

A Computational Platform for Visual Fear Conditioning

Athanasios Pavlou and Matthew Casey

Abstract—Computational models of fear conditioning help us understand the sensory pathways and neural structures underlying fear elicitation in the brain. The majority of the existing models have focused on conditioning on auditory stimuli by simulating the processing of the amygdala, which is the main brain structure implicated in processing fearful stimuli. However, there is now a growing understanding of how fear is elicited from visual stimuli, but as yet we do not have sufficiently capable techniques that can be used to model visual fear conditioning. Masking experiments are a key psychophysics technique that can help us understand these pathways by observing the behavior of the amygdala when presented with visual input that is not consciously perceived (masked). The amygdala’s response is indicative of whether it is influenced more by the proposed sub-cortical pathway, than by the cortical pathway. In this paper, we present a computational platform for visual fear conditioning. We use the platform to model the visual pathways leading to the amygdala and with them simulate masking experiments to explore the hypothesis that a sub-cortical pathway exists. The platform uses a modularized Hebbian learning architecture that can organize inputs topographically and condition on multiple stimuli representing visual inputs. We evaluate the properties and behavior of the platform and its capability in simulating masking experiments by comparing our simulation results with those observed for human behavior. Our results provide computational evidence for the influence the sub-cortical pathway has on the amygdala.

I. INTRODUCTION

FEAR is an innate defense mechanism that prepares organisms to deal with a potential threat [1]. It is not surprising that the expression of fear is observed in animals and humans alike [2], since ensuring survival is a matter of crucial importance for any organism. The universal character of fear has enabled scientists to perform a variety of experiments ranging from behavioral studies on animals [3] to human brain lesion and functional magnetic resonance imaging (fMRI) studies [4] to understand fear. From these we have developed a good understanding of the underlying mechanisms, particularly for the auditory system [3], [5]. Computational models have helped us understand the auditory fear mechanism [6], however, similar models for the visual pathways have not been developed despite the importance of visual stimuli in provoking fear responses [7], [8].

Our understanding of fear originally came from Pavlovian (or classical) conditioning experiments [1]. In these experiments, an otherwise neutral input, called the conditioned stimulus (CS) is temporally paired with an aversive input, called the unconditioned stimulus (US). The US may have

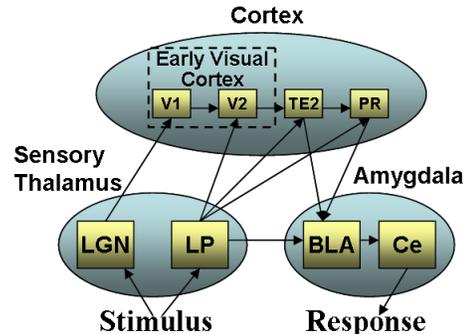


Fig. 1. Visual fear conditioning pathways as described in [5], showing the Lateral Geniculate Nucleus (LGN), Lateral Posterior Nucleus (LP), Basolateral Amygdala Complex (BLA), Central Amygdaloid Nucleus (Ce), Primary (V1) and Secondary (V2) Visual Cortices, Temporal Cortical Areas (TE2), and Perirhinal Cortex (PR).

the form of a loud burst of white noise or a foot shock, while the CS has varied from seeing food [1], lights [3] or emotional faces [9]. What these studies have shown is that the evaluation of potentially threatening stimuli is performed by a neural structure called the amygdala, located in the telencephalon.

The amygdala appears to receive input from two distinct visual pathways (Fig. 1). The cortical pathway for the visual modality follows the route of the lateral geniculate nucleus (LGN) of the thalamus, to visual cortical areas and extrastriate regions. The sub-cortical (direct) pathway is proposed to go through the lateral posterior nucleus (LP) (in particular through the superior colliculus and pulvinar) of the thalamus to the amygdala [5], [10]. Although the existence of the cortical pathway has been established [3] there is still some debate as to whether the sub-cortical pathway exists [11]. A typical approach to exploring this issue is through masking experiments, which use short presentations of stimuli to control conscious versus sub-conscious perception and reveal the participating neural structures, through imaging of active brain areas [4], [9].

