MEMORY IN EMBODIED AGENTS

TECHNICAL FIELD

[0001] Embodiments described herein relate to the field of artificial intelligence, and systems and methods for implementing and using Memory in Embodied Agents. More particularly, but not exclusively, embodiments described herein relate to unsupervised learning.

BACKGROUND ART

[0002] A goal of Artificial Intelligence (AI) is to build computer systems with similar capabilities to humans, including human-like learning and memory. Most contemporary machine learning techniques rely on "offline" learning, wherein AI systems are provided with prepared and cleaned data to learn on, limited to a specific domain. An outstanding challenge in the prior art remains in creating AI systems which experience objects and events in the world in a human-like way and learn from embodied interaction. By virtue of their embodiment and sensorimotor feedback loops with their environment, such AI agents may influence and guide their own learning. Such agents would make sense of streams of multimodal data from the world and retain information in a meaningful and useful way. A further outstanding challenge is to create a flexible AI Embodied Agent which can both learn from its own experience, as well as have its memories authored or altered by an external source (such as a human user). Hierarchical Temporal Memory (HTM) is an approach to replicating human memory which is based on a computational structure with multiple registers as analogues to cortical layers. HTM is configured to replicate patches of cerebral cortex. Nonetheless, HTM fails to provide Memory in Embodied Agents which allows Embodied Agents to learn and develop in real-time from sensorimotor experience.

SUMMARY

[0002a] In a first aspect, the present disclosure provides a method for training a SOM, including a plurality of Neurons, each Neuron associated with a weight vector, and a training record; including the steps of: receiving an input vector; determining if the input vector is "new"; if the input vector is not new: selecting a first Winning Neuron, favouring higher similarity between the Input Vector and the Winning Neuron, and modifying the weight vector of the first Winning Neuron towards the input vector; if the input vector is new: selecting a second Winning Neuron, favouring neurons with lower training records, and modifying the weight vector of the second Winning Neuron towards the input vector.

[0002b] In a second aspect, the present disclosure provides a method for targeted forgetting in a SOM including a plurality of Neurons, each Neuron associated with a weight vector, including the steps of: creating a Reset Mask comprising a plurality of Mask Values, each Mask Value associated with a Neuron of the SOM; applying a Reset Function to each Neuron of the SOM associated with a non-zero Mask

Value, wherein the Reset Function includes: a Forgetting Component, for resetting the Neuron's weights to an untrained state; and a Mask Component for modifying the output of the Reset Function as a function of the Forgetting Component and the Mask Value; modifying SOM Neuron weight vectors according to the output of the Reset Function.

BRIEF DESCRIPTION OF DRAWINGS

Figure 1: a schematic diagram of a CDZ architecture.

Figure 2: an ASOM.

Figure 3: eligibility signals for different modalities.

Figure 4: how an Eligibility Trace creates an Eligibility Window.

Figure 5: the phases of a learning event.

Figure 6: a user interface for setting the eligibility of different modalities.

Figure 7: a display of ASOM training.

Figure 8: a query viewer Input Field.

Figure 9: a user interface for specifying query patterns.

Figure 10: a display of LTM and STM.

Figure 11: a Working Memory System (WM System).

CLAIMS

1. A method for animating an Embodied Agent, the method including the steps of:

   receiving sensory input corresponding to a first representation of an Experience in a first modality;

   querying an Experience Memory Store to retrieve a second representation of the Experience in a second modality; and

   using the second representation in the second modality to animate the Embodied Agent.

2. The method of claim 1 wherein the Experience Memory Store stores representations of the Experiences in neural network weights.

3. The method of claim 1 wherein the Embodied Agent is animated using a neurobehavioural model.

4. The method of claim 1 wherein Experiences in the Experience Memory Store include Experiences previously experienced by the Embodied Agent in the course of live operation of the Embodied Agent.

