and are able to improve user experience in virtual environment.

Keywords: Cognitive models, Virtual agents, Self-Organizing neural networks, Autonomy, Personality, Interactivity.

1 Introduction

Three key requirements of realistic characters or agents in virtual world can be identified as autonomy, interactivity, and personification [8]. However, most virtual worlds tend to constrain agents' actions to a very coarse level, dictated by hard coded rules [18,10]. In recent years, there has been growing interest in creating intelligent agents in virtual worlds that do not follow fixed scripts predefined by the developers, but instead react accordingly to actions performed by the players during their interaction. In order to achieve this objective, there have been approaches attempting to model the dynamic environments and user's immediate context [28,5,7]. However, they typically ignore a significant component of making the virtual world experience more intense and personalized for players, namely the capability for the agents to adapt over time to the environment and to the habits as well as eccentricity of a particular player.

Indeed, it has been a great challenge to develop intelligent learning agents that are able to adapt in real time and improve the interactivity and playability in virtual worlds. Learning in a virtual world, just like in the real world, poses many challenges for an agent, not addressed by traditional 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.

In this paper, we present a self-organizing neural model, named FALCON-X (Fusion Architecture for Learning and Cognition - eXtension), for creating intelligent learning agents in virtual worlds. By incorporating FALCON-X, an agent is able to learn from sensory and evaluative feedback signals received from the virtual environment. 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 also able to adapt itself to the variations in the virtual environment and changes in the user behavior patterns.

The FALCON-X model is proposed based on an integration of the Adaptive Control of Thought (ACT-R) architecture [1] and the fusion Adaptive Resonance Theory (fusion ART) neural model [23]. Fusion ART is a generalization of self-organizing neural models known as Adaptive Resonance Theory (ART) [4]. By expanding the original ART model consisting of a single pattern field into a multi-channel architecture, fusion ART unifies a number of network designs supporting a myriad of learning paradigms, including unsupervised learning, supervised learning and reinforcement learning. A specific instantiation of fusion ART known as Temporal Difference-Fusion Architecture for Learning and Cognition (TD-FALCON) has shown to have competitive learning capabilities, compared with gradient descent based reinforcement learning systems [24]. While retaining the structure of the visual, manual, intentional and declarative modules of ACT-R, FALCON-X replaces the symbolic production system with a fusion ART neural network serving as the core inference area for fusing and updating the pattern activities in the four memory buffers. In addition, a critic channel is incorporated to regulate the attentional and learning processes of the core inference area.

FALCON-X may potentially be used to model a wide range of cognitive processes. In this paper, we describe how behavior models can be learned as sensory-motor mappings through reinforcement learning. We have developed learning personal agents using FALCON-X in a 3D virtual world called Co-Space. In this application, the learning personal agents are designed to befriend human users and proactively offer personalized services. Our experiments show that the agents are able to learn player models that evolve and adapt with player during run time. More importantly, the user study shows that the use of intelligent agents can improve user experience in the virtual world.

The rest of this paper is organized as follows: After a brief review of related work in section 2, we present the FALCON-X architecture in section 3. Section 4 describes the generic FALCON-X dynamics, followed by how it may be used to learn procedural knowledge and behaviour model in section 5. Section 6 presents the embodiment of FALCON-X in an integrated agent architecture. The evaluative experiments on the Co-Space simulated domain is reported in section 7. The final section concludes with a highlight of our future direction.

2 Related Work

2.1 Intelligent Virtual Agents

Intelligent agents have been popularly used for improving the interactivity and playability of virtual environment and games. However, most such agents are based on scripts or predefined rules. For example, in the Virtual Theater project, synthetic actors who portray fictive characters are provided by improvising their behaviors. The agents are

based on a scripted social-psychological model which can define personality traits that depend on the values of moods and attitudes [18]. Agents in Metaverse, which was built using Active Worlds, are capable of taking tickets for rides, acting as shopkeepers or other tasks typically associated with humans. However, these agents are basically reactive agents which work in a hard-wired stimulus-response manner. Virtual psychotherapist ELIZA [26], although not even trying to understand its 'patients', often managed to make them feel taken care of, thus demonstrating the effects achievable with rule-based, adeptly modelled 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 [10]. However, the design remains rule-based.