Morris, Öhman and Dolan [9] conducted a series of masking experiments to explore the existence of the visual sub-cortical pathway. In the first phase of their experiments, the subjects were habituated on a number of stimuli (in this case random, sequential presentations of two neutral and two angry faces at intervals of 15-20 seconds). Second, the subjects were presented with the same stimuli, but with one selected as the CS (one of the angry faces), which they were then conditioned on (using 100-dB of white noise). Third, all possible pairings of the four inputs were presented to the subjects at 5 second intervals. The first image in a pair (the target) was presented for 30ms, followed by the second

Athanasios Pavlou and Matthew Casey are with the Department of Computing, University of Surrey, Guildford, Surrey, GU2 7XH, UK (phone: +44 (0) 1483 689635; fax: +44 (0) 1483 686051 email: {a.pavlou, m.casey}@surrey.ac.uk).

image immediately (the mask) for 45ms. The aim of this approach was to block the conscious perception of the target. In this way the effect of the target would be restricted to the proposed sub-cortical pathway, since the cortical pathway requires longer for the signals to propagate. Although such psychophysics experiments provide important evidence on the participating brain structures, it is harder to gain further evidence without more invasive or lesioning studies. Here, computational modeling can help by allowing us to test, for example, the influence of connectivity in an abstract model under controlled conditions.

In this paper we present a computational model of visual fear conditioning that includes both a cortical and sub-cortical pathway to the amygdala. We build directly upon the model used by Armony et al, who used a competitive algorithm to explore auditory fear conditioning [6]. In their model, Armony et al used a series of interconnected modules, each consisting of a single layer of neurons, which were capable of exhibiting simple receptive field properties for stimuli representing single auditory frequencies. We first extended this model by incorporating topographic organization, which is a property met in thalamic and early cortical areas both for the visual and auditory modalities (a preliminary evaluation of this modified model was reported in [12]). In the work we report here, we have further modified the model to change the input to represent the presentation of a target and a mask and then evaluated this as a suitable platform for masking experiments. Whereas Armony et al’s work reproduced the required conditioning behavior for abstract auditory stimuli represented in one dimension, conditioning on just one example, our extended model can cope with more complex two-dimensional representations of the visual space and condition on multiple stimuli. This increase in capability is needed to represent visual processing.

In our experiments, we first demonstrate the basic properties of our model, such as developing topographic organization, conditioning on multiple stimuli, and performing coordinate alignment of the outputs from thalamic modules in higher layers. Second, we conduct simulated masking experiments to explore the difference in activation observed between our model of cortical processing and that of the amygdala upon presentation of a mask-target pair. In particular our experiments show that the model can capture some of the basic characteristics of the masking experiments such as non-conscious perception of target stimuli as well as a dependence of the amygdala activations upon thalamic input. Although our findings are not enough to claim computational verification of the visual sub-cortical pathway, our results demonstrate its potential importance, while the model provides a computational platform for further experimentation of visual fear conditioning.

II. METHOD

The modeling approach followed in this study provides an abstract representation of the anatomical structures participating in visual fear conditioning (Fig. 3). We note that our model does not deal with visual pathway development

as other models have provided such studies (such as for the visual cortex [13], [14]). Synaptic plasticity between these structures is a property theoretically formulated by Hebb [15]. Many studies [16], [17], [18] used variations of synaptic plasticity rules in order to study structural and behavioral properties of humans and animals alike. Such a model was developed by Armony et al [6], drawing inspiration from Rumelhart and Zipser’s competitive algorithm [16], and working on similar principals to other Hebbian learning models [13], [14]. Armony et al’s model explicitly deals with fear conditioning of the auditory modality. For this reason our work extends this for visual fear conditioning by introducing topographic representations of the input corresponding to features in an abstract visual space. We draw inspiration for topographic organization from Kohonen’s Self-organizing Map (SOM) [18], which uses the concept of a neighborhood to represent lateral inhibition between neurons. Our extension of Armony et al’s model implements a competitive neighborhood in which a set of winning neurons are rewarded with a higher degree of weight change compared to those outside the neighborhood, the activations of which are inhibited.

A. Input Representation

We use an abstract representation of visual input for our model. In developing a model of visual fear conditioning, we are attempting to determine whether it can topographically organize distinct stimuli, and then condition on one or more of these. During initial training, our input representation must therefore aid the establishment of topographic maps. The input must then still be sufficiently flexible to be used to test conditioning and at the same time allow us to draw clear conclusions of the model’s properties. Our input represents a simple visual scene corresponding to a series of spatial locations with azimuth $[-90, 90]$ and elevation $[-65, 65]$, corresponding approximately to a human’s visual field. Within this, we allow the positioning of an object at a discrete interval of 10, so that we can encode 19 different positions for azimuth and 14 for elevation (Fig. 2a), so for example an input at azimuth -90 and elevation 25 is encoded as the co-ordinates $(1, 10)$.

