

Configuring the Stochastic Helmholtz Machine for Subcortical Emotional Learning

Chi-Yung Yau, Kevin Burn and Stefan Wermter

Abstract—Emotional learning involves two stages. The first is to acquire reinforcers from stimuli and the second is to associate such reinforcers with emotional responses. Both stages can be found occurring in the amygdala. LeDoux’s fear circuit model [1] suggests two routes, a subcortical route and a cortical route, for emotional information entering the amygdala for associative learning. It can be used to explain how the actual recognition of emotions from facial expressions can be processed in the brain. Based on the model, a neural architecture is proposed using the stochastic Helmholtz machine (SHM) with the wake-sleep algorithm. In this paper, the results of three experiments about the subcortical emotional learning are reported, where different configurations of SHMs are involved. The first two experiments are to identify a suitable way to allow behavioural responses entering the central nucleus of the amygdala for association. However, both experiments show symptoms of overfitting, where some weights and biases of neurons are observed that will unusually increase during training. Therefore, the final experiment is designed to maintain the range of weights between -1 and $+1$ in order to solve the overfitting problem. The last experiment shows that the neural architecture with the new weight policy holds a lot of potential for modelling subcortical learning.

I. INTRODUCTION

A computational system for online-learning of emotional meanings from facial expressions is an anticipated capability for future interactive robots [2-4]. However, it is a difficult task because human facial expressions are not only used for expressing emotions but also for communication [5]. Thus, a facial expression can be generated without relating to any emotion. In addition, daily experienced emotions are social emotions, which are more sophisticated than basic emotions and usually are mixtures of multiple emotions. This idea is supported by Plutchik’s theory [6] to use colours to explain emotions. One emotion can be mixed with another emotion to become a new and ambiguous emotion.

Such an idea is also supported by a biological theory from Rolls [7]. He explained emotional learning as a form of stimulus-response associative learning, which is related to two types of learning called classical conditioning [8] and instrumental learning [9]. There is an emotional stimulus called a ‘reinforcer’ is used as an intermediate emotional stimulus to bridge the sensory stimuli and emotional

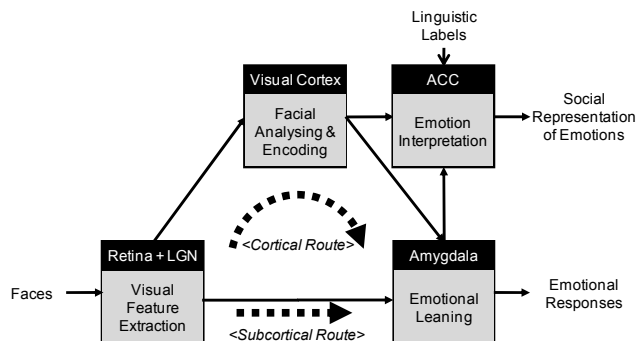


Figure 1. A block diagram shows the core structure of the neural architecture. (1) Visual features of faces are extracted at the retina & LGN module. (2) The extracted features are sent to the amygdala to recall reactive responses and the visual cortex to obtain detailed analysed and encoded results of faces respectively. (3) The encoded facial information is sent to the amygdala to produce higher level responses. (4) All information are integrated and interpreted at the ACC.

responses. Some sensory stimuli can recall primary reinforcers and elicit autonomic emotional responses directly without learning, such as an unexpected scream causing a shock face. Such stimulus-response associations are innately present in the brain. Then, a new non-emotional stimulus, such as a horrified face, can dynamically associate with primary reinforcers, thereby becoming secondary reinforcers in order to trigger similar emotional responses as original primary emotional responses. As a result, humans own various reinforcers and different levels of emotional responses.

The recognition of emotions from facial expressions in the brain can be considered as the interpretation of reinforcers [10], where the reinforcers are recalled by facial expressions. In this case, the recognition is more likely to acquire and memorise several associations between facial expressions and reinforcers, instead of categorising and grouping facial expressions to a particular emotional category.

Therefore, the choice of machine learning approach must support the dynamic associations and be configured for long term continuous running in robots. In this case, it is unlikely that most typical categorisation-based pattern classification algorithms are appropriate to use. In addition, because emotional learning is involved in many autonomic functions, an unsupervised learning method is desired. With all requirements considered, the stochastic Helmholtz machine (SHM) [11], which is a type of bi-directional associative memory, is examined in our research.