In view of the limitations of static agents, some researchers have adopted learning methods into agents in virtual environment. For example, Yoon et.al. present a Creature Kernel framework to build interactive synthetic characters in the project Sydney K9.0 [28]. 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. To name the most elaborated one, ALICE [25] utilizes a knowledge base containing 40000 input response rules concerning general categories, augmented with knowledge modules for special domains like Artificial Intelligence. This approach has also been employed in other domains, e.g., to simulate co-present agents in a virtual gallery [5]. More recently, an embodied conversational agent that serves as a virtual tour guide in Second Life has been implemented by Jan [7]. Although it learns from past experience, it does not adapt over time according to the habits of a particular player or the changes in the environment.

All the work described above have developed a wide range of agents in virtual world with specific motivations. However, to the best of our knowledge, there have been very few, if any, agents that perform reinforcement learning in real time and can adapt their actions and behaviour during their interaction with the user and environment in virtual world. Our work is motivated by these considerations.

2.2 Cognitive Models

In the fields of artificial intelligence and cognitive science, there has been a debate over symbolic and sub-symbolic (connectionist) representation of human cognition [9], motivating two parallel streams of research directions. The symbolic field holds the view that, the human cognitive system uses symbols as a representation of knowledge and intelligence is through the processing of symbols and their respective constituents. Soar [11], ACT-R [1], and ICARUS [12], for example, are representative systems taking the symbolic approach.

On the other hand, the sub-symbolic camp argues that the human cognitive system uses a distributed representation of knowledge and is capable of processing this distributed representation of knowledge in a complex and meaningful way [6]. Sub-symbolic or connectionist systems are most generally associated with the metaphor of neural models, composing of neural circuits that operate in parallel. The key strengths of sub-symbolic systems lie in their learning abilities and allowance for massively parallel processing.

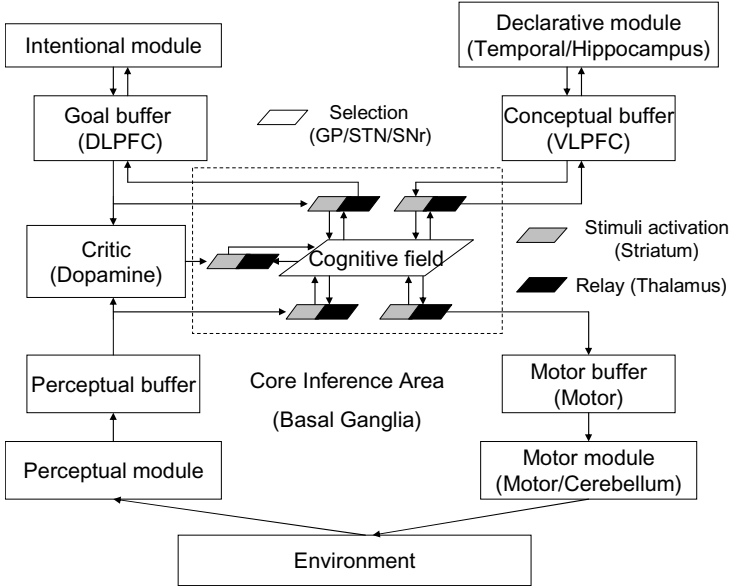


Fig. 1. The FALCON-X architecture

In view of their complementary strengths, there have been great interests in hybrid architectures that integrate high level symbolic systems with sub-symbolic massively parallel processes. Some examples are CLARION [20] and ACT-R with sequence learning [13]. Among the hybrid systems, temporal difference learning using gradient descent based function approximator has been commonly used. However, gradient descent methods typically learn by making small error corrections iteratively. In addition, instability may arise as learning of new patterns may erode the previously learned knowledge.

3 The FALCON-X Architecture

The FALCON-X architecture is presented herein, based on an integration of the ACT-R cognitive architecture and the fusion ART neural model (Figure 1). While retaining the structure of the visual, manual, intentional and declarative modules of ACT-R as the peripheral memory modules, the proposed architecture replaces the symbolic production system with a fusion ART neural network serving as the core inference area for fusing and regulating the pattern activities in the four memory buffers. Furthermore, the visual and manual modules are renamed as the perceptual and motor modules respectively, for the purpose of generality. As a key departure from ACT-R, an explicit critic module is also incorporated, which provides reward signals to the core inference area.

