Keynote at The11th Multi-disciplinary International Workshop on Artificial Intelligence (MIWAI 2017)

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

Towards Self-Awareness in Artificial Intelligence Systems

From Universal Learning to Memory Modelling and Self-Awareness in Human-Like AI Systems

Ah-Hwee Tan

(http://www.ntu.edu.sg/home/asahtan)

School of Computer Science and Engineering
Nanyang Technological University

Brunei
20 November 2017

Based on Joint work with

Collaborators:
Gail A. Carpenter, Stephen Grossberg, Jacek Zurada, Janusz A. Starzyk, Fiona Nah, Yuan-Sin Tan, Gee-Wah Ng, Loo-Non Teow, …

Research Staff: *Di Wang, Teck-Hou Teng, Budhitama Subagdja, Yilin Kang*

Students: *Dan Xiao, Wenwen Wang, Su Feng, Hui-Qing Chong, Ning Lu, …*

Outline of Talk

- Background/Motivations
- Our Approach to Brain-Inspired AI
- Neural Models for Universal Learning
- Neural Modelling of Autobiographical Memory
  - *Non-player Characters (NPC) in Unreal Tournament*
- A preliminary model for Self-Awareness
- Challenges ahead

Research Motivations

Two Key Research Questions

1. *How do our brains work?*
2. *How to develop systems/agents with high-level cognitive capabilities, based on computational but biologically-plausible neural networks?*

Main capabilities of interests

- *Learning*
- *Memory*
- *Situation awareness*
- *Reasoning*
- *Self-Awareness*
- *…*

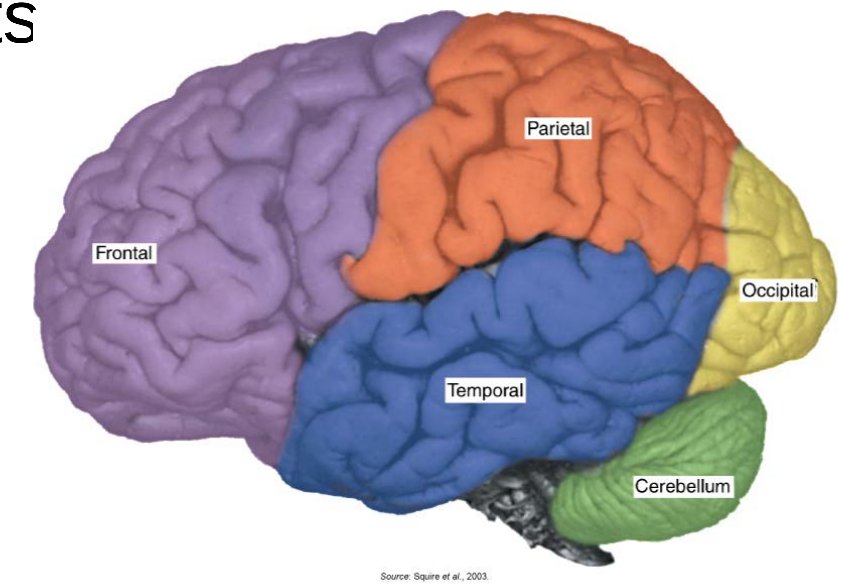

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

Embodied Cognition (Anderson, 2003)

- *Cognition is a process deeply rooted in the body's interaction with the world*

i.e. "Intelligence through interaction"

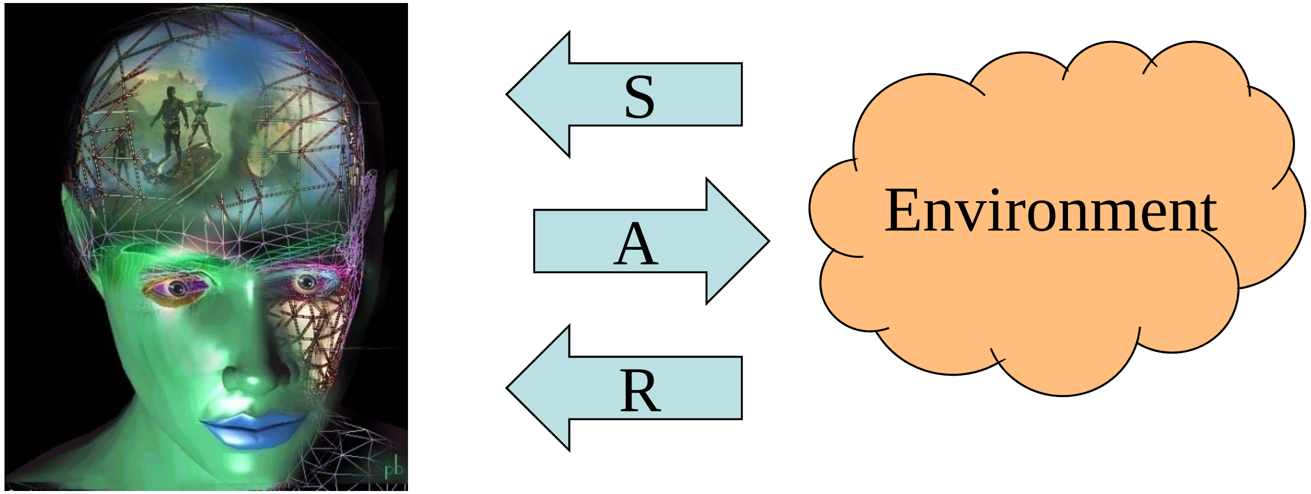

Sense, Act, and Reward Cycle

cf: MDP (Markov Decision Processes)

Stability-Plasticity Dilemma (Grossberg, 76a,b)

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

- Real world presents a challenging situation, where (sensory) data is continuously changing

- How can we continue to quickly learn new things about the environment (plasticity) and yet not forgetting what we have already learned (stability)?

Adaptive Resonance Theory (ART)
(Carpenter & Grossberg, 80's, 90's)

- Self-organizing systems – Unsupervised Learning
- Design to handle Stability-Plasticity dilemma

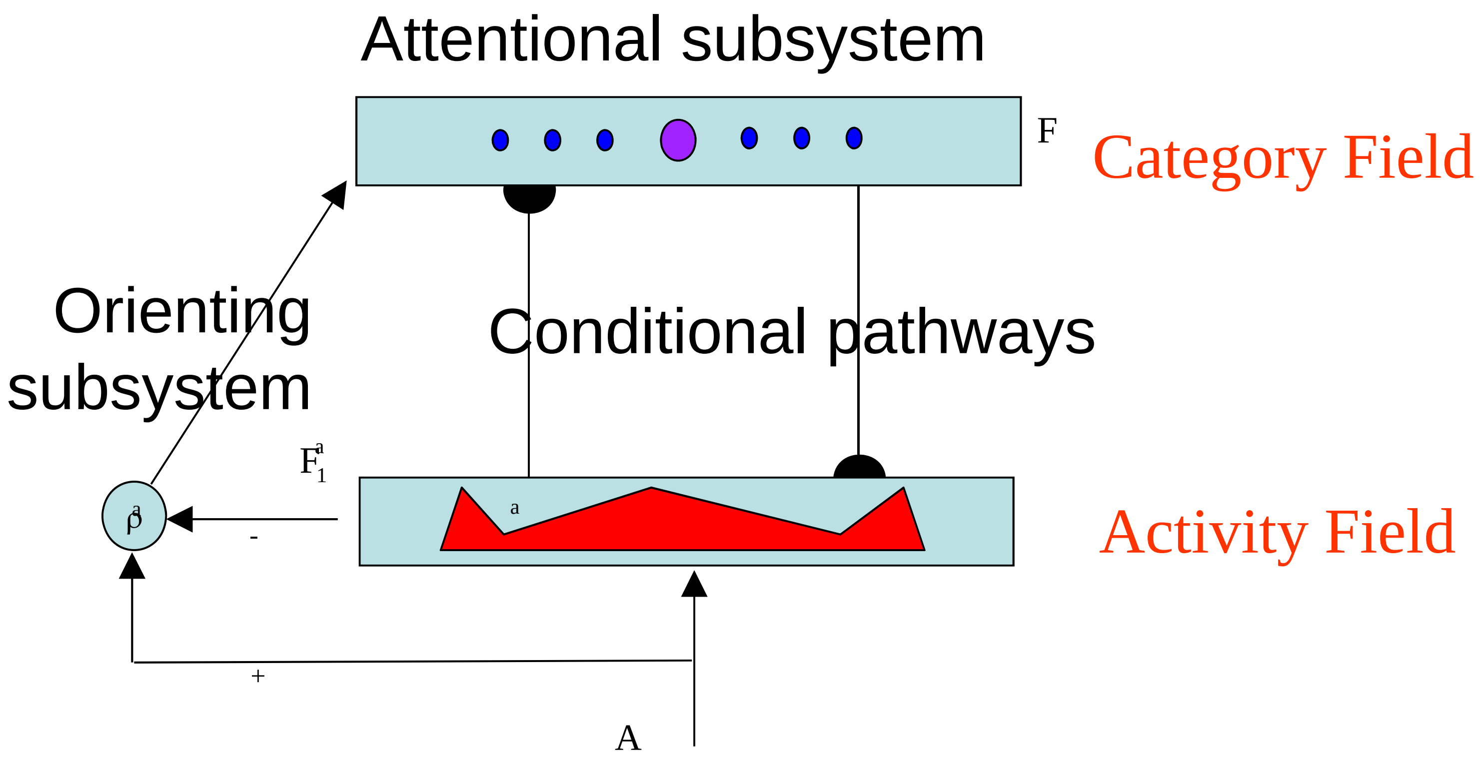

Life is like a jungle!!!

Many forms of learning required

- Unsupervised learning
- Supervised learning
  - Example-based
  - Rule-based
- Reinforcement Learning
- Case-based Learning
- Learning by Imitation
- ...

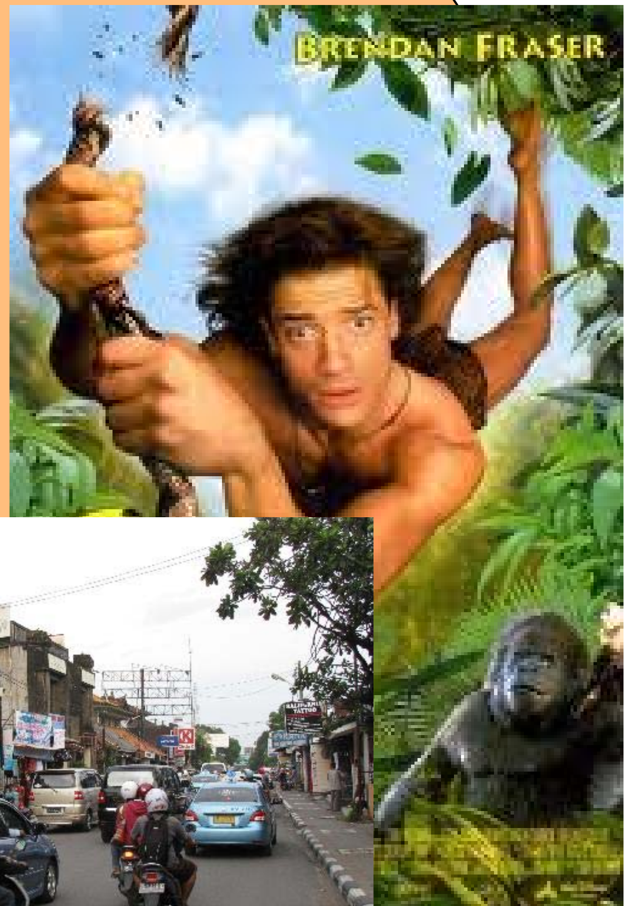

Machine Learning Paradigms

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

Unsupervised Learning

K-means,
FCM, SOM, ART,

Can we have a unified learning theory encompassing the three?

DT, BP, RBF,
LVQ, SVM, …

DP, Q-learning,
Adaptive Critic Design,
…

Supervised Learning

Reinforcement Learning

Supervised Learning: Adaptive Resonance Associative Map

cf: ARTMAP (Carpenter et al 92)

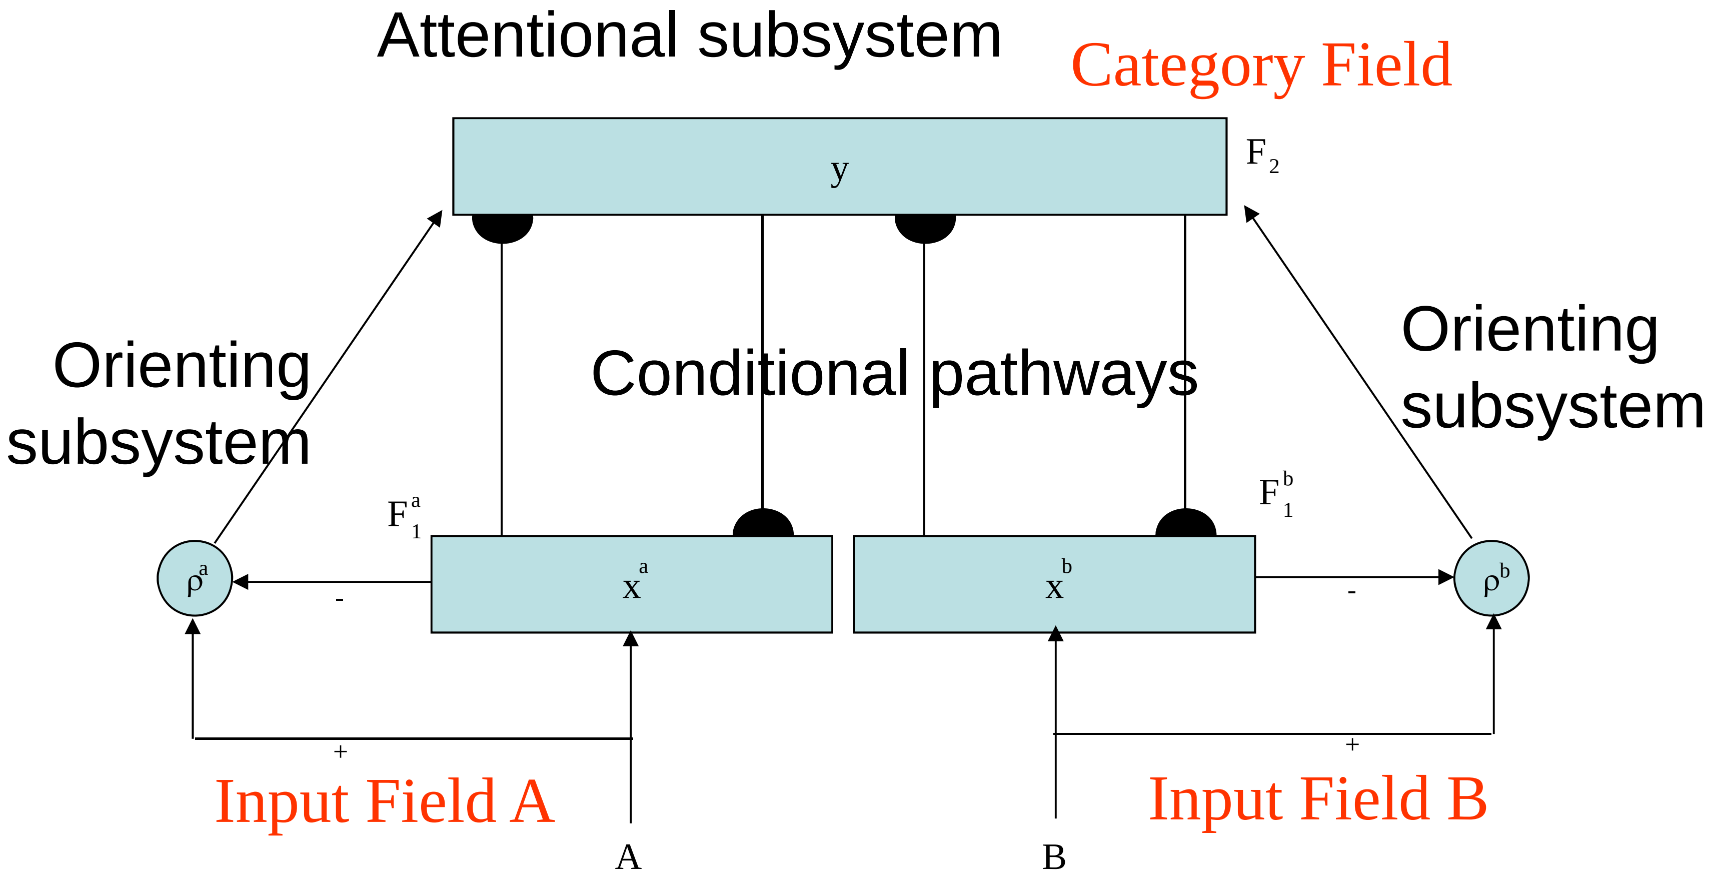

Tan (1995) Adaptive Resonance Associative Map
Neural Networks 8 (1995) 437-446

Predictive ART for Gene Expression Data Analysis

NANYANG TECHNOLOGICAL UNIVERSITY
SINGAPORE

Predictive Neural Networks
for Gene Expression Data Analysis

Ah-Hwee Tan & Hong Pan
Neural Networks 18 (2005) 297-306

A*STAR Institute for Infocomm Research

A*STAR Genome Institute of Singapore

Fusion Architecture for Learning and Cognition (FALCON) (Tan, IJCNN'04)