Mr. Chi-Yung Yau, and Dr. Kevin Burn are at the Faculty of Applied Sciences, Department of Computing, Engineering and Technology, University of Sunderland, U.K. (e-mail: alan.yau@sunderland.ac.uk; kevin.burn@sunderland.ac.uk). Professor Dr. Stefan Wermter is at the Institute for Knowledge Technology, Department of Informatics, University of Hamburg, Germany (wermter@informatik.uni-hamburg.de)

II. OVERVIEW OF NEURAL ARCHITECTURE

Although the focus of this paper is to report the emotional learning results obtained from the subcortical route, it is useful to give an overview of the biologically inspired neural architecture [12] as shown in Figure 1.

Emotion recognition from facial expressions is considered as a process to integrate and interpret two outputs, from emotional learning and facial perception. Following the idea, the architecture is designed based on the model of LeDoux's fear model [13,14], which describes how the amygdala receives sensory inputs via the subcortical route and cortical route to perform emotional leaning and generate different levels of fear response. The architecture comprises four separated regions, which are responsible for particular cognitive and emotional functions. The retina and LGN are responsible for colour and basic edge extraction for all modules. The visual cortex is specialised for analysing visual inputs, separating mixture of signals and encoding them as compact representations. The amygdala is specialised for multiple levels emotional learning. Finally, the anterior cingulate cortex (ACC) acts as an emotion and cognition interface to receive and centralise different sensory, cognitive and emotional information at one place, to enable the social interaction and emotional problem-solving [15].

The difference between the subcortical route and the cortical route is the processing speed and the level of detail of information. The subcortical route allows visual information to bypass the visual cortex entering the amygdala as a rapid channel, and the amygdala can directly recall reinforcers and elicit emotional responses. However, the quality is rough and assumes no sophisticated emotional response needs to be generated. This route is especially important to enable humans to avoid danger because it supports 'sense danger' rather than 'recognise danger'. Conversely, the cortical route allows the visual information to be processed by the visual cortex before being sent to the amygdala for emotional learning. Thus, the speed is relatively slow but more detail and accurate information are provided.

III. STOCHASTIC HELMHOLTZ MACHINE

The stochastic Helmholtz machine (SHM) is an unsupervised learning associative memory designed to simulate the cortical bottom-up and top-down pathways of the human perceptual system [11,16,17]. It works like a statistic inference to sample input sensory data and




Figure 2. The structure of the SHM

stochastically infer a binary pattern for representing original input data. It can be considered as an alternative approach to the Restricted Boltzmann machine (RBM) [18] for deep belief nets [19] and deep learning [20,21]. Thus, it can be used for compressing and encoding a sensory feature input.

The original SHM is a three layer model; however, it is replaced by a two layer model (as shown in Figure 2) in order to advance the implementation of the neural architecture. The first layer (A) is the input layer and the second layer is the output layer (B). There are two separated connection models: a recognition model and a generative model. A recognition model refers to the recognition weights (R^{ab}) that connect layer A to layer B. A generative model refers to the generative weights (G^{ba}) that connect layer B back to layer A.

All neurons are stochastic binary neurons containing two states, 0 or 1. However, when the input data is non-binary such as a natural image, the probability distribution of neurons at the input layer is used instead. The mathematical expressions for the neurons are:

$$P(s_y^b = 1) = \frac{1}{(1 + \exp(-\sum_x s_x^a \cdot R_{xy}^{ab} + b_y^R))} \quad (1)$$

$$P(s_x^a = 1) = \frac{1}{(1 + \exp(-\sum_y s_y^b \cdot G_{yx}^{ba} + b_x^G))} \quad (2)$$

where s^b is the actual state of a neuron in layer B with index y ; s^a is the actual state of a neuron in layer A with index x ; R^{ab} is the recognition weight with bias b^R ; G^{ba} is the generative weight with bias b^G .

The recognition model is a bottom-up model responsible for estimating an economical representation for describing the input data and separating the generative model. The generative model is a reverse top-down model responsible for reconstructing the input data based on the estimated representation. Two models connected as a circular loop allow information exchange between layers.

For learning, the wake-sleep algorithm was chosen. It is an unsupervised learning algorithm consisting of two phases, the 'wake' phase and the 'sleep' phase, to update the recognition weights and the generative weights separately using the delta rule.

A. Wake Phase Learning

The details of training in the wake phase are illustrated in Figure 3, where the different processing steps are shown. It can be considered in two steps: step 1 and step 2.