The roles and functions of the various peripheral modules are briefly elaborated as follows.

- The **Perceptual Module** receives input signals from the external environment. In actual applications, some preprocessing of the input signals may be necessary. The input signals are typically represented as a set of vectors of values in the perceptual buffer, taken from the sensors.
- The **Motor Module** receives and executes the actions, produced by a readout action from the core inference area. The actions are typically represented as a set of discrete values in the motor buffer, each of which denotes one of the possible actions.
- The **Intentional Module** consists of the task-relevant goals serving as the context. Each goal is represented as a target state vector in the goal buffer, representing the active goals of the agent.
- The **Declarative Module** consists of middle-term and long-term memories, relevant to the tasks. The memory can be represented in many ways. For example, it can be a look-up table or a neural network.
- The **Critic Module** computes reward signals that indicate the goodness of the actions taken. Generally, there can be two type of critics, namely, reward signals received from the external environment; and estimated payoff computed based on the current states and the target states.
- The **Core Inference Area** receives activations from the five memory modules and acts as a key driver of the inference process. In ACT-R, the production system operates in three processing steps: matching, selection and execution. In FALCON-X, the inference mechanism is realized via a five-step bottom-up and top-down neural processes, namely code activation, code competition, activity readout, template matching and template learning, described in the next section.

The design of FALCON-X is motivated by the neural anatomy of human brains. The core inference area of FALCON-X can be related to *basal ganglia* [16], which are a group of nuclei in the brain interconnected with the cerebral cortex and brainstem. Basal ganglia are important as they have been found to be associated with a variety of cognitive functions, including motor control, cognition, emotions and learning. The main components of basal ganglia includes striatum, globus pallidus (GP), subthalamic nucleus (STN), substantia nigra pars reticulata (SNr) and dopaminergic (DA) neurons.

The cognitive field in FALCON-X, employed for code selection, corresponds to the combined functionality of GP, STN and SNr, as supported in the literatures [17,3]. While ACT-R relates the pattern matching function of the production system to striatum, FALCON-X identifies striatum as the memory fields for stimuli presentation and pattern matching. While ACT-R associates thalamus to the execution function, thalamus is deemed to serve as a relay for motor commands in FALCON-X. Each pattern field of the FALCON is thus considered as a functional combination of striatum and thalamus. The neural substrates of the perceptual, motor, intentional and declarative modules have been discussed extensively in the context of ACT-R [1]. The new critic module in FALCON-X mirrors the dopamine neurons, whose phasic responses are observed when an unexpected reward is presented and depressed when expected reward is omitted [19].

4 The FALCON-X Dynamics

As a natural extension of ART, FALCON-X responds to incoming patterns in a continuous manner. In each inference cycle, the core inference area of FALCON-X receives input signals from the perceptual, intentional and declarative modules, and selects a cognitive node based on a bottom-up code activation and competition process. Whereas the intentional buffer maintains the active goals, the declarative module provides the relevant conceptual memory for code selection. The inference engine may also receive reward signals from the critic module. It is important to note that at any point in time, FALCON-X 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.

Upon activity readout, a template matching process takes place to ensure that the matched patterns in the four memory modules satisfy their respective criterion. If so, a state of resonance is obtained and the template learning process encodes the matched patterns using the selected cognitive node. Otherwise, a memory reset occurs, following which a search for another cognitive node begins. During prediction or action selection, the readout patterns typically include the actions to be executed in the motor module. In other cases, the conceptual memory buffer is updated and the goals may change as a result of inference.

The detailed dynamics of the inference cycle, consisting of the five key stages, namely code activation, code competition, activity readout, template matching, and template learning, are presented as follows.

Input vectors: Let $\mathbf{I}^{ck} = (I_1^{ck}, I_2^{ck}, \dots, I_n^{ck})$ 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 $k = 1, \dots, K$. Let \mathbf{y} denote the F_2 activity vector.

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

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

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

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

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

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

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

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

$$\mathbf{x}^{ck} = \mathbf{I}^{ck} \wedge \mathbf{w}_J^{ck}. \quad (3)$$

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:

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

If any of the vigilance constraints is violated, mismatch reset occurs in which the search process selects another F_2 node J 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:

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

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

5 Learning Procedural Knowledge

In this section, we illustrate how FALCON-X, specifically the core inference area together with the perceptual, motor and critic modules, can acquire procedural knowledge through reinforcement learning in a dynamic and real-time environment.