- Self-organizing neural network for learning cognitive nodes across multi-modal pattern fields
- Compatible with rule-based representation

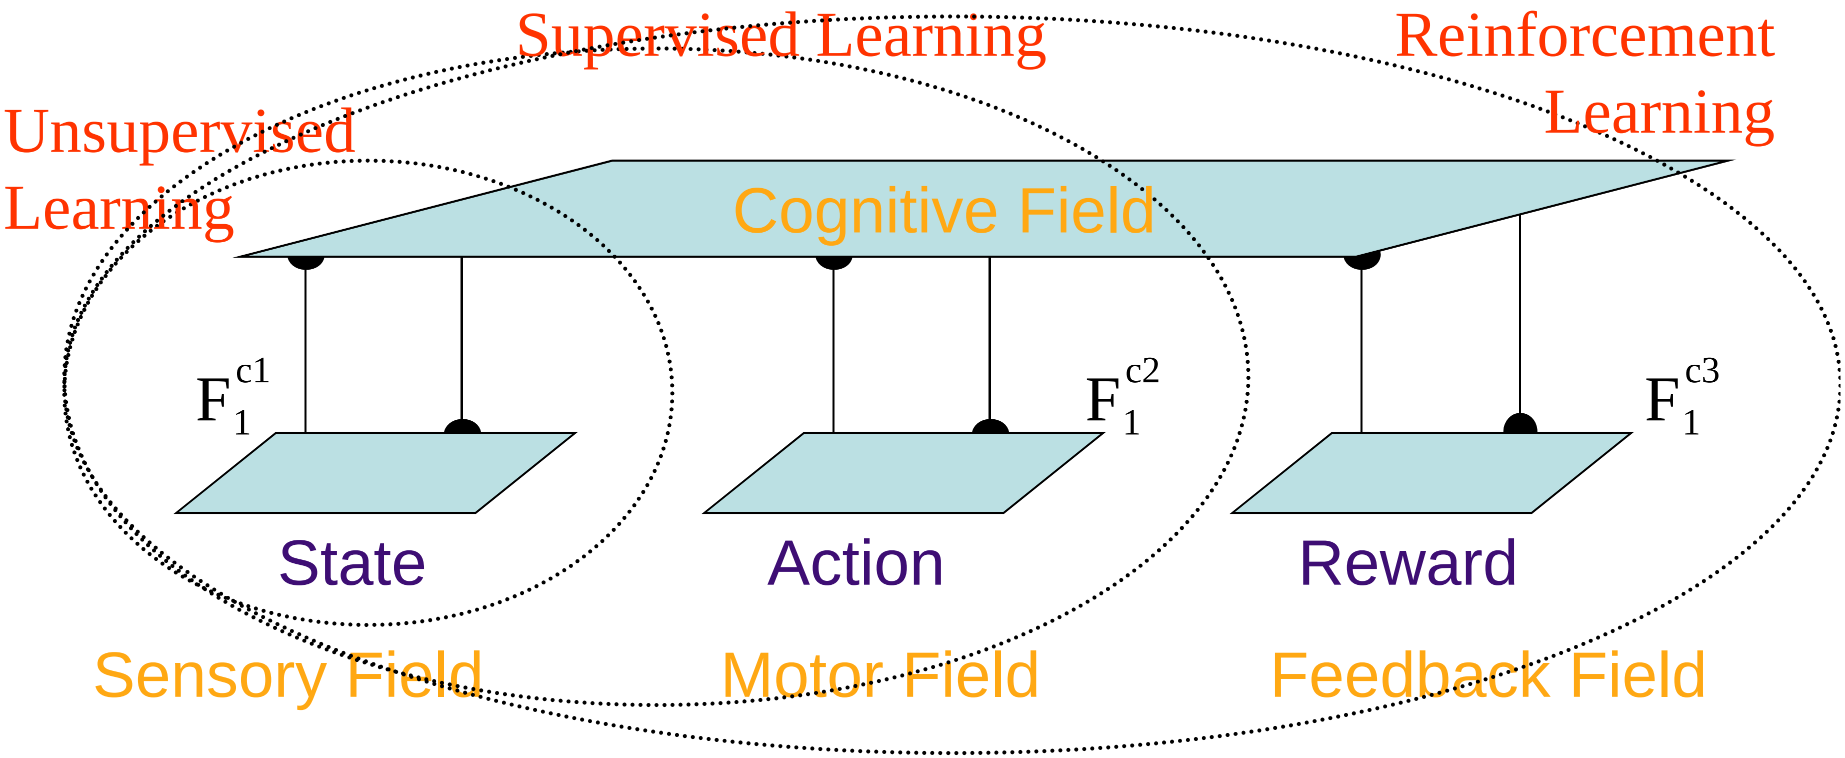

Temporal Difference - FALCON

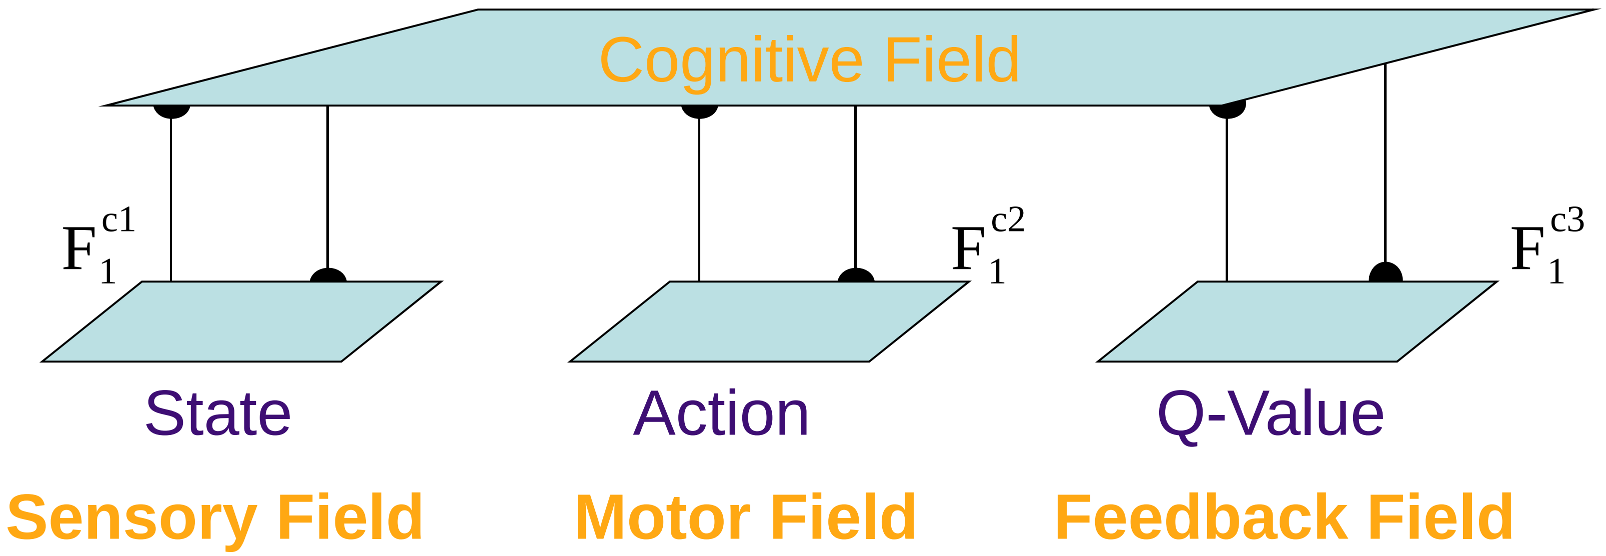

- *Use temporal difference learning rule to estimate future value of performing an action in a state*
- *Useful for situations without immediate rewards*
- *Fast action searching through direct code access*

Ah-Hwee Tan, Ning Lu and Dan Xiao. ***Integrating Temporal Difference Methods and Self-Organizing Neural Networks for Reinforcement Learning with Delayed Evaluative Feedback*****.** *IEEE Transactions on Neural Networks, Vol. 9 (2008), No. 2, 230-244.*

TD-FALCON with Direct Code Access
(Tan, IJCAI'07)

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

1. Initialize the FALCON network
2. Sense the environment and formulate a state representation *s*.
3. Following an action selection policy
   - If exploring, take a random action.
   - If exploiting, present the state vector to TD-FALCON to identify action *a* with max *Q(s,a)* for situation *s*.
4. Perform action *a*, observe next state *s'*, and receive a reward *r*.
5. Estimate revised value function *Q(s,a)*.
6. Present the state, action, and reward (Q-value) vectors to TD-FALCON for learning.
7. Update the current state by *s=s'.*
8. Repeat from Step 2 until *s* is a terminal state.

Ah-Hwee Tan. ***[Direct Code Access in Self-Organizing Neural Architectures for Reinforcement Learning](#)*** ***. IJCAI 2007****, pp. 1071-1076, Hyderabad, 2007*

Direct Code Access for Action Selection

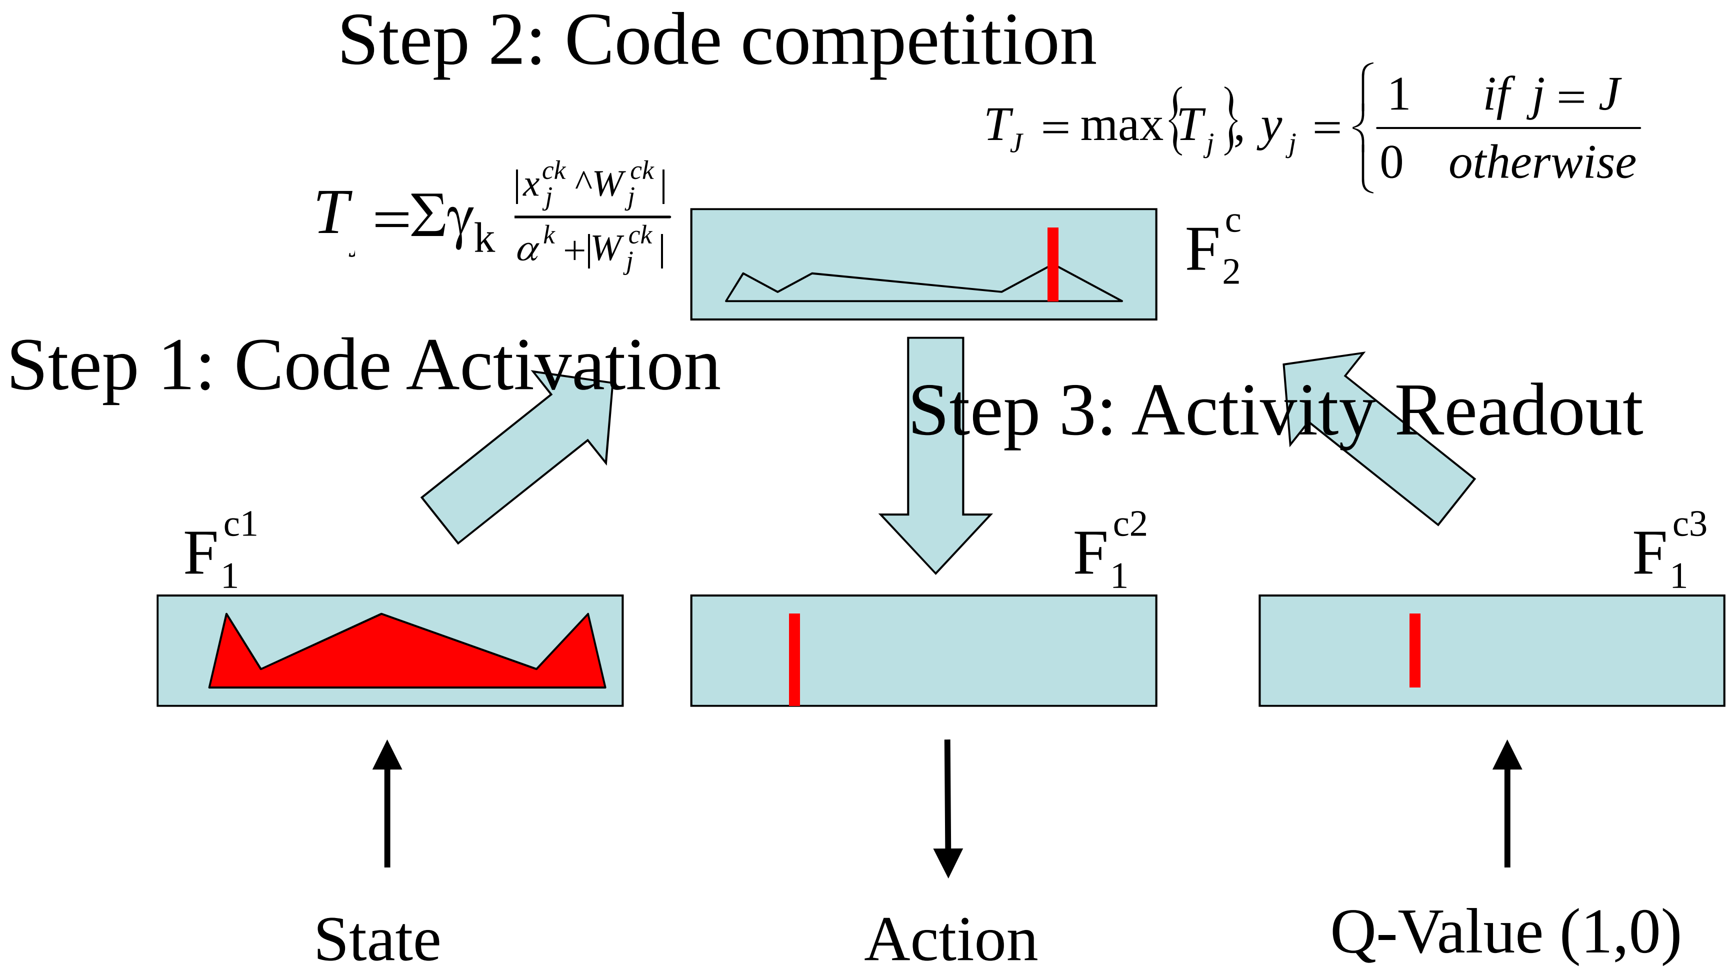

Learning Value Function

Step 4: Template learning

$$w_j^{ck(new)} = (1-\beta^a)w_j^{ck(old)} + \beta^a(x^{ck} \wedge w_j^{ck(old)})$$