In step 1, the input data is entered into the input layer (A). The recognition model then calculates a probability distribution over the neurons $P_R(s^b)$ at the output layer (B), where the probability of each neuron $P(s^b=1)$ can be calculated using equation 1. The purpose is to detect features from input data and recognise them as a binary representation.

In step 2, the generative model calculates a probability distribution over the neurons $P_G(s^a)$ at the reconstruction layer (A2), which can be calculated using equation 2. The difference between the original input data vector at layer (A) and the reconstructed data vector at layer (A2) can then be used to adjust the generative weight and its bias by using equations 3 and 4 respectively.

$$G_{yx}^{(n+1)} = G_{yx}^{(n)} + \gamma(s_x^a - s_x^{aG})s_y^b \quad (3)$$

$$b_x^{G(n+1)} = b_x^{G(n)} + \gamma(s_x^a - s_x^{aG})(-1) \quad (4)$$

where n is the iterative count; $(s^a - s^{aG})$ is the difference between the original input data vector and the reconstructed data vector; γ is the learning rate.

There are two outputs: the binary representation is outputted to upper SHMs for further processing and the reconstructed input data is used as feedback for guiding the learning of lower SHMs.

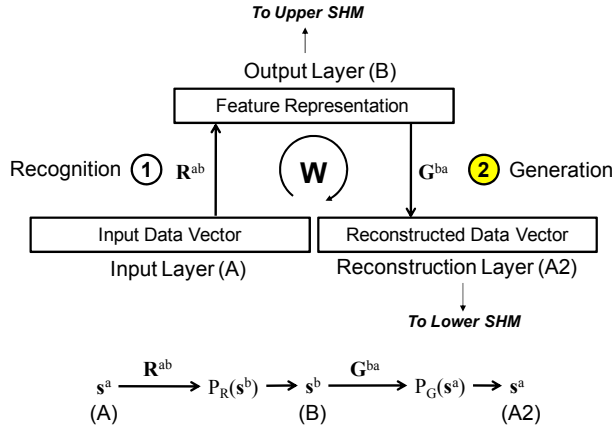


Figure 3. An unfolded view shows the learning sequences of wake phase

B. Sleep Phase Learning

Sleep phase learning is important to separate the generative model. The details of training are illustrated in Figure 4. It can also be considered in two steps: step 3 and 4.

In step 3, the input of fantasy layer (B3) is a binary selected feature randomly sampled from either the feedback of other SHMs or a stochastic pattern. The generative model then calculates the probability distribution of neurons $P_R(s^a)$ in the approximation layer (A3) using equation 2.

In step 4, the recognition model calculates the probability distribution of neurons $P_R(s^b)$ in the simulation layer (B4) using equation 1. Finally, the difference between layer B3 and layer B4 can be used to adjust the recognition weight and its bias by using equations 5 and 6 respectively.

$$R_{xy}^{(n+1)} = R_{xy}^{(n)} + \gamma(s_y^b - s_y^{bR})s_x^a \quad (5)$$

$$b_y^{(n+1)} = b_y^{(n)} + \gamma(s_y^b - s_y^{bR})(-1) \quad (6)$$

where n is the iterative count; $(s^b - s^{bR})$ is the difference between the sampled feature and the simulated feature; γ is the learning rate.

Basically, the update of the generative weight is driven by the recognition model in the wake phase for increasing the chance of reconstructing an accurate input. Alternatively, the update of the recognition weight is driven by the generative model in the sleep phase for improving the estimation of the input data.

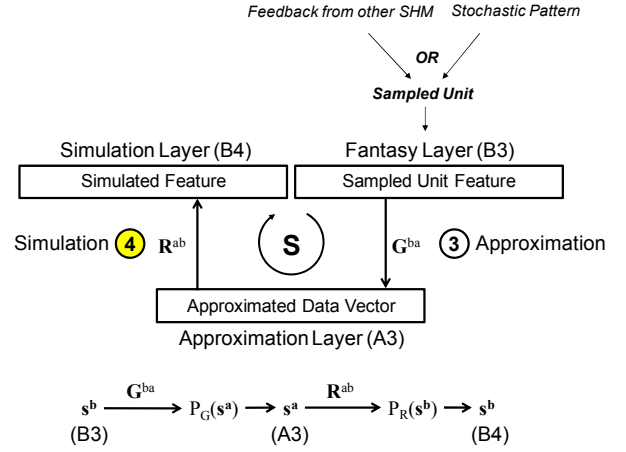


Figure 4. An unfolded view shows the learning sequences of sleep phase

IV. EXPERIMENTS

The experiments presented here are related to testing of different configurations of SHM for the amygdala. The results obtained will become the basis for the further experiments of the neural architecture in the future as shown in Figure 1.