FALCON-X learns mappings simultaneously across multi-modal input patterns, involving states, actions, and rewards, in an online and incremental manner. Various strategies are available for learning in FALCON-like architectures. We highlight two specific methods, namely reactive learning and temporal difference learning as follows.

5.1 Reactive Learning

A reactive learning strategy, as used in the R-FALCON (Reactive FALCON) model [21], performs fast association between states and actions, based on reward signals. Given a reward signal (positive feedback) in the critic buffer, FALCON associates the current state in the perceptual buffer with the selected action represented in the motor buffer. If a penalty is received, it learns the mapping among current state, the complement pattern of the action taken and the complement value of the given reward.

Table 1. The TD-FALCON algorithm with direct code access

-
1. Initialize the FALCON network.
 2. Sense the environment and formulate a state vector \mathbf{S} based on the current state s .
 3. Following an action selection policy, first make a choice between exploration and exploitation. If exploring, take a random action. If exploiting, identify the action a with the maximal $Q(s, a)$ value by presenting the state vector \mathbf{S} , the action vector $\mathbf{A}=(1,\dots,1)$, and the reward vector $\mathbf{R}=(1,0)$ to FALCON.
 4. Perform the action a , observe the next state s' , and receive a reward r (if any) from the environment.
 5. Estimate the revised value function $Q(s, a)$ following a Temporal Difference formula such as $\Delta Q(s, a) = \alpha(r + \gamma \max_{a'} Q(s', a') - Q(s, a))$.
 6. Formulate action vector \mathbf{A} based on action a and reward vector \mathbf{R} based on $Q(s, a)$.
 7. Present the corresponding state, action, and reward vectors \mathbf{S} , \mathbf{A} , and \mathbf{R} to FALCON for learning.
 8. Update the current state by $s=s'$.
 9. Repeat from Step 2 until s is a terminal state.
-

5.2 Temporal Difference Learning

A key limitation of reactive learning is the reliance on the availability of immediate reward signals. TD-FALCON [24,22] is a variant of FALCON that incorporates Temporal Difference (TD) methods to estimate and learn value functions of action-state pairs $Q(s, a)$ that indicates the goodness for a learning system to take a certain action a in a given state s . Such value functions are then used in the action selection mechanism, also known as the *policy*, to select an action with the maximal payoff. The temporal difference learning algorithm is summarized in Table 1.

Given the current state s , TD-FALCON first decides between exploration and exploitation by following an action selection policy. For exploration, a random action is picked. For exploitation, TD-FALCON performs instantaneous searches for cognitive nodes that match with the current states and at the same time provide the highest reward values using a direct access procedure. Upon receiving a feedback from the environment after performing the action, a TD formula is used to compute a new estimate of the Q value of performing the chosen action in the current state. The new Q value is then used as the teaching signal for TD-FALCON to learn the association of the current state and the chosen action to the estimated Q value.

6 The Integrated Cognitive Agent Architecture

For modelling intelligent virtual agents, FALCON-X needs to be integrated with the necessary peripheral modules for interaction with the environment. As shown in Figure 2, the integrated agent architecture consists of a *Perception Module* receiving situational signals from the environment through a set of sensory APIs and an *Action Module* for performing actions through the various actuator APIs. If the sensory signals involve a text input, the *Chat Understanding Module* interprets the text for the player's

intention. The outputs of *Situational Assessment* and *Chat Understanding Modules* then serve as part of the working memory content providing conditional attributes to the *Inference Engine*. The *Inference Engine* based on the FALCON-X model then identifies the most appropriate action, by tapping a diverse pool of knowledge, in accordance to the desire, intention and personality of the virtual agent. The knowledge learned and used by the Inference Engine include declarative knowledge of self, players, and environment, as well as procedural knowledge of goal-oriented rules, which guide an agent in fulfilling goals, and social rules, for generating socially appropriate behavior. The decision of the *Inference Engine* again forms part of the *Working Memory*, which throughout maintains the context of the interaction. For actions involving a verbal response, the *Natural Language Generation Module* translates the chosen response into natural text for presentation.