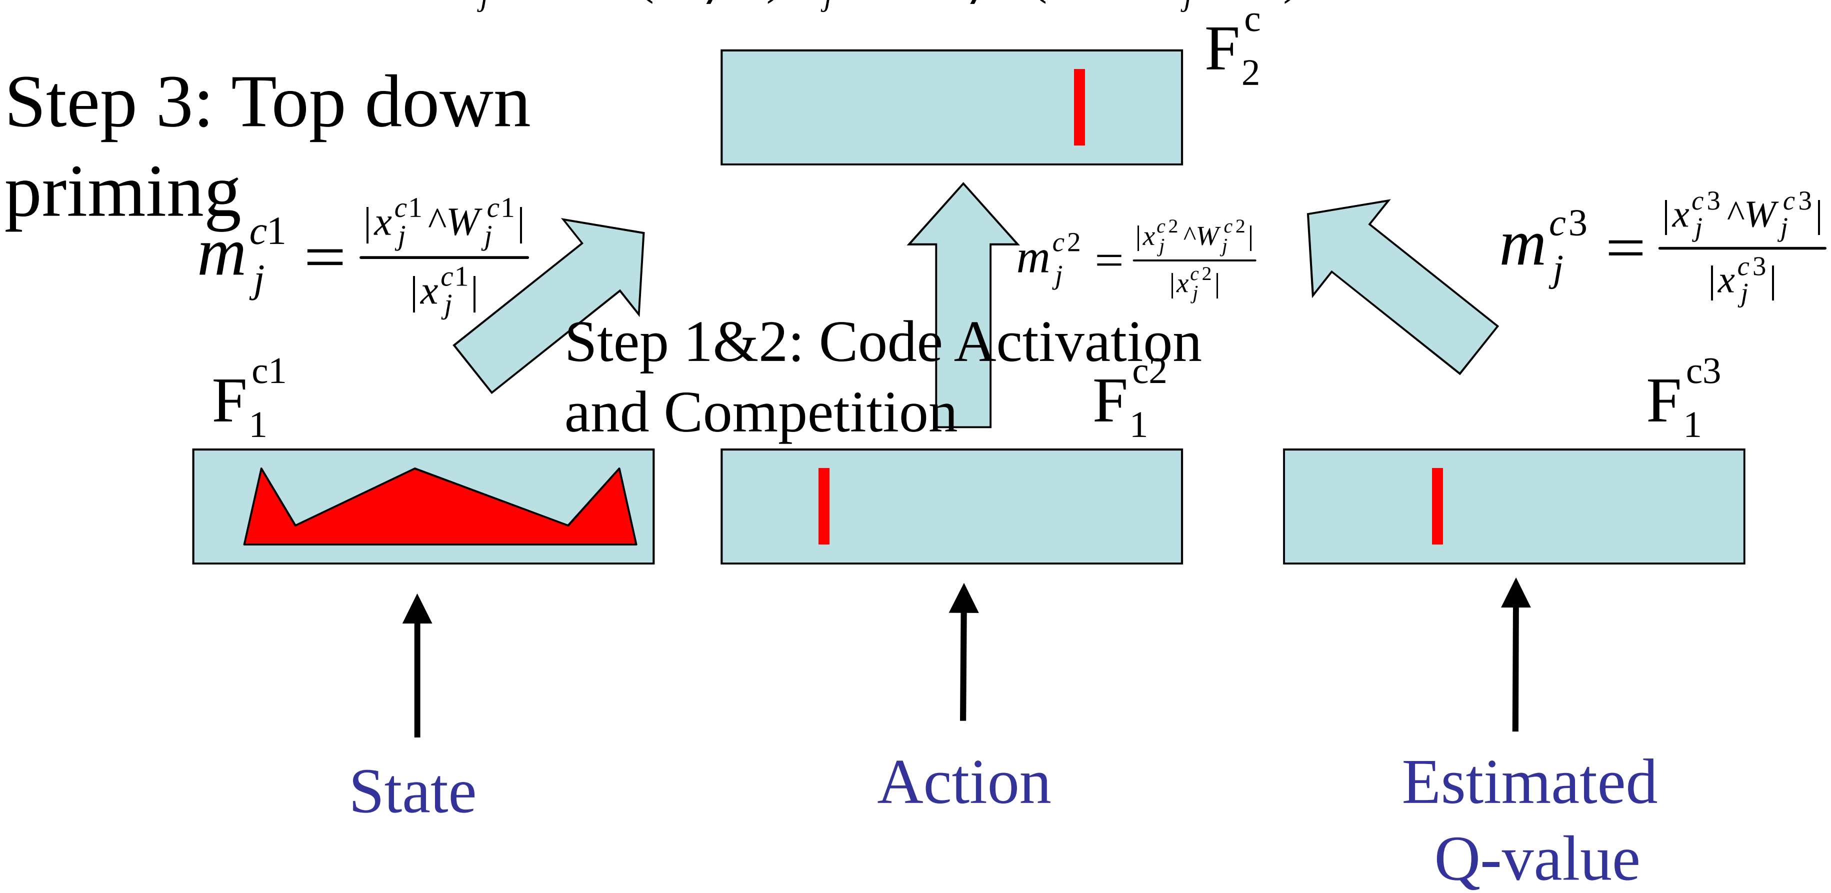

Minefield Navigation Task

- N x N field with M mines
- Different configuration for each run
- Objective:
  Reach target from a random starting point
- Reach target – Success
  Hit Mine – Failure
  Out of Time – Failure

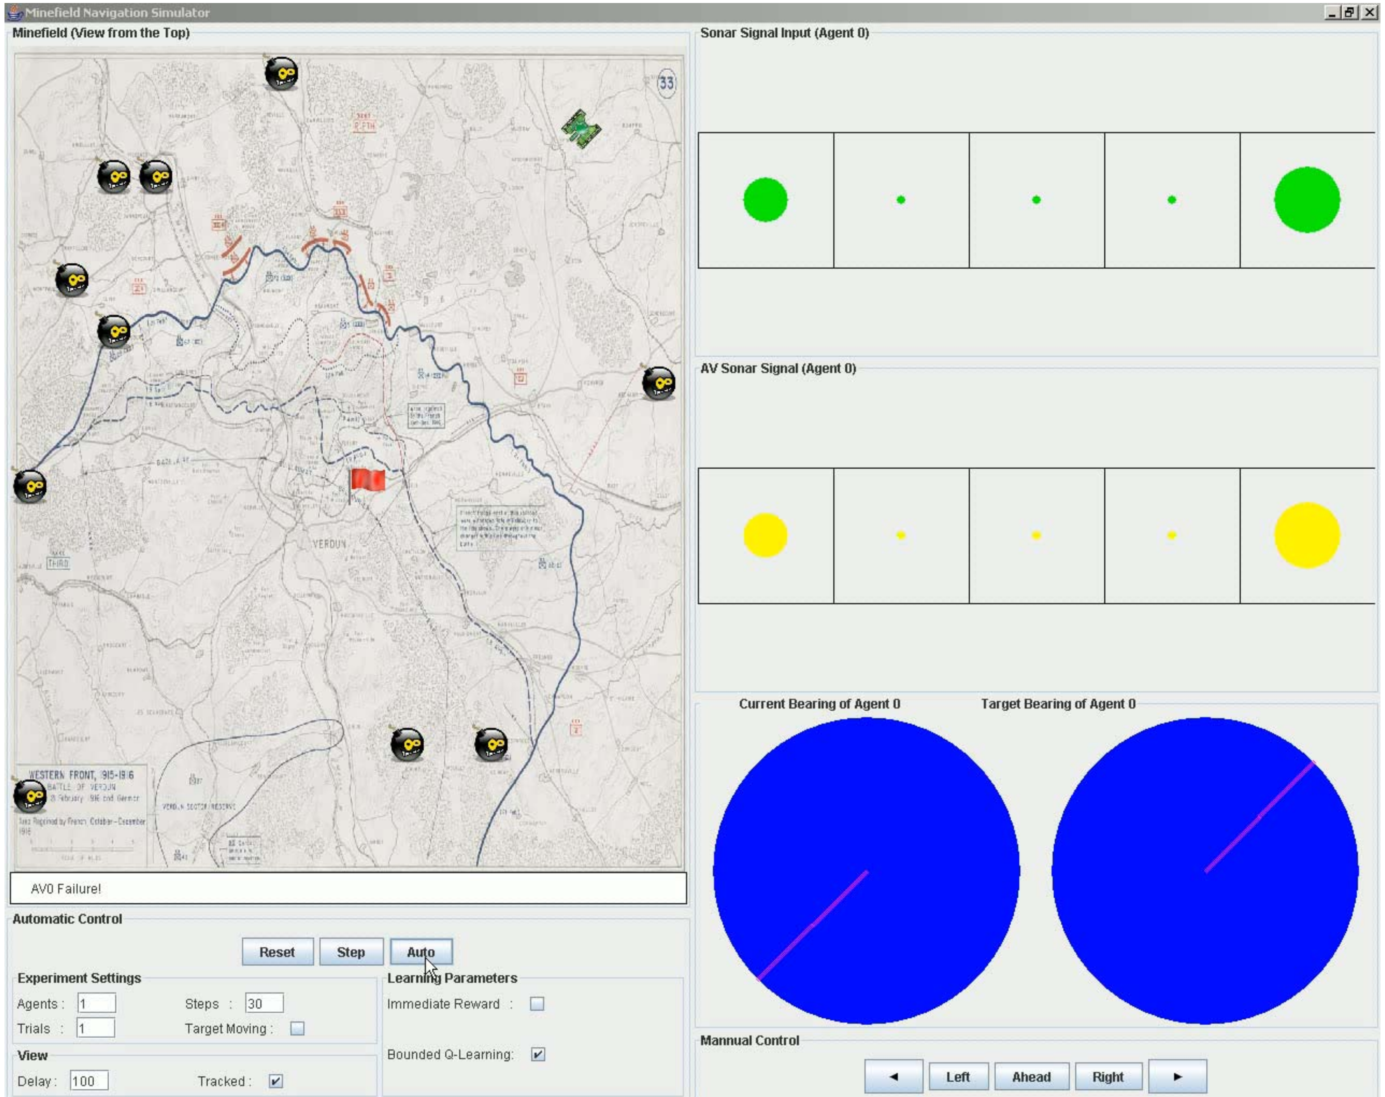

Autonomous Adaptive Agents in UT2004 FPS Game

- UT2004 is a commercial First-Person Shooter (FPS) computer game that allows embodiments of virtual agents for combats.
- Pogamut is an IDE for rapid development and provides sample bots.

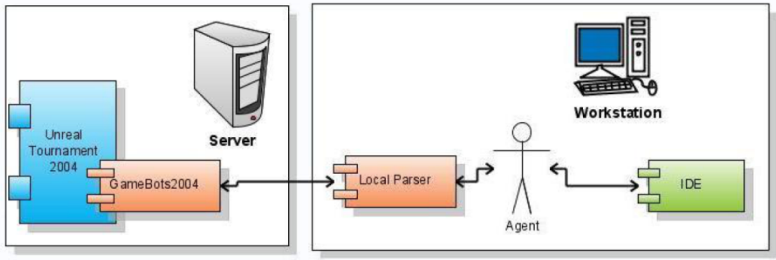

Creating Autonomous Adaptive Agents in a Real-Time First-Person Shooter Computer Game

Di Wang & Ah-Hwee Tan,
IEEE Transaction on Computational Intelligence and AI in Games, Vol. 7, No. 2, 2015

Using FALCON Networks to Store, Retrieve, and Adapt Knowledge

- A TD-FALCON network is used for ***behavior modeling***:
  1. **Running around**: Explores randomly in the neighborhood;
  2. **Collecting item**: Runs to particular locations to collect useful items;
  3. **Escaping from battle**: Flees and collects nearby health boosts;
  4. **Engaging fire**: The bot tries to kill the opponent and avoids being hit;

  To learn appropriate behaviour mode in different situations.

- A FALCON network is used for ***weapon selection***:

  To learn appropriate fighting strategies for different situations/opponents

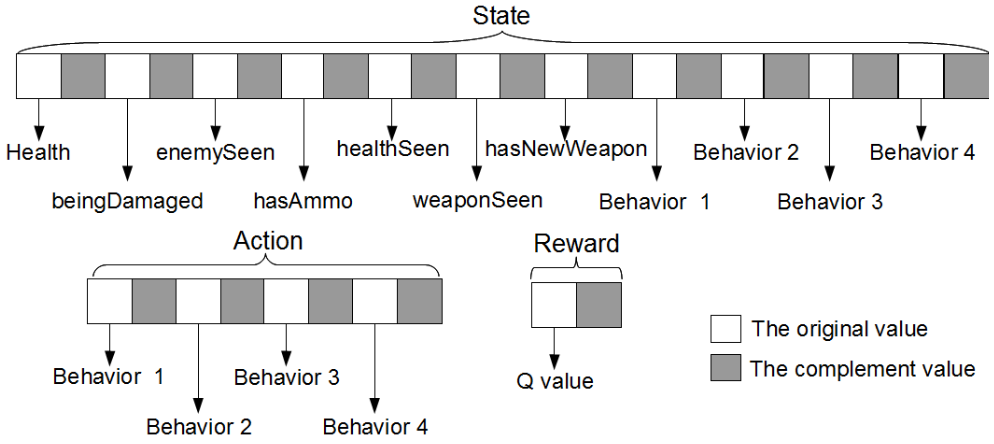

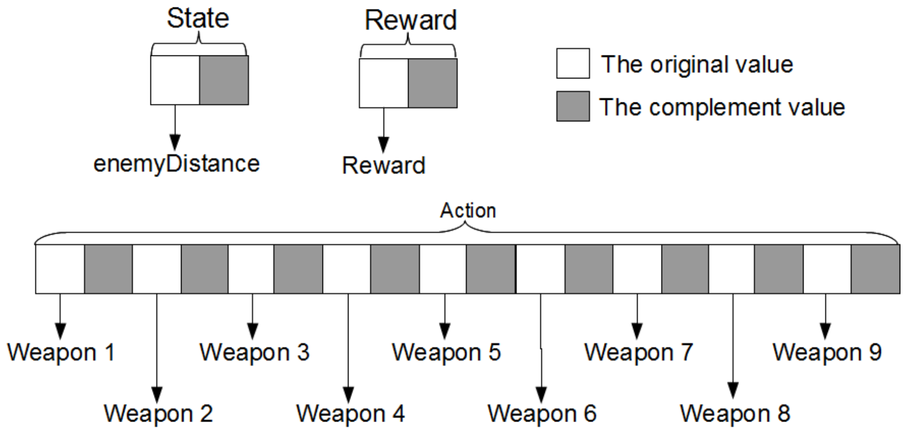

Highlight of Experimental Results

- Left plot shows how FALCONBot gradually learns from scratch (Config. 5)

  - 45 game trials: game score difference in the 1st game trial is below 0

- Right plot shows how FALCONBot quickly adapts with transferred knowledge (Config. 6)

  - 15 game trials: game score difference in the 1st game trial is above 7

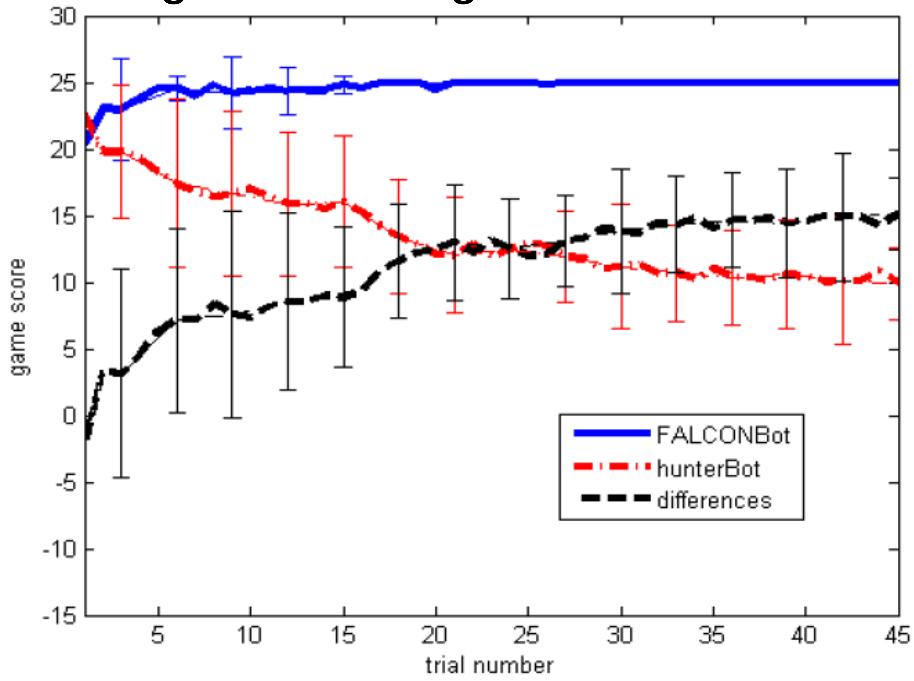

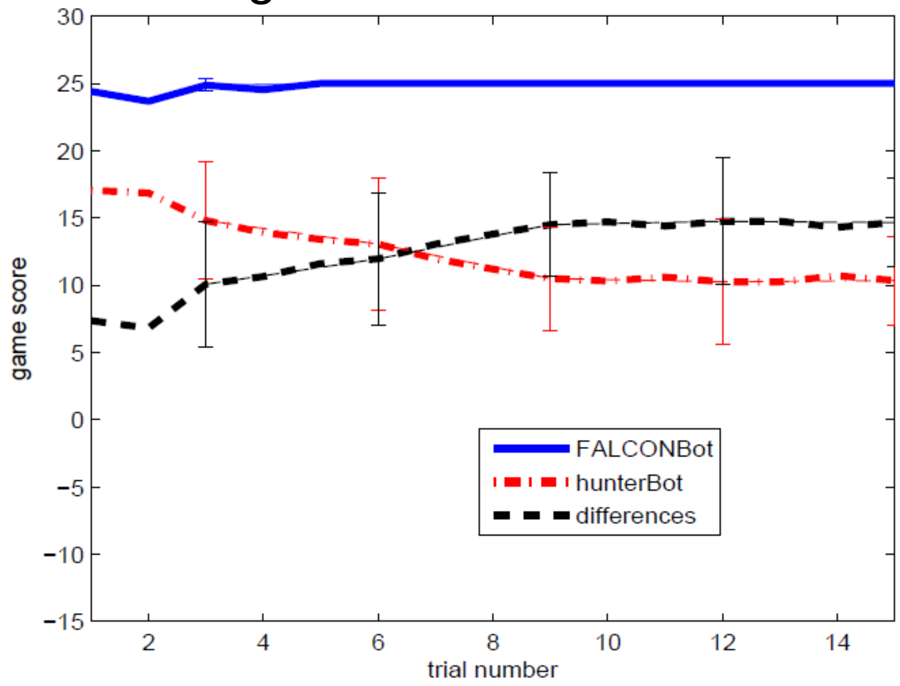

Autobiographical Memory

- *memory of autobiographical events,*

consisting of times, places, associated emotions,

and other contextual knowledge, that can be explicitly stated.