A. Training Dataset

In order to benefit the analysis of SHM, a simple robotic expression image dataset was used. The dataset contains 144 robotic expression images, including all combinations of facial motor positions. There are 36 images for each position of left and right eyebrows and 48 images for each position of upper and lower lips. Therefore, the probability of each motor position being trained is quite even.

The robotic images are then manually assigned into one of four emotion categories: excitement, stress, depression and calm, with the meaning positive-high arousal, negative-high arousal, negative-low arousal and positive-low arousal respectively. These four categories are also used as the behaviour responses. As shown in Table I, the quantity of image data within each category is uneven because the categorisation is performed according to the primary feeling to images of expressions and no deliberation has been involved.

TABLE I
THE DISTRIBUTION OF FOUR EMOTION CATEGORIES

Category	Image
Excitement (57)	[1,2,3,4,5,6,7,8,9,13,14,15,17,22,23,24,26,31,32,40,41,42,44,49,50,51,53,58,59,67,68,69,71,72,76,77,78,80,81,85,86,94,95,96,98,99,103,104,112,113,114,116,117,121,122,130,131]
Calm (10)	[10,19,28,37,46,55,64,73,91,109]
Depression (20)	[11,12,20,21,29,30,38,39,47,48,56,57,65,66,74,75,92,93,110,111]
Stress (57)	[16,18,25,27,33,34,35,36,43,45,52,54,60,61,62,63,70,79,82,83,84,87,88,89,90,97,100,101,102,105,106,107,108,115,118,119,120,123,124,125,126,127,128,129,132,133,134,135,136,137,138,139,140,141,142,143,144]

B. Subcortical Emotional Learning Model

Figure 5 shows the subcortical emotional learning model used in different experiments. There are two modules involved: the retina and LGN module and the amygdala module. The retina and LGN module is used to extract the edged features from images, where a basic canny edge detector algorithm is applied. Two SHMs are used in the amygdala module for each sub-region, the basolateral complex (BLA) and the central nucleus (CeA). BLA directly receives the visual features and recognises them as different primary reinforcers at high speed as it is assumed that they are an innate connection. Its output is then sent to CeA for eliciting different emotional responses.




Figure 5. A subcortical emotional learning model

TABLE II
THE CONFIGURATION OF NETWORKS

Exp	Behavioural Response	BLA	CeA	Weight Limit
1	Top	(4980 – 64) [0.1, 0.001]	(64 – 4) [0.1, 0.1]	[0,+∞]
2	Side	(4980 – 64) [0.1, 0.0001]	(68 – 64) [0.1, 0.01]	[0,+∞]
3	Side	(4980 – 64) [0.001, 0.01]	(68 – 64) [0.01, 0.1]	[-1,+1]

* (Size of layer A – Size of layer B) [lr_w, lr_s]

* 'lr_w' is the learning rate for the wake phase and 'lr_s' is the learning rate for the sleep phase

C. Experimental Setup

Three experiments have been carried out with different configurations, as shown in Table II. Experiment 1 is designed to test the configuration with the behavioural responses entering CeA from the 'top'. Both BLA and CeA

are configured to receive feedback from CeA and CeA receives feedback (or supervision) from other cortical areas. Also, no negative weight is permitted. In contrast, experiment 2 is to test another configuration, where the behavioural responses are entered from the 'side', pairing with the output of BLA. Finally, experiment 3 is used to test the effect of limiting the value of connection weights between -1 and +1.

V. RESULTS

In order to evaluate the learning performance of the SHM, the average reconstruction error for every SHM is calculated based on the difference between the input data at layer A and the reconstructed data at layer A2.

A. Experiment 1

The average error of BLA is shown in Figure 7(a). The first half of the curve continuously drops in steps. Every step represents a new feature being discovered from the input data. Later, overfitting occurs in the second half of the learning curve. In a practical situation, however, this should not occur, as robots will be continuously exposed to new faces. A network will be selected only if it can satisfy the requirement of long-term learning without overfitting. Therefore, the current configuration of SHM for BLA is not yet desired.

In order to address this issue, a deeper analysis is necessary. Three points have been measured from the learning curve for representing different learning stages: on-going training, best training and overtraining. Their subsets of weights are plotted in Figure 6. The weights indicate that lips features are easier to discover than eyebrows features. The reason is that some non-overlapped features, such as lips, involve more input units (as pixels) resulting in easier and faster learning. In contrast, overlapped features, such as eyebrows, involve relatively fewer units; therefore, BLA need more iterations to separate them.