Consistent with the view in the state of the art [8], we outline three key characteristics of realistic characters in virtual worlds, namely autonomy, interactivity, and personification, described as follows.

Autonomy. Based on a family of self-organizing neural models known as fusion Adaptive Resonance Theory (ART) [23], the *Inference Engine* of the proposed agent architecture performs a myriad of cognitive functions, including recognition, prediction and learning, in response to a continual stream of input signals received from multiple pattern channels. As a result, an agent makes decisions not only based on the situational factors perceived from the environment but also her mental states characterized by desire, intention and personality. By modelling the internal states of individual agents explicitly, the virtual humans can live a more complete and realistic life in the virtual world.

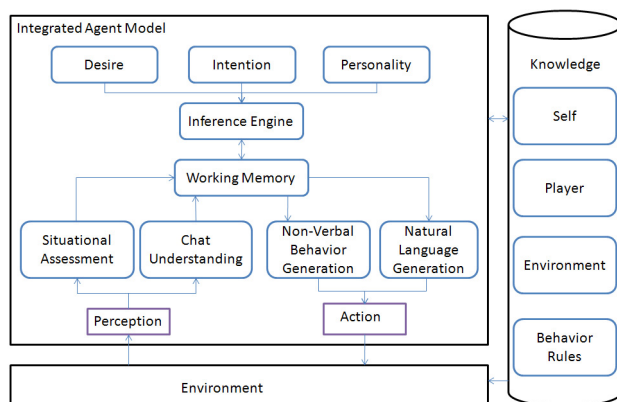


Fig. 2. A schematic of the integrated agent model

Interactivity. For interaction between the agents and the players, an intuitive user interface is provided, through which a player may ask typical questions and provide quick responses by button clicks. The player may also enter free-text sentences via the chat box. The dual communication mode provides the players both ease of use and flexibility. While interacting with player, the agent builds an internal model of the player, with

his/her profile, interests and preferences. The player model in turns allows the agent to make intelligent conversation on topics relevant to the player.

Personification. For improving the believability of virtual humans, our agents adopt the Five Factor Model (FFM) [14], which characterizes personality in five trait dimensions. By giving a weighage to each dimension, a unique personality can be formed by a combination of the traits. Comparing with traditional pattern-matching-based conversational agents, our agents with strong *openness* and *extroversion* personality are warmer and friendlier as they do not stay idle and wait for input queries. Acting pro-actively, they approach the players, offer help, and make conversation.

7 Evaluative Experiments

7.1 Research Methodology

We developed three versions of NTU Co-Space, each with a distinct type of virtual agents in the form of Non-Player Characters (NPCs). The first environment (E1) provides the baseline control condition, wherein the NPCs are only able to display static messages but do not have the capability to interact with the users. The second environment (E2) is the first treatment condition, wherein the virtual humans are designed as embodied conversational agents using the Artificial Intelligence Mark-Up Language (AIML) [25]. AIML is an XML-compliant language that was developed by the Alice-bot free software community. It is considered as a rule based repository of knowledge where the engine performs pattern matching to select the most appropriate output based on the input utterance. We have encoded as many AIML patterns as possible to enhance the conversational abilities of the agents. The third environment (E3) is the second treatment condition, wherein autonomous agents using our proposed fusion ART-based agent model are populated. Although the agents we described in the previous sections have different personalities, for the purpose of this study, we remove the variation in personality by deploying only friendly agents.

The subjects were recruited from an Introduction to Management Information Systems class at a large US university. The scenario given to the subjects is that they were looking for an overseas university for an exchange program and were visiting NTU Co-Space to help them in making the decision. Each subject was asked to complete a quest in the form of a mini-game, where they would experience the key places of the NTU campus through the quest. The quest involves finding five check-points on campus where the clue to each check-point was given at the previous check-point.

The objectives of the experiment are two-fold. First, we observe whether deploying virtual humans in the virtual world will benefit the player's experience. Second, we assess how virtual humans with different levels of intelligence may affect the player's experience, especially in terms of the following constructs, namely *Telepresence* (TP), *Social Presence* (SP), *Perceived Interactivity* (PI), *Perceived Usefulness* (PU), *Flow* (FLW), *Enjoyment* (ENJ) and *Behavioral Intention* (to return to NTU Co-Space) (BI).