- *It is a special part of human mind, as it relates to who we are and things we remember during the course of our lifetime.*

In a way, autobiographical memory allows us to "travel mentally back in time" and re-experience specific events from our personal past (Tulving & Markowitsch, 1998).

Research Goal

To investigate the neural dynamics of autobiographical memory so as to develop computational models,

which are able to simulate its key functions, notably the storage and retrieval of real-life autobiographical experience in real time.

Potential Applications

- *To build autonomous systems with memory*
- *Human-like and user friendly systems*
- *Possible explanation for memory disorders*
- *Assist people with memory deficiency in retaining and refreshing their precious memory*

Neural Basis of Memory

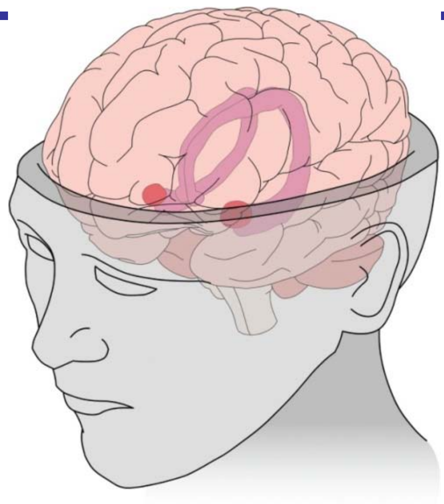

Source: Drawn by Shawn Fu.

- MTL (perirhinal cortex) include two hippocampi and olfactory area.
- MTL interacts with the higher level visual area: inferior temporal lobe (IT)
- Close to MTL is auditory cortex and amygdala responsible for emotions

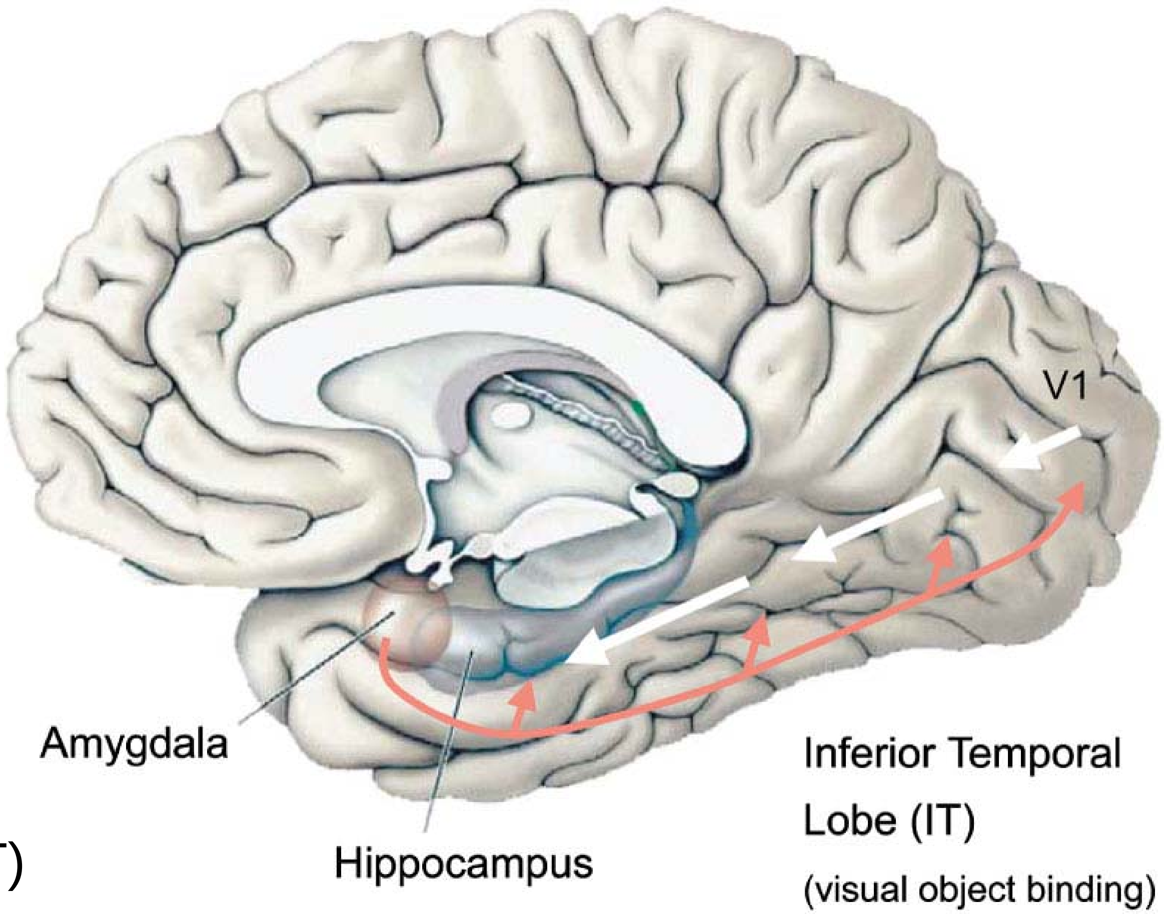

Source: Vuilleumier, 2005.

Neural Basis of Memory

- MTL (perirhinal cortex) integrates multiple brain inputs.
- It is a “hub of hubs”.
- Hippocampus combines cognitive information from neocortex with emotional information from limbic areas and binds this information into memory that codes consciously experienced events.

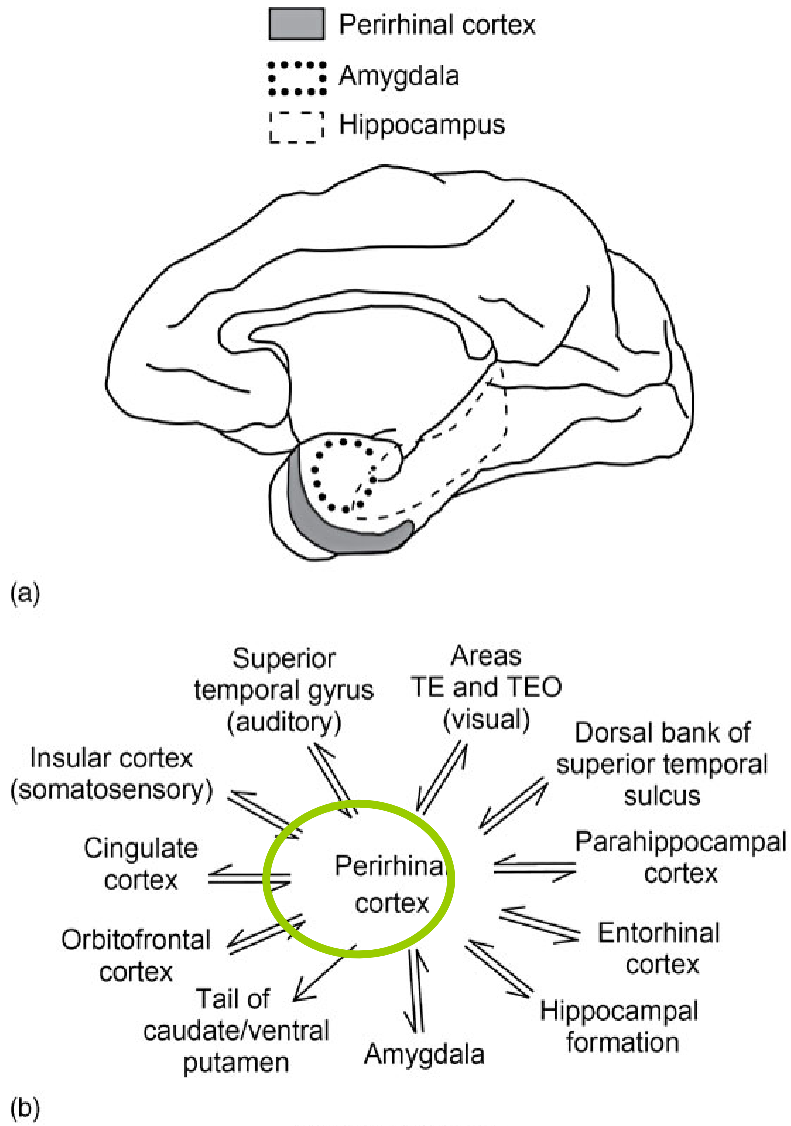

Source: Murray and Richmond, 2001.

Mechanism for Memory Formation

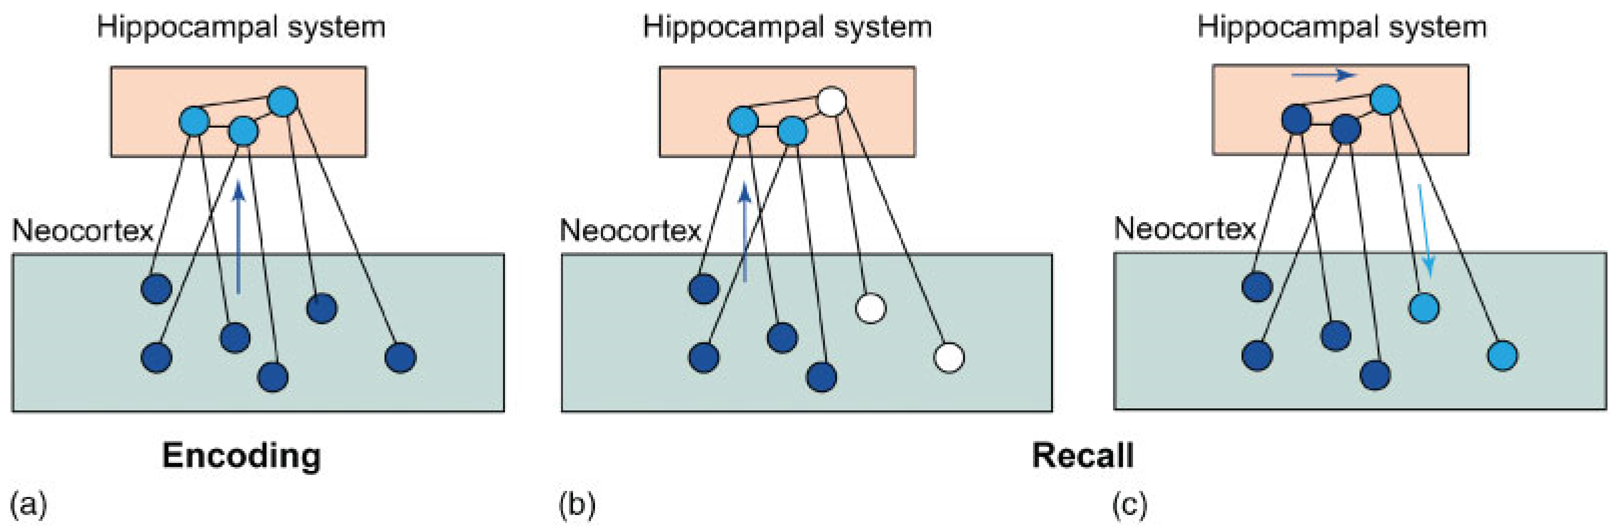

Source: Gluck et al., 2003.

Formation and recall of memories involves interplay between temporal lobe and neocortex:

- During the encoding or learning process (a), information from cortex is transferred to the hippocampal system.
- During recall, (b) a part of an neocortical event serves to evoke an overlapping pattern of neural activation in the MTL (the blue dots).
- (c) The hippocampal system responds by activating neocortical regions that provide the experience of recall of missing parts of the original event.

Fusion ART
(Tan, Carpenter & Grossberg, 2007)

- A generalized multi-channel ART architecture for pattern fusion and binding
- Capable of supporting many distinct learning paradigms as well as symbolic knowledge extraction and integration

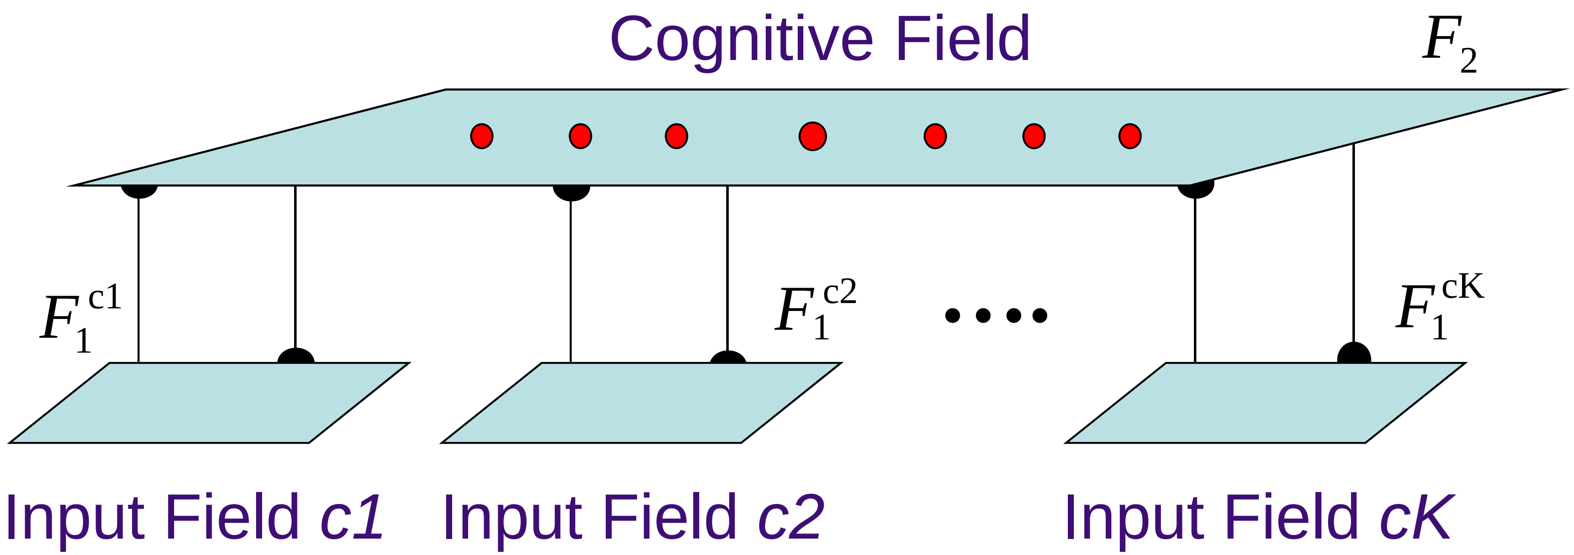

Fusion ART for Memory Encoding

- learn a snapshot of an event as distributed patterns across multiple channels
- growing nodes in response to new experience
- Control the growth of the network and the level of generalization with adjustable vigilance parameters