Figure 6. An example of weights extracted from three different learning stages: ongoing training best training and overtraining

The average error of CeA is shown in Figure 7(b). The learning curve of CeA drops rapidly at the beginning. However, it then rises and drops again at the end. The reason is that BLA and CeA were learnt together. At the beginning, not all features have been discovered by BLA and this

incompletely detected feature was sent to CeA for learning. After that, BLA continuously discovered features and CeA also adapted the change and started learning so that the curve is slightly rising in the middle stage and dropping again as CeA is learnt.

Table III summarises the range of parameters in BLA and CeA at the best training stage and the overtraining stage. Comparing the best training stage and overtraining stage, it indicates that some weights and biases at the overtraining stage are boosted to an unusual large value. In principle, the bias of neuron is an adaptive threshold of the total of the weighted inputs for keeping the weights changed within an effective region of a sigmoid curve. Therefore, the weights should not be very large and the bias should be just large enough needed. However, the result indicated that both of them increase continuously.

TABLE III
THE MINIMUM AND MAXIMUM RANGE OF PARAMETERS (EXP 1)

Best training (1,000)	Overtraining (2,000)
BLA.Rab [0, 0.58646]	BLA.Rab [0, 5.8483]
BLA.br [5.7828, 7.4531]	BLA.br [5.9083, 8.8663]
BLA.Gba [0, 27.6475]	BLA.Gba [0, 20.5761]
BLA.bg [-12.3183, 28.222]	BLA.bg [-12.3149, 55.2897]
CeA.Rab [0, 22.8772]	CeA.Rab [0, 22.9526]
CeA.br [9.4139, 11.4182]	CeA.br [11.3525, 12.3263]
CeA.Gba [0, 20.4676]	CeA.Gba [0, 22.7251]
CeA.bg [-2.0622, 16.1244]	CeA.bg [-4.1517, 16.8109]

B. Experiment 2

Due to the overfitting of BLA occurring in experiment 1, it was suggested that a slower detailed update of the recognition weights can improve the detection of features and reduce the overtraining. Therefore, the learning rate of the sleep phase is decreased.

The results of BLA and CeA are shown in Figure 8(a) and (b) respectively. Basically, the shape of the learning curves of BLA obtained in this experiment is almost the same as the curve obtained in experiment 1, but it is slower. However, the overtraining still occurred at the end. One possible reason is that the reduction of learning rate of the sleep phase only reduces the speed of learning, but cannot improve the estimation. It implies that the optimisation of SHM requires a balance of the actual learning speed between the recognition model and generative model.

Similarly, the result of CeA obtained here is close to the result obtained in the experiment 1. This experimental result indicates that whether the behavioural response is entered from the top of CeA or from the side pairing with the output of BLA, the learning curves are similar. Table IV summarises the range of the parameters in BLA and CeA, which shows that some weights are still dominant.

To summarise, although no significant performance improvement has been observed, the ‘side’ configuration is still useful if only four behavioural responses are included.