Subjects participated in the experiment in a computer lab. They were asked to fill out a pre-questionnaire (used to assess the players' profile, covering demographics information and 3D virtual world skills) and then carry out the experiment by completing the

quest given to them. For subjects in the E1 (control) condition, they completed the quest by using the map in the system to navigate the virtual world, check up information on different parts of campus and teleport to the respective checkpoints without receiving any help from NPCs. For subjects in the E2 (i.e., first treatment) condition, they were not only provided with the interactive map in the E1 condition, but they could also talk to the embodied conversational agents to ask for assistance before teleporting through the interactive map. For subjects in the E3 (i.e., second treatment) condition, in addition to being provided with the interactive map, they were also offered the assistance of fusion ART-based NPCs that have the ability of performing autonomous behaviors both in proactive and responsive ways; moreover, since these NPCs are embedded with a Natural Language Processing module, they can understand input sentences in a flexible way. Hence, subjects were able to interact with the intelligent autonomous NPCs to request for and obtain the information they needed. Because the NPCs are autonomous, they could even offer teleport service to the specific locations requested. After the subjects have completed the quest, they filled out a post-questionnaire which assessed their experience. The subjects were first asked for the acronym followed by the full name of the university to assess their level of recall. The main part of the questionnaire then captured the subjects' assessment of the seven constructs (as described earlier) related to NPC functions. In addition, 3D virtual world skills were also captured to examine their perceived improvement of skills after experiencing NTU Co-Space.

Among the various constructs, *Perceived Usefulness* could be objectively assessed through the time taken to complete the Amazing Quest. The less time spent to complete the Amazing Quest, the more useful the agents are. *Flow* was captured through providing the description of *Flow*, and then asking subjects to rate the degree of *Flow* they experienced and the frequency in which they experienced *Flow*. The other constructs were captured using measurement items. At least five items were used to measure each construct. The items used to measure *Telepresence*, *Enjoyment* and *Behavioral Intention* were derived from Nah [15]. The scale for measuring *Social Presence* and *Interactivity* were adopted from Animesh et al [2]. Given the limited space, we present a sample set of these items in Table 2.

Table 2. A sample set of Post-Questionnaire

Item	Measurement
TP 1	I forgot about my immediate surroundings when I was navigating in the virtual world.
SP 1	During the virtual tour, the interaction with the virtual humans were warm.
ENJ 1	I found the virtual tour to be fulfilling.
BI 1	I would consider visiting this virtual world site again.

7.2 Data Analysis

Overall Performance: Table 3 shows the overall performance in the three environments. We observe that subjects in all the three environments were able to complete the quest successfully. However, subjects in E3 spent the least amount of time and the percentage of subjects who could correctly recall the acronym of the campus is higher

than those of the other two environments. This indicates that the autonomous agents deployed in E3 are more useful in helping the subjects than those of the other two environments.

Table 3. overall performance

Evaluation Measures	E1	E2	E3
% of players complete the quest	100%	100%	100%
Time to complete the quest	20m 31s	25m 32s	16m 56s
% of Players recall the acronym of the campus	44%	25%	45%

Descriptive Statistics: Table 4 shows the means, standard deviations (SD), and confidence intervals (CI, with a confidence level of 95%) of ratings in E1, E2, and E3 in terms of *Telepresence* (TP), *Social Presence* (SP), *Perceived Interactivity* (PI), *Enjoyment* (ENJ), *Flow* (FLW) and *Behavior Intention* (BI). All of the constructs were measured using the seven point Likert scale. From the table, we observe that for E3, the rating of *Telepresence*, *Social Presence*, *Perceived Interactivity*, *Flow* and *Behavior Intention* have better results than E1 and E2. This means by employing autonomous agents, these factors are perceived to be stronger than those in the environment with dummy and AIML based agents. However, we note that for enjoyment, the rating in E2 is the best of the three environments. Referring to the 3D virtual world skills assessed in the pre-questionnaire, subjects in E2 appear to have the best 3D virtual world skills compared to the other two. As prior work have shown that a higher level of skill is likely to enhance the feeling of enjoyment [27], we believe the rating obtained in E2 could be affected by the higher level of 3D virtual world skills.