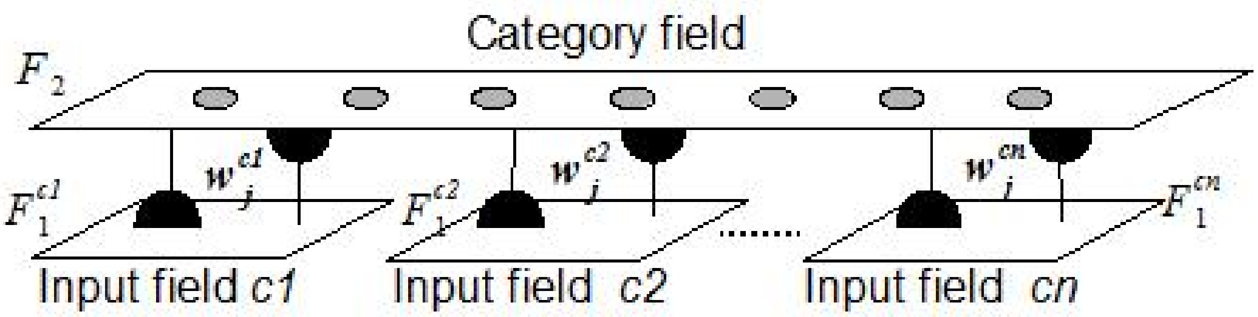

Resonance Search

$$T_j^c = \sum_{k=1}^{n} \gamma^{ck} \frac{|x^{ck} \wedge w_j^{ck}|}{\alpha^{ck} + |w_j^{ck}|}$$

- Code Activation/Competition: $T_J^c = \max\{T_J^c : \forall F_2^c \text{ node } j\}$
- Template Matching: $m_J^{ck} = \frac{|x^{ck} \wedge w_J^{ck}|}{|x^{ck}|} \geq \rho^{ck}$
- Readout: $x^{ck(new)} = w_J^k$
- Template Learning:

$$w_J^{ck(new)} = (1-\beta^{ck})w_J^{ck(old)} + \beta^{ck}(x^{ck} \wedge w_J^{ck(old)})$$

Fusion ART for Memory Binding

Step 4: Template learning

$$w_j^{ck(new)} = (1-\beta^a) w_j^{ck(old)} + \beta^a (x^{ck} \wedge w_j^{ck(old)})$$

Step 3: Top down priming

$$m_j^{c1} = \frac{|x_j^{c1} \wedge W_j^{c1}|}{|x_j^{c1}|} \qquad m_j^{c2} = \frac{|x_j^{c2} \wedge W_j^{c2}|}{|x_j^{c2}|} \qquad m_j^{c3} = \frac{|x_j^{c3} \wedge W_j^{c3}|}{|x_j^{c3}|}$$

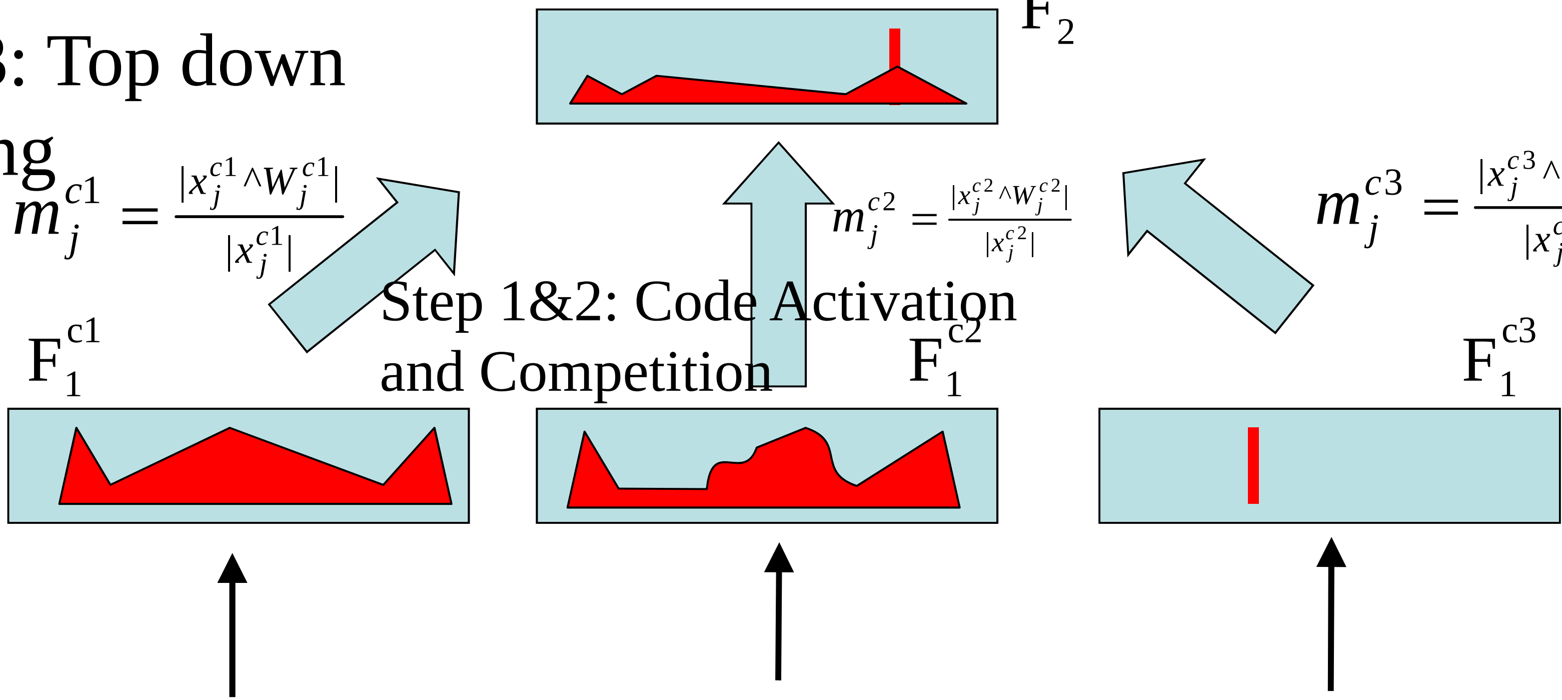

Context

Visual

Auditory

When, Where, Who, What, …

Fusion ART for Memory Recall

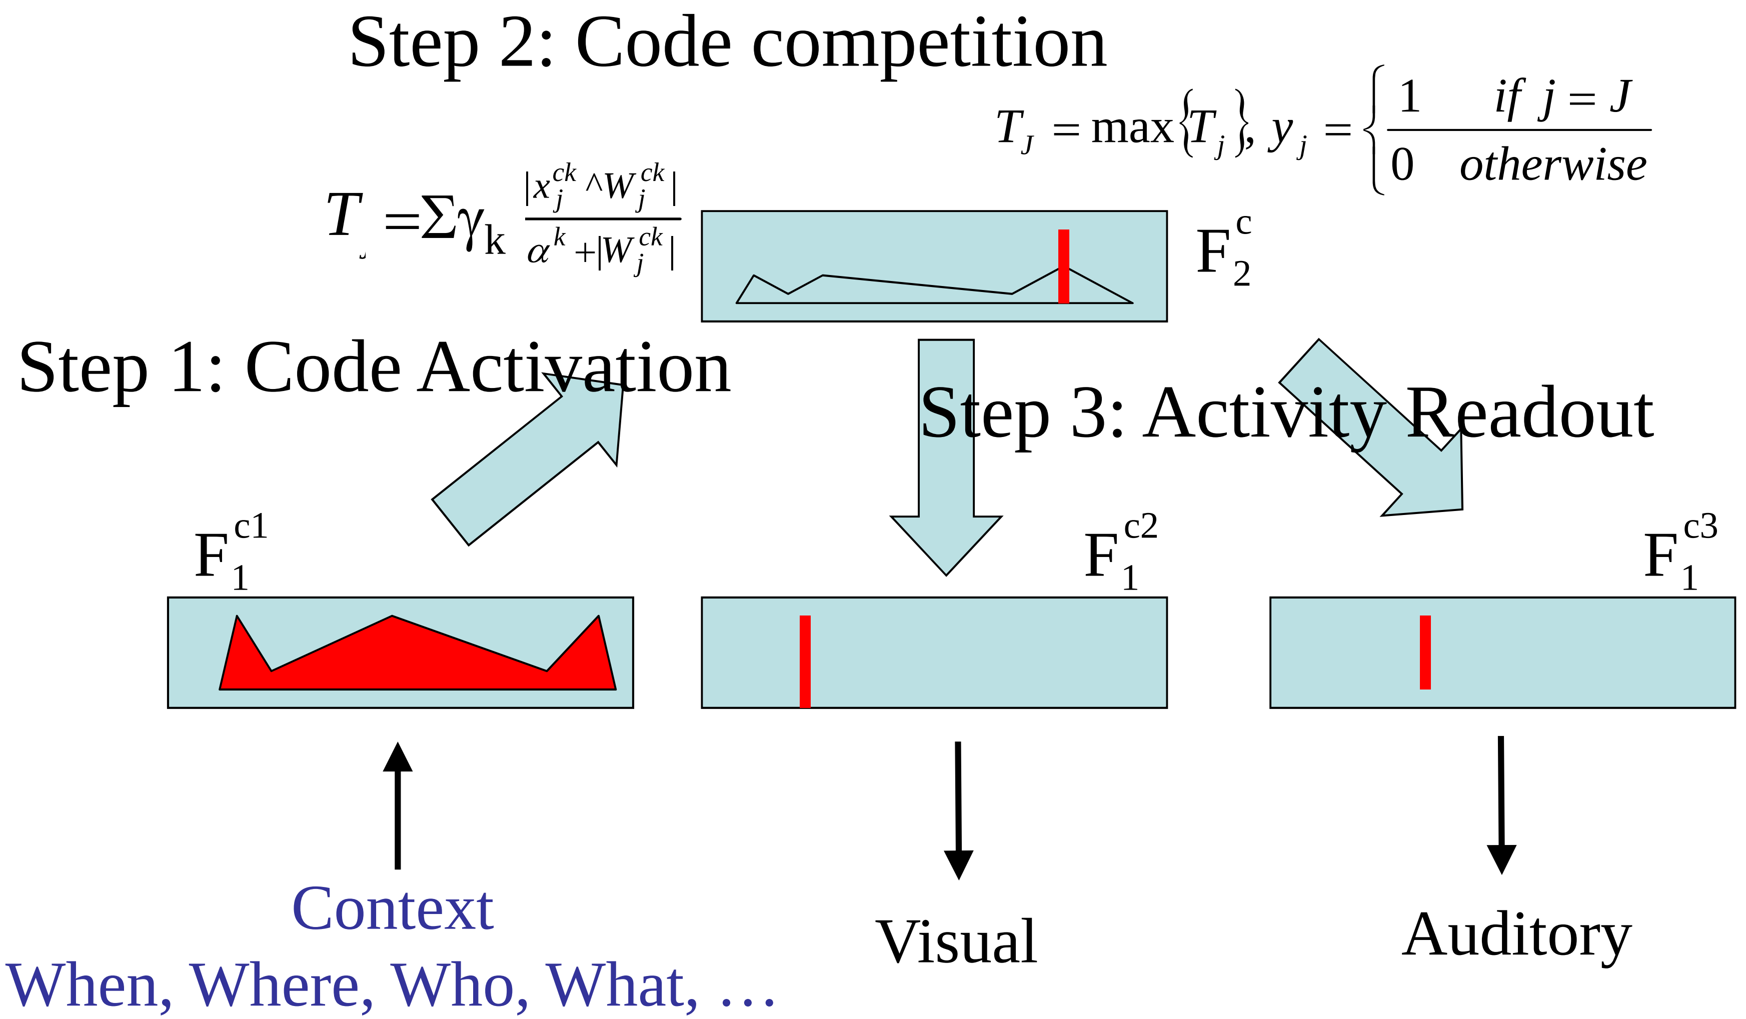

Modeling Spatiotemporal Memory

- Two basic elements: ***event*** (a snapshot of experience) and ***episode*** (a temporal sequence of events that one experiences)
- The proposed episodic memory model, namely ***EM-ART*** , is built by hierarchically joining two fusion ART network
- F2 learns individual events; F3 learns episodes (sequences)

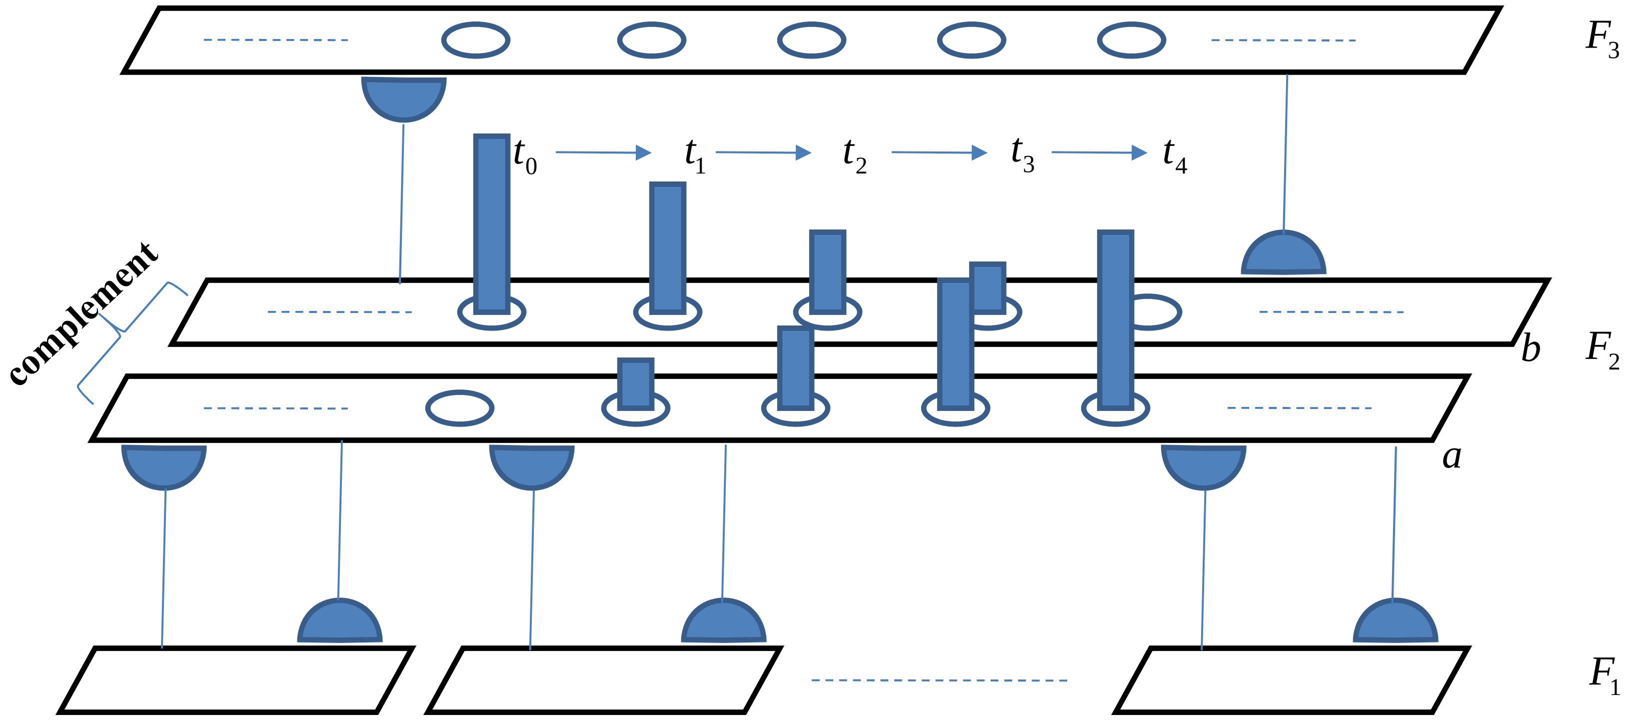

Process for Event Encoding

Given an incoming pattern of event in F_1

Activate and select a node (winner-take-all) in F_2

While the node is not in *resonant* condition (template matching)

or has been selected previously

Do reset the current node activation and choose another node in F_2

If no F_2 node can be found

then recruit an uncommitted node and learn it as a novel event

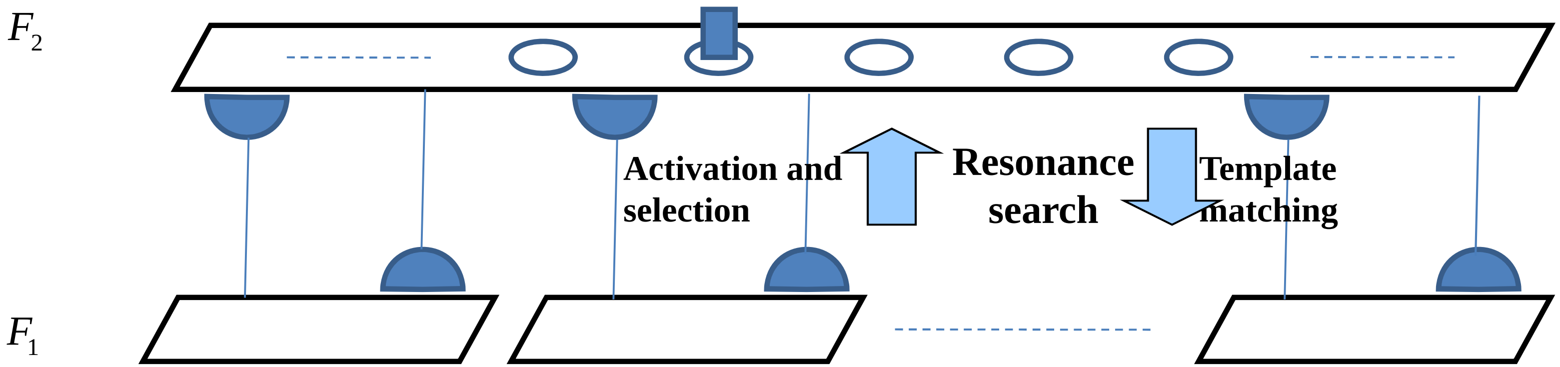

Process for Episode Encoding

Representing/Storing Sequences of Patterns

- Gradient Encoding: a real activation value indicates a time point or a position in an ordered sequence.
- Allows sequences to be learned and matched by stored patterns using the same mechanism of fusion ART neural network