To represent a stimulus at each location, we use a simple Gaussian pattern of activity, such that each stimulus overlaps. For a stimulus at azimuth p and elevation q , we have:

$$x_{pq} = \lambda e^{-\left(\frac{p^2+q^2}{\sigma^2}\right)} \quad (1)$$

where λ is the maximum amplitude and σ the radius. Here we use $\lambda = 1$ and $\sigma = 10$.

B. Simulating Masking

For the simulation of the masking experiments, we present the model with a pair of inputs at the same time. Since our model is rate-coded, it assumes the synchronous presentation of input for a feedforward pass of the network. In order to encode a change in stimulus temporally, such that the propagation of activity for one input is still ongoing whilst a new input is presented, we use a simple time delay technique

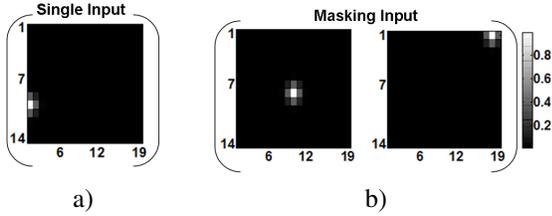


Fig. 2. The representation of a) a single input at (1, 10), and b) a masking input, mask (10, 8) and target (18, 1). Note, during habituation and conditioning, the pair of inputs are the same.

whereby two inputs from different time steps are presented at the same time (Fig. 2b and Fig. 3b).

These pairs of images represent the target, which in the psychophysics experiments is only presented for 30ms so that it is not consciously perceived, and the mask, which is presented for 45ms to allow for conscious perception. Therefore the distinct states of our input consist of a) the presentation of an identical pair (mask, mask) and b) the introduction of the target which is then masked (mask, target). During habituation we have a continuous presentation of stimuli without masking, and therefore we show pairs of identical inputs (mask, mask). During conditioning the same situation applies, except that one of the stimuli (presented as a pair) is conditioned on by turning on the US for that input only. Having conditioned, we explore masking by presenting two different stimuli (mask, target). For example in Fig. 2b, the mask is an input with stimulus at location (10, 8), whereas the target is (18, 1).

C. Cortical and Sub-cortical Pathway Representation

The model is formed from a series of feedforward neural modules that are trained using a competitive learning algorithm. Fig. 3a) and b) show a schematic of the model, which consists of four modules representing the LGN, LP, early visual cortices (VC) and the amygdala (AM). The cortical pathway is represented as connections from the LGN and LP to the VC, which then feed to the AM. The sub-cortical pathway feeds the output of the LP direct to the AM. The visual stimulus is input to the LGN and LP. To condition the model on a stimulus, we use an additional input to the LP and AM (equivalent to the US). When conditioning on the CS, the value of the US is set to 1 just for this input, and at all other times the US is set to 0. The US always has a fixed weight value associated with it, which for us is 0.7.

Each module consists of a lattice of neurons that are fully connected to the input, such that a neuron (i, j) has an output y corresponding to an m -dimensional input x :

$$u_{ij} = \sum_{k=1}^m x_k w_{kij}(t) \quad (2)$$

$$y_{ij} = \begin{cases} f(u_{ij}) & \text{if } \|c_{ij} - c_{win}\| < h(t) \\ f(u_{ij} - y_{win}) & \text{otherwise} \end{cases} \quad (3)$$

$$f(u) = \begin{cases} 1 & u \geq 1 \\ u & 0 < u < 1 \\ 0 & u \leq 0 \end{cases} \quad (4)$$

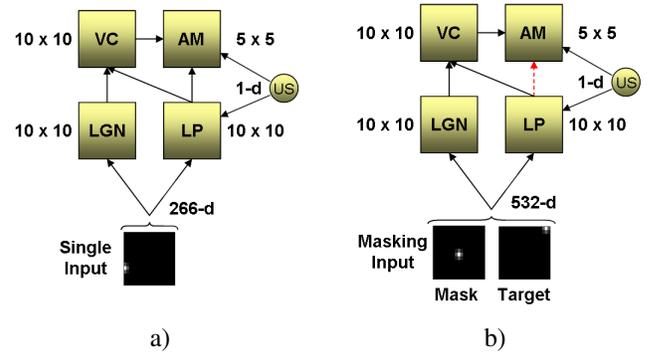


Fig. 3. Schematic of the model showing the Lateral Geniculate Nucleus (LGN), Lateral Posterior Nucleus (LP), early visual cortices (VC) and the amygdala (AM) for a) a single Gaussian input, and b) the mask-target paired Gaussian input. Note that the masking experiments are conducted on a model with and without the connection from the LP to the AM (red arrow).

where $w_{kij}(t)$ is the weight from input k for neuron (i, j) in the lattice at time step $t \geq 0$, initialized with uniformly distributed small random values. Note in equation 3 that a neuron is considered to be in the winner area if the distance from the neuron (i, j) to the winner in the lattice is less than the current radius value $h(t)$. Here we use c_{ij} and c_{win} to denote the lattice co-ordinates of the two neurons. All neurons outside this area are inhibited by the activation value of the winning neuron $y_{win} = \max_{ij} f(u_{ij})$.