Table 4. Descriptive Statistics

Constructs	E1			E2			E3		
	Mean	SD	CI	Mean	SD	CI	Mean	SD	CI
TP	3.86	0.34	3.64-4.08	4.08	0.49	3.76-4.40	4.41	0.42	4.14-4.68
SP	3.63	0.44	3.28-3.98	3.82	0.43	3.47-4.16	3.84	0.65	3.32-4.36
PI	4.66	0.67	4.13-5.19	4.37	0.58	3.90-4.83	5.02	0.66	4.49-5.54
ENJ	4.42	0.44	4.04-4.81	4.50	0.33	4.11-4.90	4.31	0.72	3.68-4.94
FLW	4.47	0.54	4.00-4.94	4.21	0.28	3.96-4.46	4.58	0.68	3.98-5.18
BI	4.23	0.71	3.61-4.56	4.51	0.30	4.24-4.78	4.58	0.40	4.23-4.93

Furthermore, a one-way analysis of variance (ANOVA) is used to analyze the results. Specifically, the F-test is used to evaluate the hypothesis of whether there are significant differences among the statistic data means for those constructs. The F values are calculated by the variances between conditions divided by the variance within the conditions. The p values, on the other hand, represent the probability of test statistic being different from the expected values and are directly derived from the F test. A small p value thus indicates a high confidence that the values of those constructs are different. A summary

Table 5. F-test result

Constructs	E1, E2 & E3		E1 & E3		E2 & E3	
	F	p	F	p	F	p
TP	5.47	0.004	11.22	0.001	4.55	0.034
SP	0.58	0.561	1.10	0.295	0.01	0.903
PI	5.71	0.004	4.00	0.047	12.12	0.001
ENJ	0.15	0.862	0.33	0.566	0.17	0.68
FLW	1.62	0.199	0.28	0.595	3.12	0.079
BI	1.23	0.294	2.40	0.120	0.10	0.751

of the F values and p values among E1, E2 and E3, between E1 and E3, and between E2 and E3 are given in Table 5.

This data analysis revealed the significant effects of the three kinds of virtual humans in virtual worlds on *Telepresence* and *Perceived Interactivity*: $F(2, 519) = 5.47, p < 0.01$ for *Telepresence* and $F(2, 345) = 5.71, p < 0.01$ for *Perceived Interactivity*, where the two parameters enclosed in parentheses after F indicate the degrees of freedom of the variances between and within conditions respectively. Consistent with the statistics in Table 4, the fusion ART-based virtual human generates higher levels of *Telepresence* and *Perceived Interactivity* than the other two types of virtual humans, with a mean of 4.41 for E3 (versus 3.86 for E1 and 4.08 for E2) for *Telepresence*, and a mean of 5.02 for E3 (versus 4.66 for E1 and 4.37 for E2) for *Perceived Interactivity*. Although the effect of E1, E2 and E3 on *Flow* is smaller than that of *Telepresence* and *Perceived Interactivity*, the difference in *Flow* is perceived to be marginally significant between E2 and E3, with $F(1, 198) = 3.12, p < 0.1$, and a mean of 4.58 for E3 (versus 4.21 for E2). This means the *Flow* experience perceived by subjects who interacted with the fusion ART-based virtual human is stronger than those interacting with the AIML based virtual humans. No significant difference was found for the rest of the constructs.

8 Conclusion

For creating realistic agents in virtual world, this paper has proposed a cognitive agent architecture that integrates goal-directed autonomy, natural language interaction and human-like personality. Extending from a family of self-organizing neural models, the agent architecture maintains explicit mental representation of desires, personalities, self-awareness, situation awareness and user awareness.

We have built and deployed realistic agents in an interactive 3D virtual environment. We have also carried out systematic empirical work on user study to assess whether the use of intelligent agents can improve user experience in the virtual world. Our user study has so far supported the validity of our agent systems. With the virtual characters befriending and providing personalized context-aware services, players generally found virtual world more fun and appealing. To the best of our knowledge, this is perhaps one of the few in-depth works on building and evaluating complete realistic agents in virtual worlds with autonomous behavior, natural interactivity and personification. Moving forward, we wish to extend our study by completing the agent architectures with more functionalities, such as emotion and facial expressions.