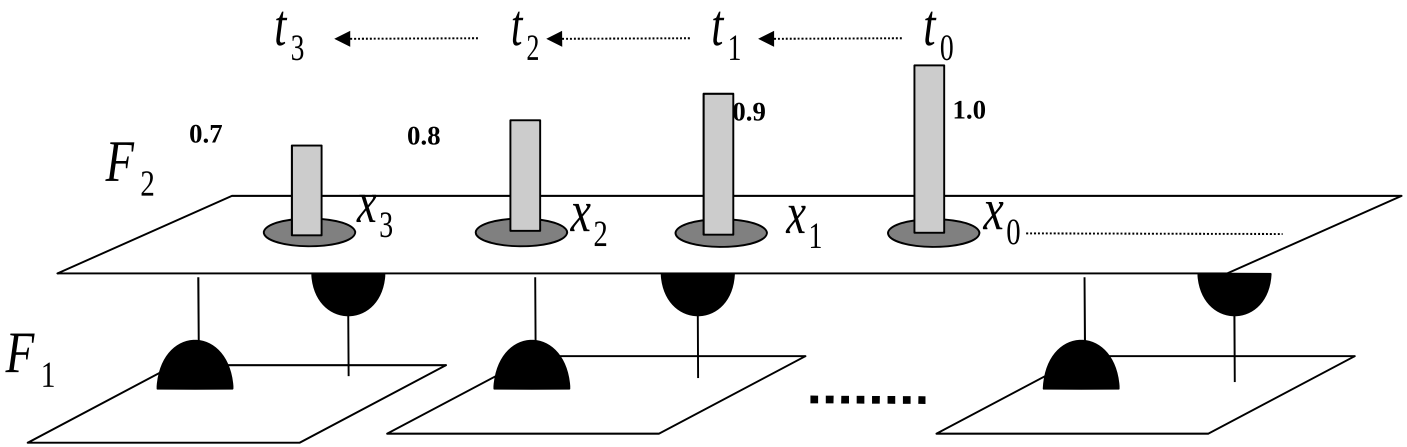

Process for Episode Encoding

For each event in an episode *S*

select a node in F_2 based on incoming event in F_1 via resonance search

set activation value of the selected node T_t in $F_2^{\,a}$ to 1

decay values of previously activated F_2 nodes at time *t-i*

$$\tau_{\text{t-i}}^{(\text{new})} = \tau_{\text{t-i}}^{(\text{old})} * (1-\nu) \ \text{ where } \ 0 \le \nu \le 1 \ \text{ and } \ 0 \le \tau_{\text{t-i}} \le 1$$

Set complementary nodes in F_2 (b) to 1-T_t

At the end of *S* activate, select and learn a node in F_3 based on the pattern of episode formed in F_2 through resonance search

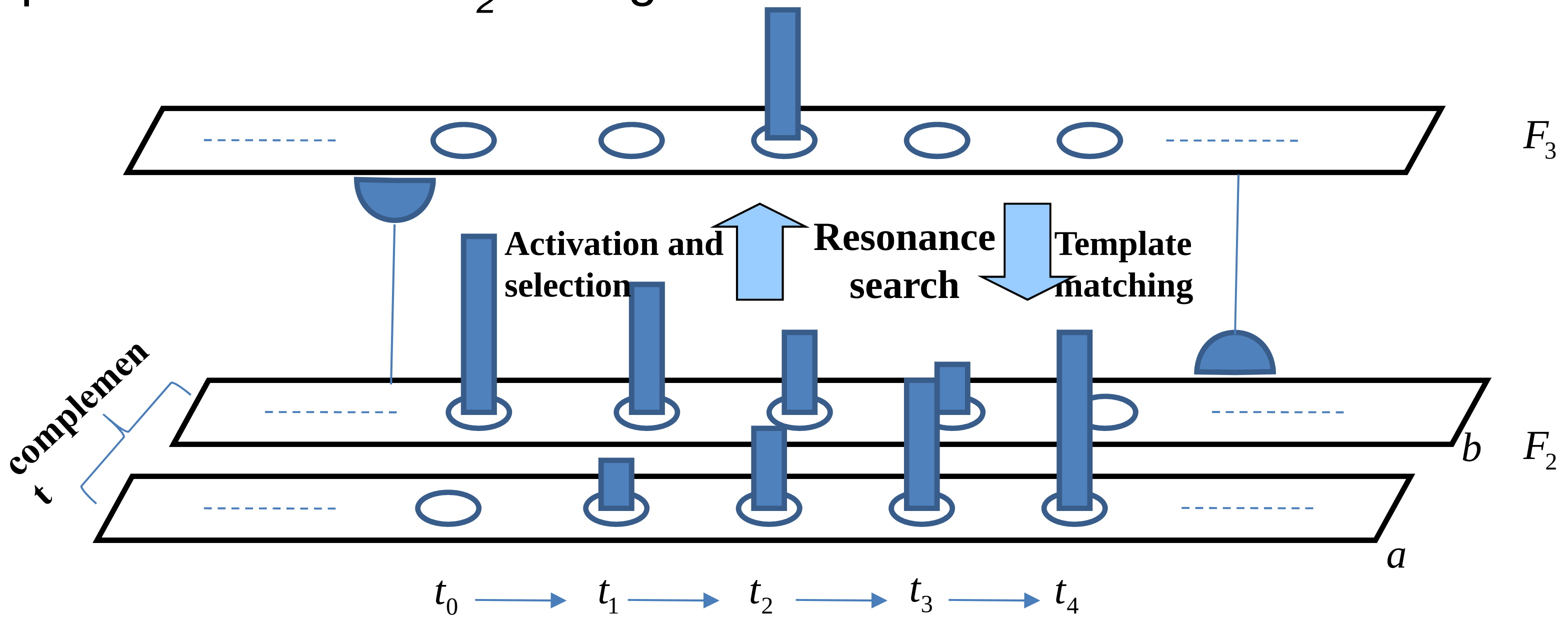

Experiments:
Episodic Memory of NPC

Unreal Tournament:

Episode Encoding

- To learn the experience of the agent from its 100 battles (i.e. episodes) played .
- There are in total 7,735 events in the data set.
- The number of events within an episode varies from 7 to over 250.

Event

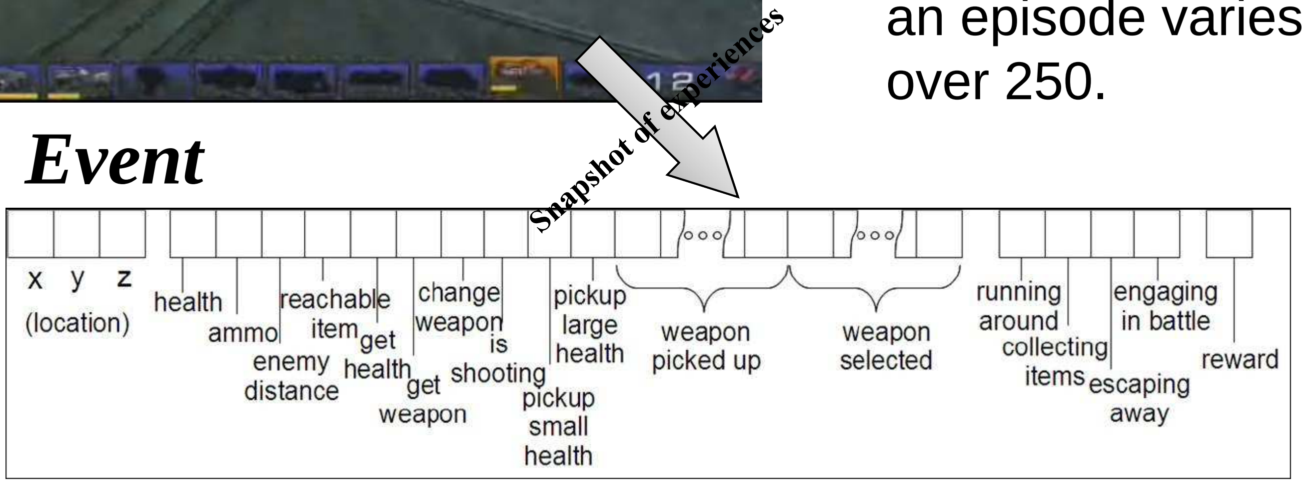

Event Encoding

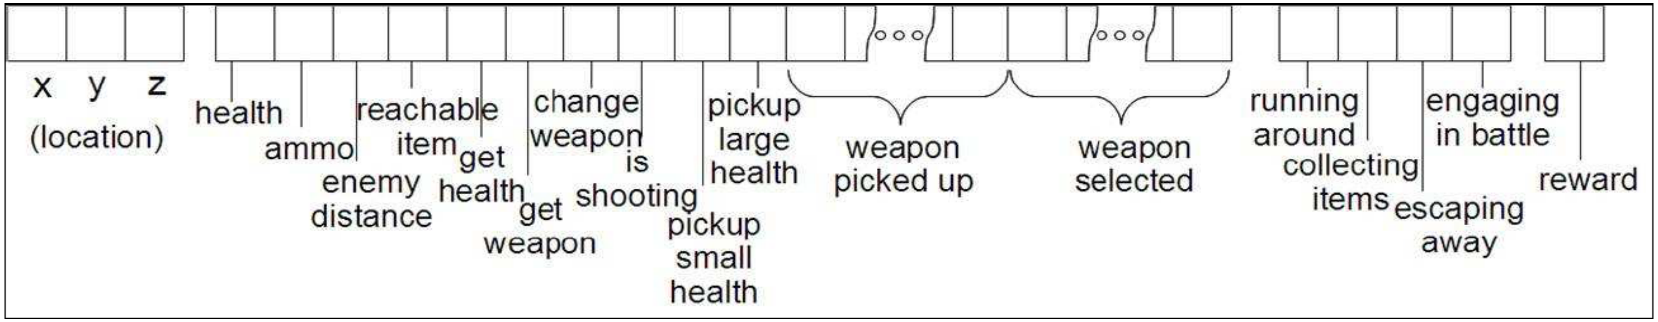

Tests are conducted to evaluate the accuracy of retrieval, given variations of cues:

- Partial length episodes of different length starting from beginning
  (Episode: A B C D E; Cue: A B)
- Partial length episodes of different length starting from end in the episode
  (Episode: A B C D E; Cue: D E)
- Partial length episodes of different length starting from different locations
  (Episode: A B C D E; Cue: C D)
- Full length episodes with event-level errors
  (Episode: A B C D E; Cue: A F C D E)
- Full length episodes with sequence-level errors
  (Episode: A B C D E; Cue: A B D C E)

Experimental Results

- EM-ART can correctly retrieve all episodes across all types of partial cues
- EM-ART can tolerate as high as 20% errors in the sequence representation
- but shows a gradual degradation of the performance as noises are added into the event representation of the cues

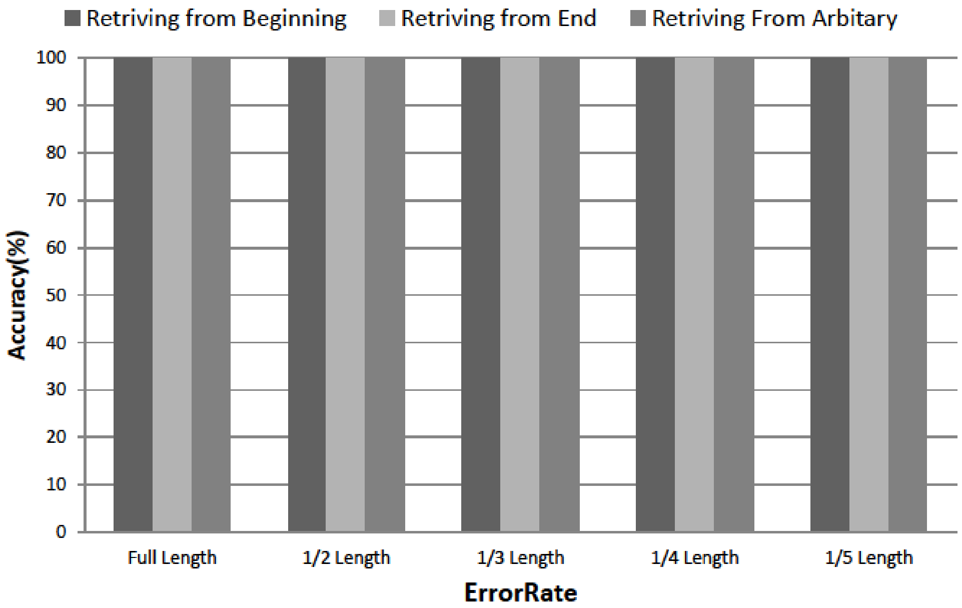

Retrieval Accuracies With Partial Cues

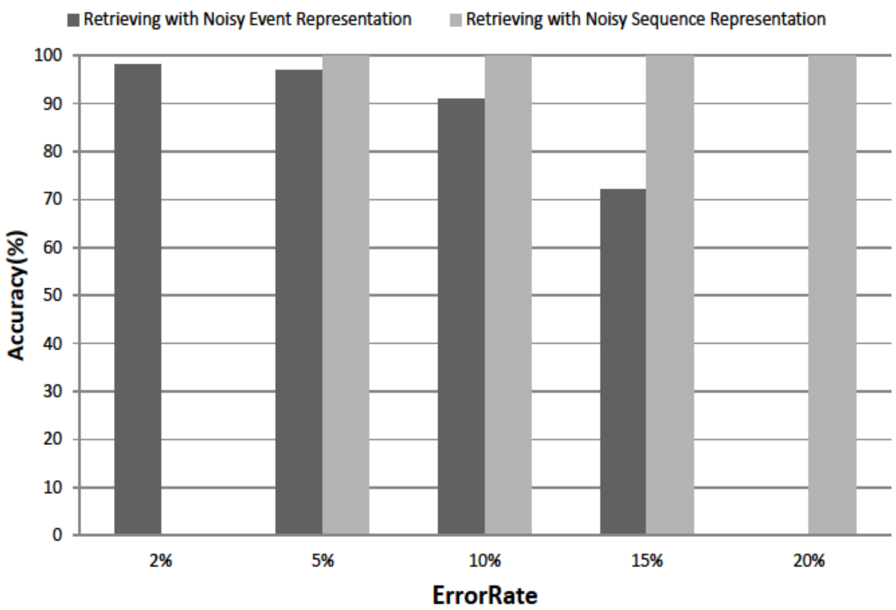

Retrieving with Noisy Cues

Comparison with LTM Model
(Starzyk & He, 2009)

- The two models in comparison show similar performance with partial cues and noisy cues on event representation
- EM-ART is better in tolerating errors in the sequence