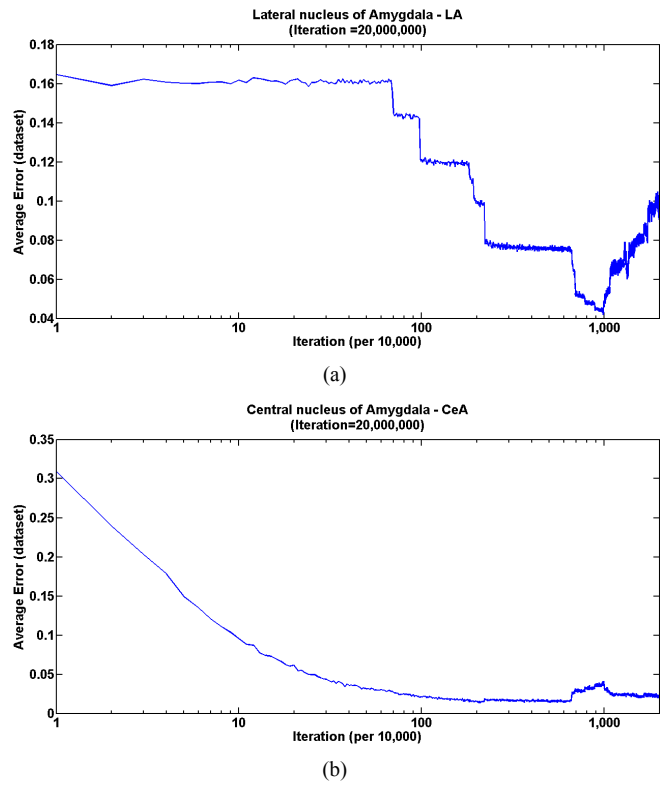


Figure 7. The average error of BLA (a) and CeA (b) for the experiment




Figure 8. The average error of BLA (a) and CeA (b) for the experiment 2

When the problem becomes more complicated, the ‘top’ configuration becomes difficult to manage; instead, the ‘side’ configuration should be able to deliver a reasonable result via increasing the size of the output layer.

TABLE IV
THE MINIMUM AND MAXIMUM RANGE OF PARAMETERS (EXP 2)

Best training (4,000)	Overtraining (5,000)
BLA.Rab [0, 1.149]	BLA.Rab [0, 2.458]
BLA.br [5.4607, 7.9958]	BLA.br [5.4684, 8.0648]
BLA.Gba [0, 30.5946]	BLA.Gba [0, 26.0096]
BLA.bg [-13.5879, 36.6932]	BLA.bg [-13.591, 36.7777]
CeA.Rab [0, 50.1305]	CeA.Rab [0, 60.1103]
CeA.br [4.4423, 12.3617]	CeA.br [4.4707, 14.1286]
CeA.Gba [0, 19.4583]	CeA.Gba [0, 23.8297]
CeA.bg [4.0958, 25.75]	CeA.bg [4.1461, 26.5465]

C. Experiment 3

Based on the previous experiences, the ‘balance’ of the learning performance of the two connection models in a SHM is important. In order to deal with this, the idea is to avoid the weights and biases to become too large. The best value of weights should always be kept moving along the effective range in a sigmoid curve. Therefore, a limit is applied to maintain the range of weights between -1 and +1 to emulate a biological neuron’s ability to inhibit and excite.

Additionally, in the wake phase, many units are updated at once. Conversely, in the sleep phase, only one unit is updated every time. Thus, the generative weights are initially updated faster than the recognition weights. For this reason, the learning rate of the wake phase is decreased to smaller than the sleep phase in order to balance the update speed of the recognition model and the generative model.

The result of the BLA in the experiment 3 is shown in Figure 9(a). The learning curve drops rapidly at the beginning and then keeps stable for a long time. The result of the CeA is shown in Figure 9(b). The curve slightly drops at the beginning when no features have been detected by BLA. However it quickly recovers since the BLA learns fast. The curve then rises back to the normal training level before continuously decreasing gradually.

Since no overtraining is observed, Table V summarises the range of the parameters in BLA and CeA at the early training stage and the on-going stage instead. The result indicates that

TABLE V
THE MINIMUM AND MAXIMUM RANGE OF PARAMETERS (EXP 3)

Early training (100)	On-going training (3,000)
BLA.Rab [-1, 0.99666]	BLA.Rab [-1, 1]
BLA.br [0.92835, 1.2641]	BLA.br [0.76513, 2.4521]
BLA.Gba [-1, 1]	BLA.Gba [-1, 1]
BLA.bg [-4.3729, 4.7831]	BLA.bg [-8.6574, 9.9374]
CeA.Rab [-1, 1]	CeA.Rab [-1, 1]
CeA.br [3.0499, 6.4116]	CeA.br [0.81233, 9.7944]
CeA.Gba [-1, 1]	CeA.Gba [-1, 1]
CeA.bg [0.011917, 8.2109]	CeA.bg [-0.80241, 10.3686]




Figure 9. The average error of BLA (a) and CeA (b) for the experiment

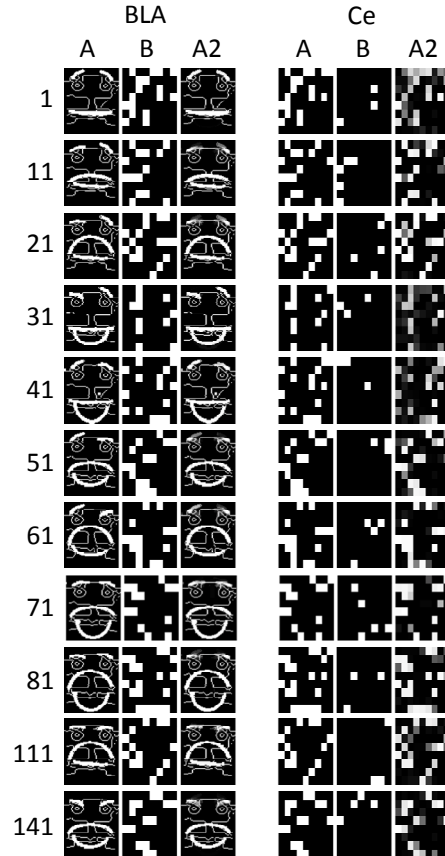


Figure 10. Some examples of BLA (left) and CeA (right) in the experiment 3 including the original input, output and reconstructed input.