5. The method of claim 1 wherein the plurality of modalities include one or more of the group comprising: visual, audio, touch, motor, neurochemical or location modalities.

6. The method of claim 1 wherein Experiences in the Experience Memory Store include authored Experiences.

7. The method of claim 1 wherein the second representation in the second modality corresponds to an internal state of the Embodied Agent, and the animation of the Embodied Agent is directed by modifying its current internal state according to the retrieved internal state.

8. The method of claim 7 wherein the second representation in the second modality corresponds to an emotional state of the Embodied Agent.

9. The method of claim 2 wherein the second representation in the second modality is re-experienced by the Embodied Agent during recall.

10. The method of claim 1 wherein the Experience Memory Store is a collection of SOMs.

11. The method of claim 10 wherein the SOMS are arranged as a CDZ.

12. A system for storing memory for an Embodied Agent, including:

    an Experience Memory Store, populated from Experiences experienced in the course of operation of the Embodied Agent; wherein each Experience is associated with a plurality of representations of the Experience in different modalities, and the Experience Memory Store stores representations of the Experiences in neural network weights;

    a Memory Database, storing copies of the Experiences stored in the Experience Memory Store, wherein the Memory Database stores raw data corresponding to the representations of the Experience

~~in different modalities.~~

~~13. A method of selectively storing Experiences experienced by an Embodied Agent in the course of live operation of the agent including the steps of:~~

~~receiving representations of input from a plurality of input streams for receiving input in a plurality of modalities; wherein each input stream is associated at least one condition which creates an eligibility trace in the input stream;~~

~~detecting simultaneous eligibility traces of two or more input streams ("Eligible" input streams); and storing and associating the representations of input from the Eligible input streams.~~

~~14. The method of claim 13 wherein the conditions which create an eligibility trace include novelty, importance, or intensity.~~

~~15. The method of claim 13 wherein the plurality of modalities include one or more of the group comprising: visual, audio, touch, motor, neurochemical or location modalities.~~

~~16. A memory system comprising:~~

~~a short-term memory Self-Organizing Map (STM) storing memories experienced since a last consolidation; and~~

~~a long-term memory structure (LTM) storing memories experienced before the last consolidation, trained during consolidation on data pairs from STM activating trained units;~~

~~wherein the LTM and STM map representations in a representational space onto a 1-hot encoding of category types;~~

~~wherein memories are stored collectively and disjunctively in the STM and the LTM.~~

~~17. The memory system of claim 16 wherein category types are object types.~~

~~18.~~1. A method for training a SOM, including a plurality of Neurons, each Neuron associated with a weight vector, and a training record; including the steps of:

receiving an input vector;

determining if the input vector is "new";

if the input vector is not new:

selecting a first Winning Neuron, favouring higher similarity between the Input Vector and the Winning Neuron, and modifying the weight vector of the first Winning Neuron towards the input vector;

if the input vector is new:

selecting a second Winning Neuron, favouring neurons with lower training records, and modifying the weight vector of the second Winning Neuron towards the input vector.

~~19.~~2. The method of claim ~~18~~ 1 wherein determining if an input vector is new includes the steps of:

determining a first Winning Neuron, favouring higher similarity the Input Vector and the Winning Neuron;

determining a match quality between the input vector and the first Winning Neuron;

determining the input vector as new if the match quality is below a match quality threshold.

~~20.~~3. The method of claim ~~19~~ 2 wherein match quality is determined as the activation of the first Winning Neuron in response to the Input Vector.

~~21.~~4. The method of claim ~~19~~ 2 wherein if the input vector is not new, adjusting the learning frequency of the SOM according to the activity of the first Winning Neuron.

~~22.~~5. The method of claim ~~19~~ 2 wherein selecting a second Winning Neuron, includes favouring neurons with lower training records, and higher similarity between the Input Vector and the Winning Neuron.

~~23.~~6. The method of claim ~~22~~ 5 wherein favouring neurons with lower training records is achieved by applying an Activation Mask comprises Mask Values inversely proportional to the amount of training received by their respective Neurons.