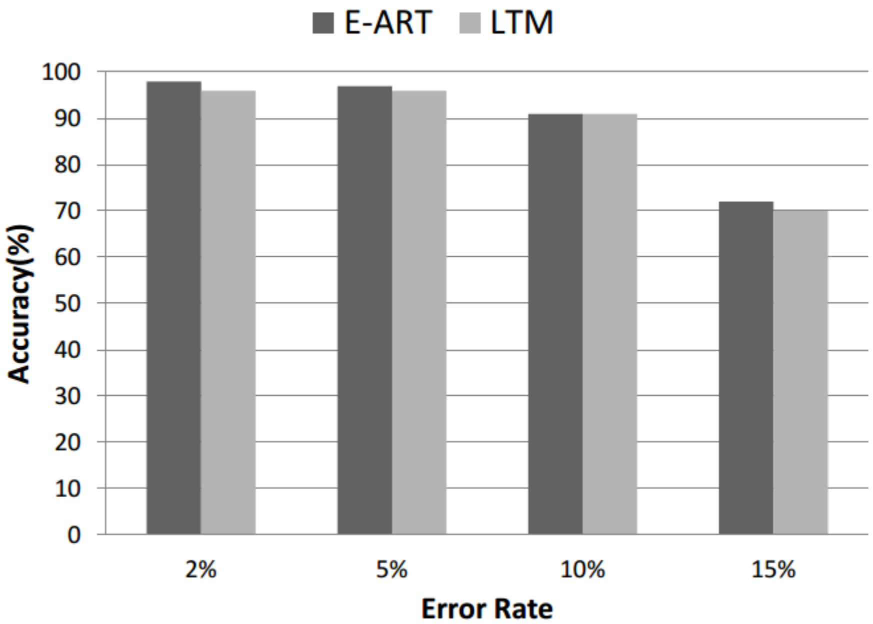

Noises in events

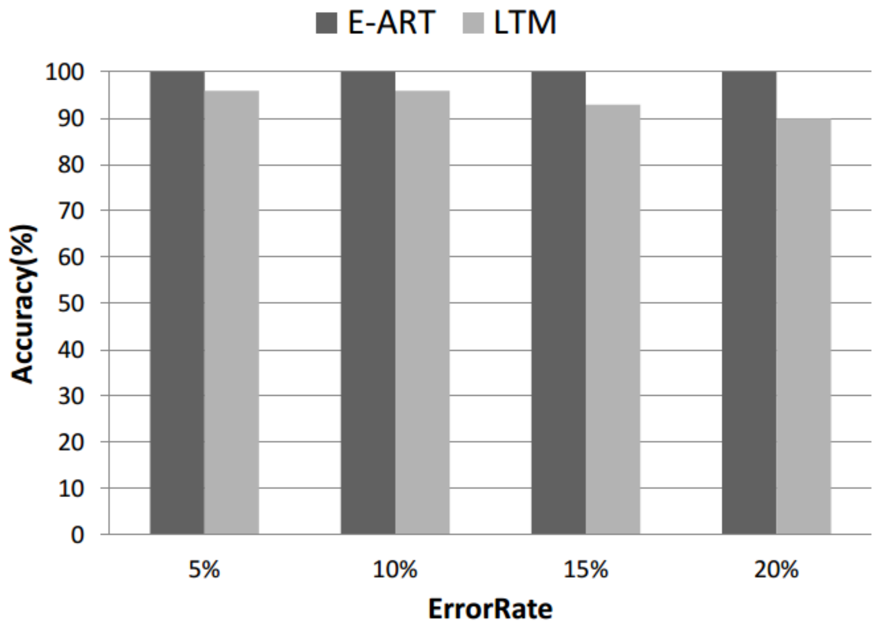

Noises in event sequences

Summary

- Our aim is to develop *biologically-inspired neural models for autobiographical memory*
- Fusion ART: a universal model for pattern binding and memory encoding, similar to that of hippocampus
- We have developed *EM-ART, a neural model for encoding and retrieval of multi-channel spatial-temporal patterns, which was evaluated on various benchmark tasks, embedded into Non-player characters in games, and implemented as MyLife Simulator*
- Working towards modelling of spatial representation and investigation into the role of emotion into memory formation/retrieval

Modeling Self-Awareness

Towards Self-Aware Sociable Humanoid Robots for Long-term Companionship

B. Subagdja and A.-H. Tan. Towards a Brain Inspired Model of Self-Awareness for Sociable Agents. In *Proceedings of The 31st AAAI Conference on Artificial Intelligence (AAAI-17)*, San Francisco, CA, USA, 4-9 February, 2017, pp 4452-4458

Background/Motivations

To develop household companion robot(s) who

- *live together and converse, give advice, plays with the occupants,*
- *learn from experiences and interaction with eople*

Requirement of Self-Aware Robots

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

- Understand its own self, the people surrounding, and how to better interact with them (self-aware and social-aware)
- Learn incrementally from experiences and can self-express what it has known from time to time
- Adapt to changes in the social environment and maintaining its role as a companion
- "..must be able to disobey in order to obey" http://www.todayonline.com/commentary/robots-must-be-able-disobey-order-obey , http://www.scientificamerican.com/article/the-conversation-why-robots-need-to-be-able-to-say-no/

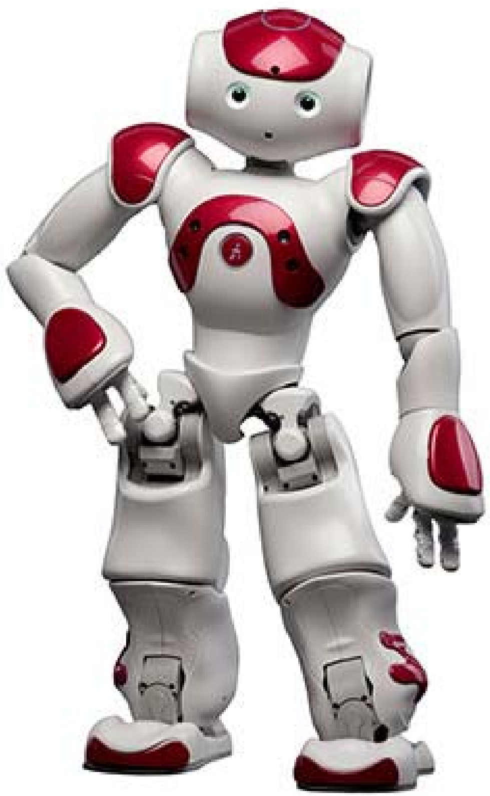

Challenge: Modeling Self of Robots

- How to make the robot knows and understands about itself? (self-awareness)
  - *The robot knows why, what, who, and how about itself*
  - *The robot can reason or think about the context related to itself, like awareness of its own experiences in the past, at the current situation, and possibilities in the future, including the possibilities of what other people (or agents) would think or feel about itself (the robot's)*
  - *The robot can adapt and/or regulate itself to maintain its functionality and integrity to be a companion robot and to better interact with others*

Modeling Self-Awareness
Aspects of Self w.r.t. levels of Self-Awareness

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

Level 1
differentiation

- **Identity**
  - Conventional identification: attributes identifying the self like name, age, place of origin, birthplace, citizenship etc.
  - Personality: Stereotypical pattern of behavior that characterizes the individual self

Level 2
Situation

- **Physical**
  - Embodiment: representation of the individual's body as whole and parts, their states (e.g posture, position, angle, location), and sensory information
  - Capability: actions that can be performed at a moment and their side effects

Level 3
Identification

- **Mental**
  - Goals: current goals (and sub-goals) the individual wants to attain
  - Intentions: the current state of goal attainments and performed actions of the individual
  - Thoughts: what the agent currently think about or pay attention to
  - Emotions/appraisal: the current state of mood, emotion, and feelings

Level 4
Permanence

- **Experiential**
  - Autobiographical Memory: past experiences
  - Learnt concepts and knowledge: knowledge and concepts (semantic) learnt from experiences

Level 5
Self-Consciousness

- **Social**
  - known people/agents and their characteristics
  - one's relations and appraisal about the other
  - the other's relations and appraisal about one's self

Modeling Self-Awareness as Memory of Self in Five Dimensions

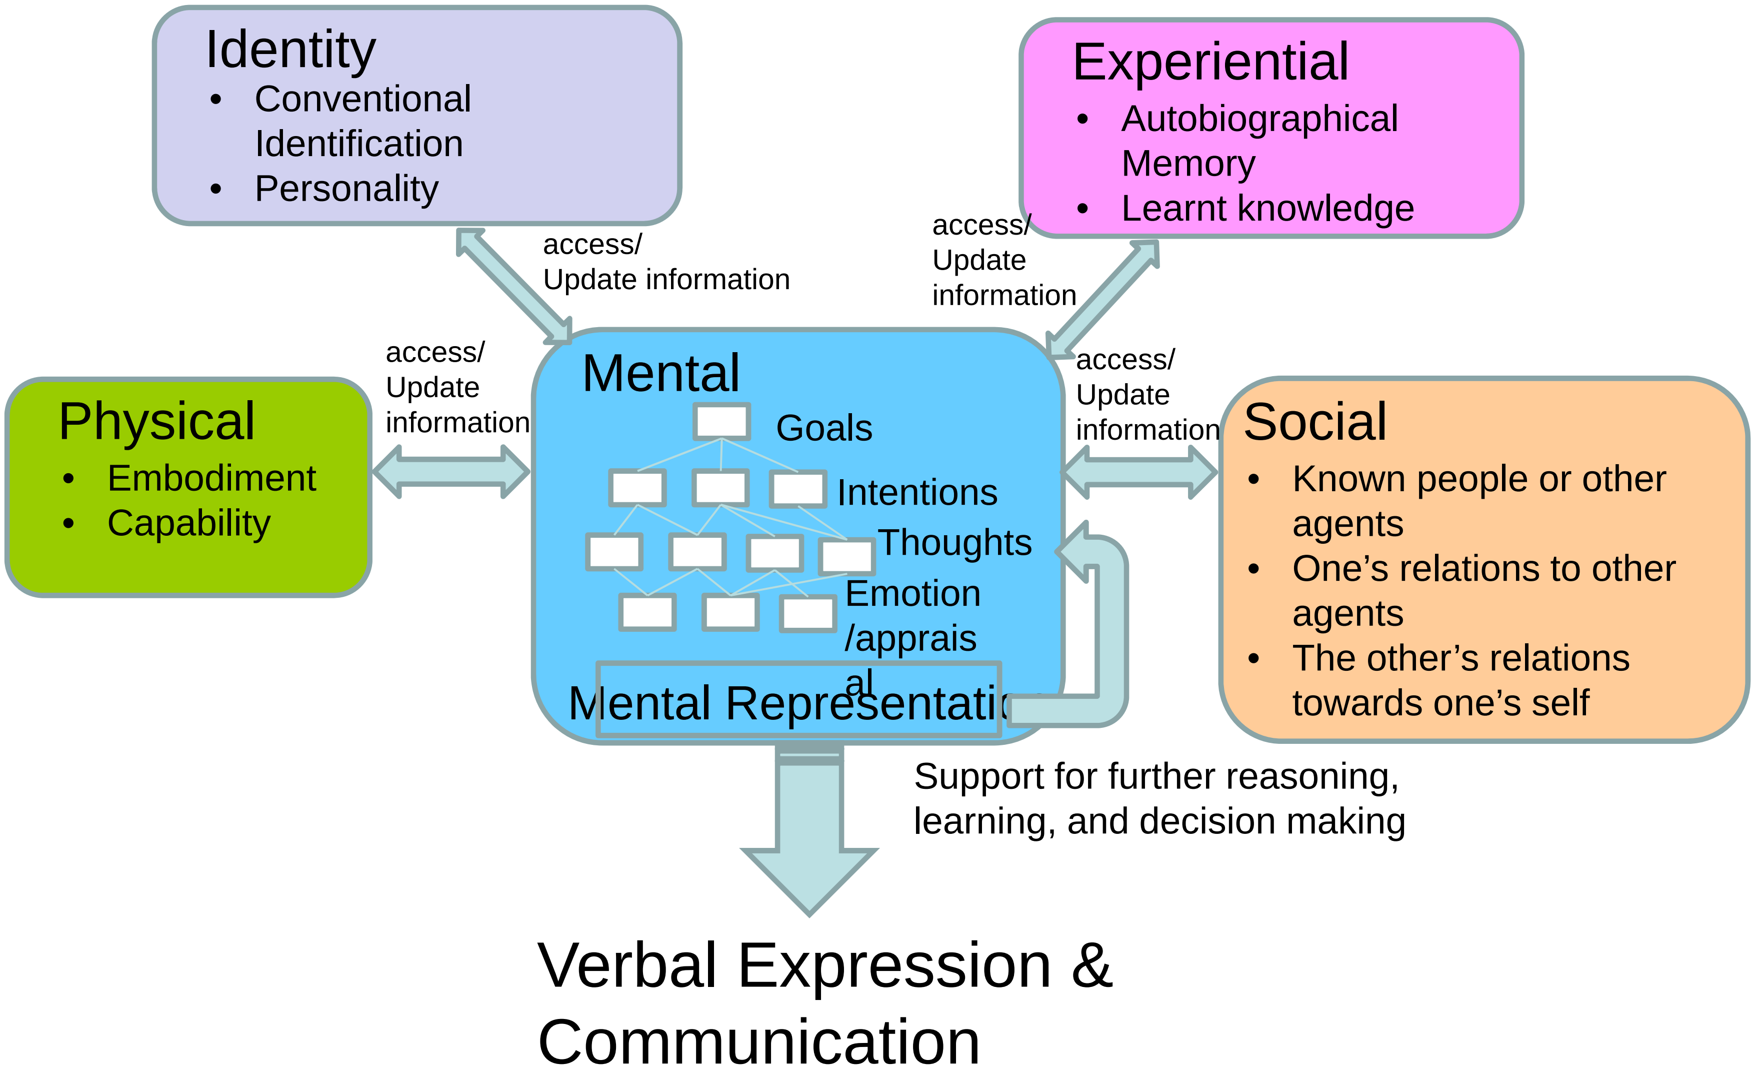

Neural Model of Self-Awareness (ARTSELF)

- Using fusion ART as building block
- Integrating aspects of individual
  - **Identity** (e.g name, Id, personality),
  - **Physical** (e.g embodiment, environment),
  - **Mental** (e.g belief, goal, feeling, imagination),
  - **Experiential** (e.g memory, learnt knowledge),
  - **Social** (e.g others, social group, relationships)

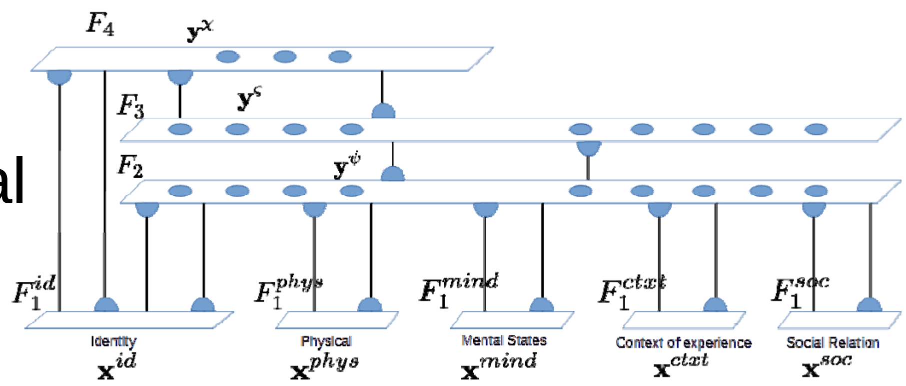

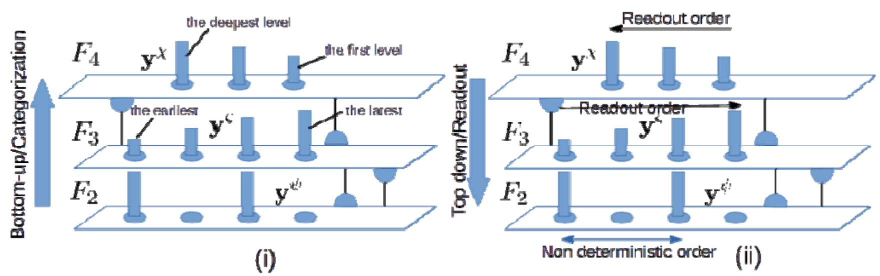

- Hierarchical layering of activations supports transient representation of multiple selves and perspectives
  - e.g. 1st person (subjective), 3rd person (objective)