VI. DISCUSSION

Several computational models of the amygdala have been suggested. They usually target the conditioning behaviours of the amygdala, using neuron-level simulation [22] or topographic conditioning maps [23]. In contrast, our research targets the recognition of emotions from facial expressions via the subcortical and cortical routes using SHM, where the amygdala is core in the emotional learning.

The study presents the experimental results for configuring the SHM for subcortical emotional learning. The main contribution is to model the subcortical route in which the amygdala receives visual information from the thalamus (LGN), recall primary reinforcers and associate reinforcers to behavioural responses. The subcortical route contains the characteristics of high speed and low details. Therefore, the internal structure of the amygdala is simple and only BLA and CeA, are involved. The edged feature of robotic expressions without any segmentation is used to emulate the low detail input. In addition, robots will be continuously exposed to new faces in a practical situation; therefore, the selected approach would support long-time continuous learning and the ability to dynamically associate stimuli with various responses.

The SHM has been used as it is a type of associative memory using unsupervised learning. Its learning speed is not very fast; however, the execution time is fast. Also, it learns based on a simple algorithm and contains two models to assist mutual learning in order to learn without supervision. This is an important characteristic to model different autonomic structure in the brain. Potentially, it can avoid the network overtraining and is suitable for use in long-term learning applications. In addition, the SHM can be practically configured for handling both encoding and association tasks.

Biologically, the association of some particular stimuli, primary reinforcers and behavioural responses can be considered as innate connections. All other associations should be formed based on learning, where an input from the cortex should be present. For this reason, experiments 1 and 2 are designed to study which configuration of SHM in CeA is better for associating the reinforcers retrieved from BLA and the behavioural responses from the cortex.

In experiments 1 and 2, both have a problem of overtraining. Addressing this problem, experiment 3 attempts to restrict the range of recognition and generative weights. The results show a faster, smoother and more stable learning behaviour and a good encoding outcome. It indicates that the limit of weights can control the overfitting problem. Finally, SHM is proved capable of encoding facial expressions without any detail and complicated analysis, and associating the encoded results (as reinforcers) with behavioural responses.

Future work should include the integration of the cortical sections in order to demonstrate parallel processing behaviours. Also, it can allow the subcortical model responsible to awake some reactive responses which can




Figure 11. A subset of the recognition weights (a) and the generative weights (b) of BLA. The gray colour represents 0 (natural), the white colour represents 1 (excite) and the black colour represents -1 (inhibit).

the value of weights and biases is much smaller compared to previous results. It can thus be assumed that the activations of neurons are controlled in an effective range.

In order to robustly confirm the performance, some examples of BLA and CeA are plotted in Figure 10. This includes the original input, output and reconstructed input at layer A, B and A2 respectively.

Different robotic facial expressions are clearly separated and encoded at BLA. The output at layer B can be considered as the representation of primary reinforcers and it can be used for further associations with the cortical inputs. With reference to CeA, the first eight columns represent the output of BLA and the rightmost column represents the behavioural responses. It can be seen that the reinforcers are successfully associated with the behavioural responses. The correct associated rate is 121/144 and it is about 84% after 30,000,000 times of learning cycles (iterations). The accuracy should further increase as the trend of the learning curve for CeA is still decreasing.

Figure 11 shows a subset of the recognition weights and the generative weights of BLA and CeA. In this case, the weights closely represent the factors of inputs. Therefore, Figure 6 shows a clear result since it only contains positive weights. However, the current weights of BLA obtained in this experiment involve both positive and negative factors, and therefore it is not contributing for tracking the learning.

assist recognition in the cortex. For example, when the subcortical model issues a fear response, the cortex can change to a particular mode that is mainly for detecting the source of fear or danger.