~~24.~~7. The method of claim ~~18~~ 1 wherein the training record is a weight and the value of the training record is proportional to the amount of training the training record's associated Neuron has received.

~~25.~~8. The method of claim ~~19~~ 2 including the step of adding random noise to the activation map of the SOM to the input vector, prior to selecting a second Winning Neuron.

~~26.~~9. The method of claim ~~19~~ 2 wherein the training record of Neurons are configured to decay with SOM training and/or time.

~~27.~~10. A method for targeted forgetting in a SOM including a plurality of Neurons, each Neuron associated with a weight vector, including the steps of:

creating a Reset Mask comprising a plurality of Mask Values, each Mask Value associated with a Neuron of the SOM;

applying a Reset Function to each Neuron of the SOM associated with a non-zero Mask Value, wherein the Reset Function includes:

a Forgetting Component, for resetting the Neuron's weights to an untrained state; and

a Mask Component for modifying the output of the Reset Function as a function of the Forgetting Component and the Mask Value;

modifying SOM Neuron weight vectors according to the output of the Reset Function.

~~28.~~11. The method of claim ~~27~~ 10 wherein the Forgetting Component creates random noise.

~~29.~~12. The method of claim ~~27~~ 10 for targeted forgetting of memories associated with an Input Vector wherein the step of creating the Reset Mask includes the steps of, for each Mask Value of the Reset Mask:

determining a similarity between the Input Vector and the Mask Value's associated Neuron, determining if the similarity is above a Reset Threshold;

if the similarity is above the Reset Threshold, setting the Mask Value to 1; otherwise if the similarity is not above the Reset Threshold, setting the Mask Value to 0.

~~30.~~13. The method of claim ~~27~~ 10 wherein the SOM is a probabilistic SOM wherein an activity of each Neuron is determined using an Activation Function mapping a distance to a probability space and weights of Neurons are updated according to Neuron activity.

~~31.~~14. The method of claim ~~30~~ 13 for targeted forgetting of memories associated with an Input Vector, including the steps of: providing the input vector as input to the SOM; and setting the Reset Mask according to the activation map of the SOM.

~~32.~~15. The method of claim ~~27~~ 10 wherein each neuron is associated with a training record.

~~33.~~16. The method of claim ~~32~~ 15 for targeted forgetting of infrequent memories, wherein the Reset Mask is set to the inverse of the training record.

~~34.~~17. The method of claim ~~32~~ 15 for targeted forgetting of memories older than a Forgetting Threshold including the step of:

eroding the training record of each neuron over time; and

setting the Reset Mask to the inverse of the training record.

~~35.~~18. The method of claim ~~32~~ 15 wherein the step of creating the Reset Mask includes the steps of, for each Mask Value of the Reset Mask:

determining if the Training Record of the Mask Value's associated Neuron is below a Reset Threshold;

if the Training Record is below the Reset Threshold, setting the Mask Value to 1; otherwise if the Training Record is not above the Reset Threshold, setting the Mask Value to 0.

~~36. A method-implemented system for directing the behaviour of an Embodied Agent, the method including:~~

~~receiving sensory input corresponding to a first representation of an Experience in a first modality;~~

~~querying an Experience Memory Store to retrieve a second representation of the Experience in a second modality; and~~

~~using the second representation in the second modality to direct the behaviour of the Embodied Agent.~~

~~A method-implemented system for selectively storing Experiences experienced by an Embodied Agent in the course of operation of the agent including the steps of:~~

~~receiving representations of input from a plurality of input streams for receiving input in a plurality~~

~~of modalities; wherein each input stream is associated at least one condition which creates an eligibility trace in the input stream;~~

~~detecting simultaneous eligibility traces of two or more input streams ("Eligible" input streams); and~~

~~storing and associating the representations of input from the Eligible input streams.~~