Proposal of a Fusion Framework for Self-Awareness

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

- *ARTSELF serving as a hub, interacting with multiple memory systems*

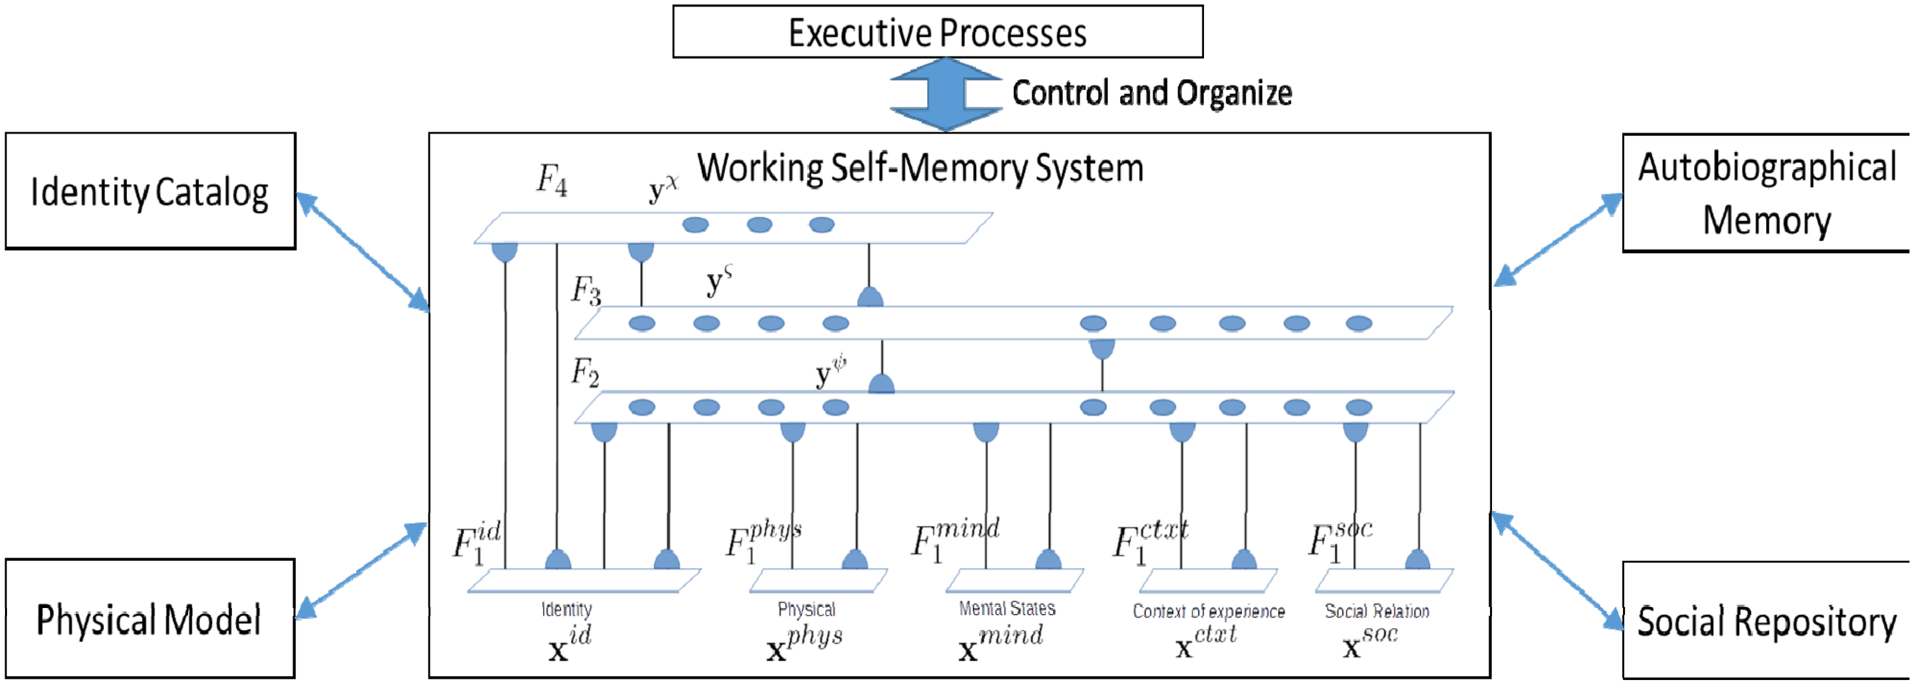

Embodying Self-Awareness in Humanoid Robot

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

- Implementation
  - Using NAO humanoid robot as the platform for ARTSELF
  - The robot can roam and explore the environment, recognize people, and have conversation with them
  - The robot can remember experiences and talk about them in a conversation

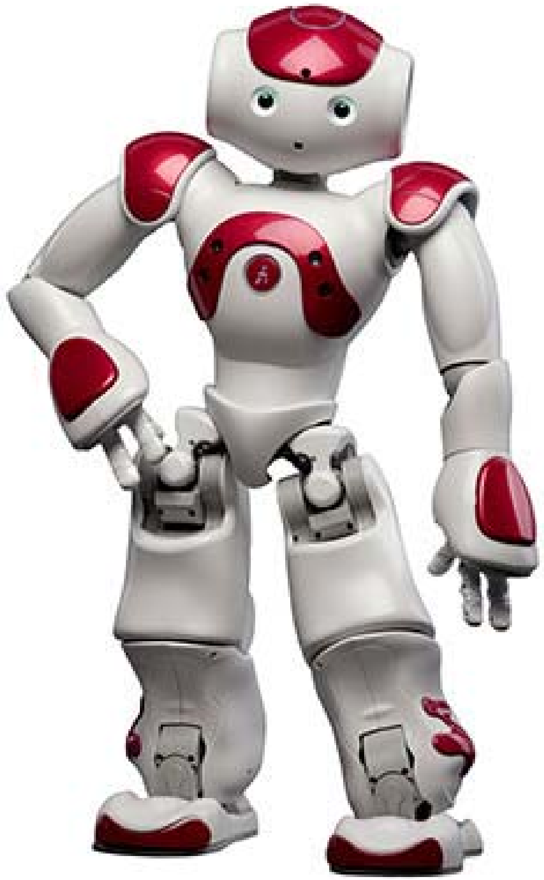

Illustrating Self-Awareness in Humanoid Robot

Example: the mind of a NAO robot during self-reflection and eye-contact

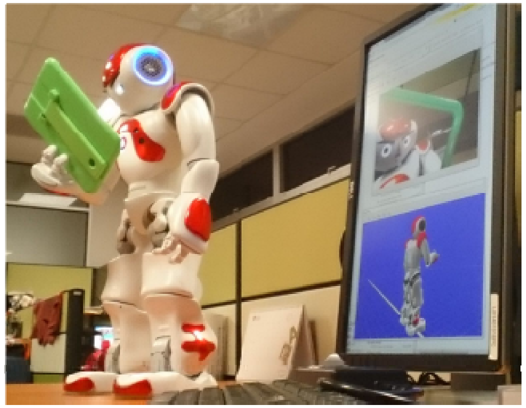

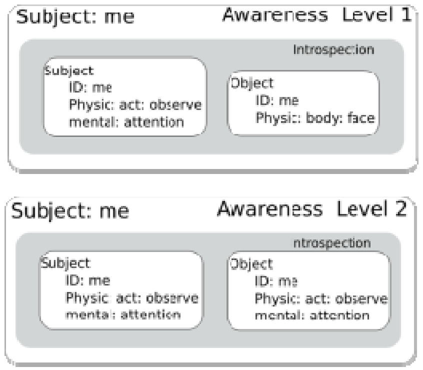

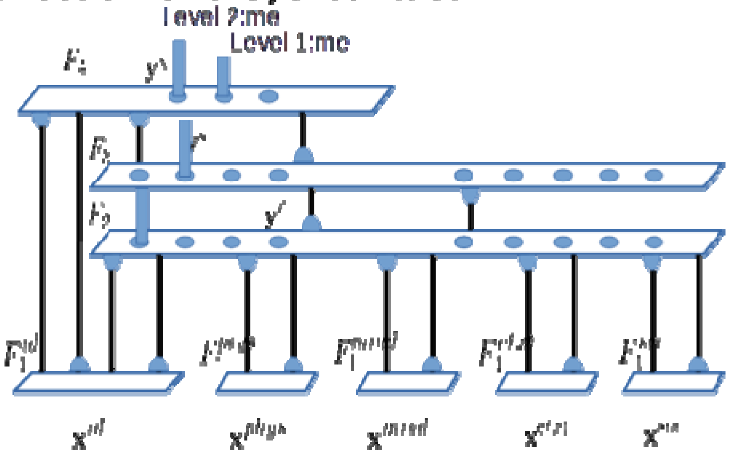

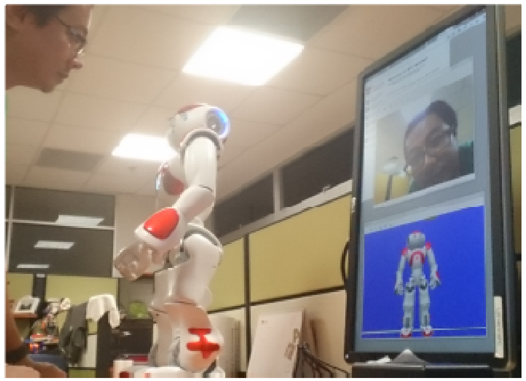

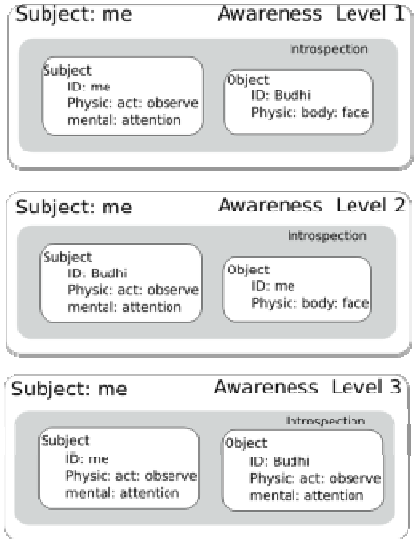

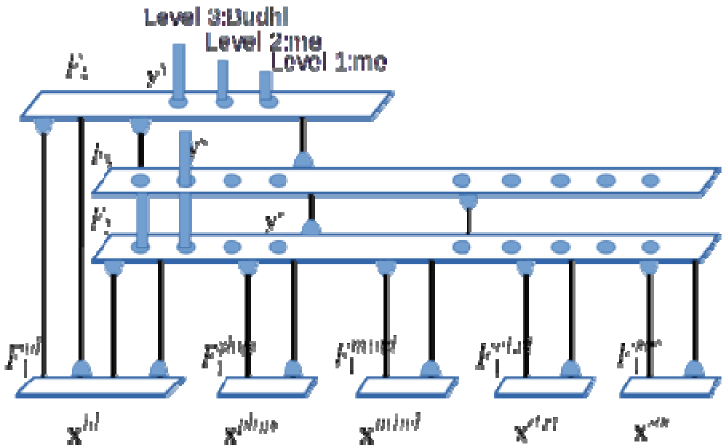

Illustrating Self-Awareness in Humanoid Robot

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

Example: Alignment and grounding in human-robot conversation

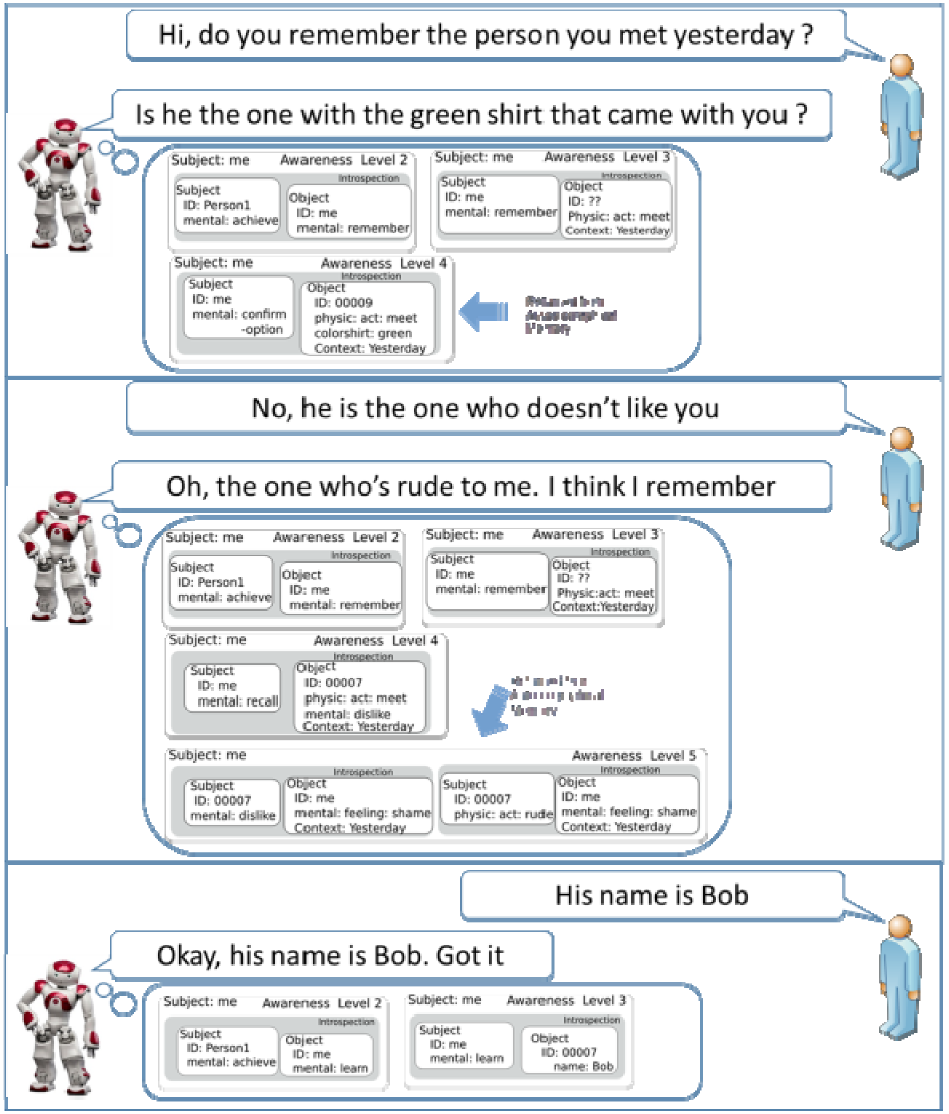

Challenges Ahead

- How to automate the formation of memories of self?
- How to model the interaction of multidimensional aspects of self?
- How to scale up the model?
- How to integrate with other knowledge?
- How to measure degree of self-awareness?
- Development of self-aware cognitive systems

References

- Ah-Hwee Tan, Gail A. Carpenter and Stephen Grossberg. Intelligence Through Interaction: Towards A Unified Theory for Learning. *ISNN 2006, LNCS 4491, Part I*, pp. 1098-1107.
- Ah-Hwee Tan. Direct Code Access in Self-Organizing Neural Architectures for Reinforcement Learning . In proceedings, International Joint Conference on Artificial Intelligence (IJCAI'07), pp. 1071-1076, Hyderabad, India, 6-12 January 2007.
- Ah-Hwee Tan, Ning Lu and Dan Xiao. Integrating Temporal Difference Methods and Self-Organizing Neural Networks for Reinforcement Learning with Delayed Evaluative Feedback. IEEE Transactions on Neural Networks, Vol. 9, No. 2 (February 2008), 230-244.
- Teck-Hou Teng, Ah-Hwee Tan, and Jacek Zurada, Self-Organizing Neural Networks Integrating Domain Knowledge and Reinforcement Learning. *IEEE Transactions on Neural Networks and Learning Systems*, Vol. 26, No. 5 (2015), 889-902.
- Di Wang, Ah-Hwee Tan. Creating Autonomous Adaptive Agents in a Real-Time First-Person Shooter Computer Game. *IEEE Transactions on Computational Intelligence and AI in Games*, Vol. 7, No. 2 (2015), 123-138.
- Wenwen Wang, Budhitama Subagdja, Ah-Hwee Tan and Janusz A. Starzyk. Neural Modeling of Episodic Memory: Encoding, Retrieval, and Forgetting. *IEEE Transactions on Neural Networks and Learning Systems*, Vol. 23, No. 10 (Oct 2012), 1574–1586.
- Di Wang, Ah-Hwee Tan and Chunyan Miao. Modeling Autobiographical Memory in Human-Like Autonomous Agents. In Proceedings, 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), pp845-853, Singapore, May 9-13, 2016.
- Budhitama Subagdja and Ah-Hwee Tan. Towards a Brain Inspired Model of Self-Awareness for Sociable Agents. In *Proceedings of The 31st AAAI Conference on Artificial Intelligence (AAAI-17)*, San Francisco, CA, USA, 2017, pp 4452-4458

Acknowledgement

NATIONAL RESEARCH FOUNDATION
Interactive Digital Media Programme

A*STAR I²R

DSO NATIONAL LABORATORIES

UBC NANYANG TECHNOLOGICAL UNIVERSITY Lily

JOINT UBC-NTU RESEARCH CENTRE

Singapore Technologies Engineering

IntelliSys