References

1. Anderson, J.R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., Qin, Y.: An integrated theory of the mind. *Psychological Review* 111, 1036–1060 (2004)
2. Animesh, A., Pinsonneault, A., Yang, S.-B., Oh, W.: An odyssey into virtual worlds: Exploring the impacts of technological and spatial environments. *MIS Quarterly* 35, 789–810 (2011)
3. Bogacz, R., Gurney, K.: The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Computation* 19(2), 442–477 (2007)
4. Carpenter, G.A., Grossberg, S.: Adaptive Resonance Theory. In: *The Handbook of Brain Theory and Neural Networks*, pp. 87–90. MIT Press (2003)
5. Gerhard, M., Moore, D.J., Hobbs, D.J.: Embodiment and copresence in collaborative interfaces. *Int. J. Hum.-Comput. Stud.* 64(4), 453–480 (2004)
6. Haykin, S.: *Neural Network: A Comprehensive Foundation*. Prentice Hall (1999)
7. 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)
8. Kasap, Z., Thalmann, N.: Intelligent virtual humans with autonomy and personality: State-of-the-art. *Intelligent Decision Technologies* 1, 3–15 (2007)
9. Kelley, T.D.: Symbolic and sub-symbolic representations in computational models of human cognition: What can be learned from biology? *Theory and Psychology* 13(6), 847–860 (2003)
10. 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)
11. Laird, J.E., Newell, A., Rosenbloom, P.S.: Soar: An architecture for general intelligence. *Artificial Intelligence* 33, 1–64 (1987)
12. Langley, P., Choi, D.: A unified cognitive architecture for physical agents. In: *Proceedings of 21st National Conference on Artificial Intelligence*, pp. 1469–1474 (2006)
13. Lebiere, C., Wallach, D.: Sequence learning in the act-r cognitive architecture: Empirical analysis of a hybrid model. In: Sun, R., Giles, C.L. (eds.) *Sequence Learning. LNCS (LNAI)*, vol. 1828, pp. 188–212. Springer, Heidelberg (2000)
14. McCrae, R., Costa, P.: An introduction to the five-factor model and its applications. *Journal of Personality* 60, 172–215 (1992)
15. Nah, F., Eschenbrenner, B., DeWester, D.: Enhancing brand equity through flow and telepresence: A comparison of 2d and 3d virtual worlds. *MIS Quarterly* 35, 731–748 (2011)
16. O’Reilly, R.C., Frank, M.J.: Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation* 18, 283–328 (2006)
17. Prescott, T.J., Gonzalez, F.M.M., Gurney, K., Humphries, M.D., Redgrave, P.: A robot model of the basal ganglia: Behavior and intrinsic processing. *Neural Networks* 19(1), 31–61 (2006)
18. Rousseau, D., Roth, B.: A social-psychological model for synthetic actors. In: *Proceedings of 2nd International Conference on Autonomous Agents*, pp. 165–172 (1997)
19. Schultz, W.: Getting formal with dopamine and reward. *Neuron* 36(2), 241–263 (2002)
20. 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)
21. Tan, A.-H.: FALCON: A fusion architecture for learning, cognition, and navigation. In: *Proceedings of International Joint Conference on Neural Networks*, pp. 3297–3302 (2004)
22. Tan, A.-H.: Direct Code Access in Self-Organizing Neural Networks for Reinforcement Learning. In: *Proceedings of International Joint Conference on Artificial Intelligence*, pp. 1071–1076 (2007)

23. Tan, A.-H., Carpenter, G.A., Grossberg, S.: Intelligence through interaction: Towards a unified theory for learning. In: Liu, D., Fei, S., Hou, Z.-G., Zhang, H., Sun, C. (eds.) ISSN 2007, Part I. LNCS, vol. 4491, pp. 1094–1103. Springer, Heidelberg (2007)
24. 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)
25. Wallace, R.S.: The anatomy of A.L.I.C.E. Tech. report, ALICE AI Foundation (2000)
26. Weizenbaum, J.: ELIZA: a computer program for the study of natural language communication between men and machines. *Communications of the ACM* 9 (1996)
27. Yi, M.Y., Hwang, Y.: Predicting the use of web-based information systems: self-efficacy, enjoyment, learning goal orientation, and the technology acceptance model. *International Journal of Computer-Human Studies* 59, 431–449 (2003)
28. Yoon, S., Burke, R.C., Blumberg, B.M., Schneider, G.E.: Interactive training for synthetic characters. In: AAI, pp. 249–254 (2000)