REFERENCES

- [1] J.E. LeDoux, "Emotion, memory and the brain," *Scientific American*, vol. 270, 1994, pp. 50-57.
- [2] R.W. Picard and J. Klein, "Computers that recognise and respond to user emotion: theoretical and practical implications," *Interacting with Computers*, vol. 14, Feb. 2002, pp. 141-169.
- [3] C. Breazeal, "Toward sociable robots," *Robotics and Autonomous Systems*, vol. 42, Mar. 2003, pp. 167-175.
- [4] M.A. Arbib and J. Fellous, "Emotions: from brain to robot," *Trends in Cognitive Sciences*, vol. 8, Dec. 2004, pp. 554-561.
- [5] R. Adolphs, "Recognizing emotion from facial expressions: psychological and neurological mechanisms," *Behavioral and Cognitive Neuroscience Reviews*, vol. 1, Mar. 2002, pp. 21-62.
- [6] R. Plutchik, "The nature of emotions," *American Scientist*, vol. 89, 2001, pp. 344-345.
- [7] E.T. Rolls, "Vision, emotion and memory: from neurophysiology to computation," *International Congress Series*, vol. 1250, Oct. 2003, pp. 547-573.
- [8] I.P. Pavlov, *Conditioned Reflexes*, Oxford University Press, 1927.
- [9] E.T. Rolls, *Emotion Explained*, USA: Oxford University Press, 2005.
- [10] A.M. Glenberg, D. Havas, R. Becker, and M. Rinck, "Grounding language in bodily states: The case for emotion," *Grounding Cognition: The Role of Perception and Action in Memory, Language, and Thinking*, D. Pecher and R.A. Zwaan, eds., Cambridge University Press, 2005, pp. 115-128.
- [11] P. Dayan, G.E. Hinton, R.M. Neal, and R.S. Zemel, "The Helmholtz machine," *Neural Computation*, vol. 7, Sep. 1995, pp. 889-904.
- [12] C.Y. Yau, K. Burn, and S. Wermter, "A neural wake-sleep learning architecture for associating robotic facial emotions," *Proceedings of IEEE International Joint Conference on Neural Networks (IJCNN), Hong Kong, China, 1-6 June 2008*, 2008, pp. 2715-2721.
- [13] J.E. LeDoux, "Emotion circuits in the brain," *Annual Review of Neuroscience*, vol. 23, 2000, pp. 155-84.
- [14] J.E. LeDoux, "The amygdala," *Current Biology*, vol. 17, Oct. 2007, pp. R868-R874.
- [15] J.M. Allman, A. Hakeem, J.M. Erwin, E. Nimchinsky, and P. Hof, "The anterior cingulate cortex: The evolution of an interface between emotion and cognition," *Annals of the New York Academy of Sciences*, vol. 935, May. 2001, pp. 107-17.
- [16] G.E. Hinton, P. Dayan, B.J. Frey, and R.M. Neal, "The Wake-Sleep Algorithm for Unsupervised Neural Networks," *Science*, vol. 268, May. 1995, pp. 1158-1161.
- [17] P. Dayan, "Helmholtz machines and sleep-wake learning," *The Handbook of Brain Theory and Neural Networks*, M.A. Arbib, ed., The MIT Press, 2002, pp. 522-525.
- [18] R. Salakhutdinov and G.E. Hinton, "Deep Boltzmann machines," *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2009, pp. 448-455.
- [19] J.M. Susskind, G.E. Hinton, J.R. Movellan, and A.K. Anderson, "Generating Facial Expressions with Deep Belief Nets," *Affective Computing: Focus on Emotion Expression, Synthesis and Recognition*, J. Or, ed., I-Tech Education and Publishing, 2008.
- [20] Y. Bengio, *Learning deep architectures for AI*, Universit'e de Montr'eal, 2007.
- [21] G.E. Hinton, "To recognize shapes, first learn to generate images," *Computational Neuroscience: Theoretical Insights into Brain Function*, P. Cisek, D. Trevor, and J. Kalaska, eds., 2007, pp. 535-547.
- [22] J. Mor'en and C. Balkenius, "A computational model of emotional learning in the amygdala," *From animals to animats 6: Proceedings of the 6th International Conference on Simulation of Adaptive Behavior (SAB)*, J. Meyer, A. Berthoz, D. Floreano, H.L. Roitblat, and S.W. Wilson, eds., The MIT Press, 2000, pp. 383-391.
- [23] A. Pavlou and M. Casey, "Identifying Emotions Using Topographic Conditioning Maps," *Advances in Neuro-Information Processing: Proceedings of the 15th International Conference on Neuro-Information Processing (ICONIP)*, Springer, 2009, pp. 40-47.