The output y from a module (consisting of the values y_{ij} for each neuron) can then be fed into another module directly as input. If the output from several modules is combined (such as to the VC) or if a module output is combined with an input (such as to the AM with the US input), then the vectors are concatenated together.

Competitive learning in a module is achieved by updating each weight, except those fixed for the US, and then normalizing all weights to prevent exponential growth:

$$w'_{kij}(t+1) = w_{kij}(t) + \epsilon(t)x_k y_{ij} \quad (5)$$

$$w_{kij}(t+1) = \frac{w'_{kij}(t+1)}{\sum_{l=1}^m w'_{lij}(t+1)} \quad (6)$$

where $\epsilon(t)$ is the learning rate at time step t , corresponding to the presentation of a single input. This not only differs from Armony et al [6] in the use of a lattice and a neighborhood, but also in equation 5 with all weights being updated, and not just those that have an input that is above average.

III. EXPERIMENTS AND EVALUATION

We evaluate our model in two ways. First, we evaluate whether the model is adequate at simulating the basic properties of the neural pathways through its capability in capturing the topographic relationships amongst the input data and conditioning on an arbitrary stimulus. Second, we use this computational platform to simulate the masking experiments conducted by Morris et al [9] and record the behavior of the modules both when our simulated sub-cortical pathway is connected or not.

A. Topographic Organization and Conditioning on a Single Stimulus

In this section we evaluate the model’s behavior in organizing and conditioning on the abstract visual stimuli being used. This evaluation is conducted in two phases. In the first phase, we train the model on the full set of inputs to determine whether each module develops an appropriate topographic organization. The data used for this training are the 266 single Gaussian activation patterns described in section II-A.

Training of the model takes place in a layered fashion so that each module is allowed to develop a topographic organization before training the modules it feeds input to. The LGN and LP modules are trained first and then their outputs are used to train the VC. Finally, the LP and VC outputs feed in, and are used to train, the AM.

A 10 by 10 lattice of neurons was used for the LGN, LP and VC modules, while a 5 by 5 lattice was used for AM. These sizes were chosen because they can adequately represent the inputs, without being too computationally expensive, with the smaller map size used for the AM corresponding with studies that attribute a lesser processing capability in the amygdala [6]. We use a Gaussian neighborhood radius and an exponential learning rate function that both decrease per epoch in a similar way as typically used for SOM (cf. [19]):

$$h(t) = r_{min} + (r_{max} - r_{min})e^{-\left(\frac{(t/t_e)^2}{2r_s^2}\right)} \quad (7)$$

$$\epsilon(t) = l_{min} + (l_{max} - l_{min})e^{-\left(\frac{(t/t_e)}{2l_s^2}\right)} \quad (8)$$

where r_{min} and r_{max} are the minimum and maximum radii for the neighborhood and r_s is the bandwidth; similarly, l_{min} , l_{max} and l_s for the learning rate. Since the values only vary per epoch, t_e defines the number of time steps per epoch (266). Values for these parameters were selected to produce stability in the organization over the required number of epochs. The same values for these parameters were used for each of the modules except the minimum neighborhood radius. Here, r_{max} was equal to the lattice width (10 or 5), r_{min} was 2 for the LGN, 3 for the LP, 1 for the VC and 4 for the AM corresponding to the differing coarseness in each module, while $r_s = 300$, $l_{max} = 0.1$, $l_{min} = 0.001$ and $l_s = 13$. Training lasted for 700 epochs, with each epoch presenting all 266 inputs in uniformly random order. This was sufficient for the modules to develop a stable topographic organization.

Fig. 4 shows the outputs from each module for stimuli (9, 1) to (9, 14), which form a strip along the middle of the input space. We observe that the data has been successfully organized topographically so that the selected strip of inputs is arranged within the map along one orientation. The spread of activation for a module for each input corresponds to the different neighborhood radii used. For example, the VC has a very narrow pattern of activation compared to the AM, in which a large radius was used. Here, the large radius for the AM was selected deliberately to ensure that topographic

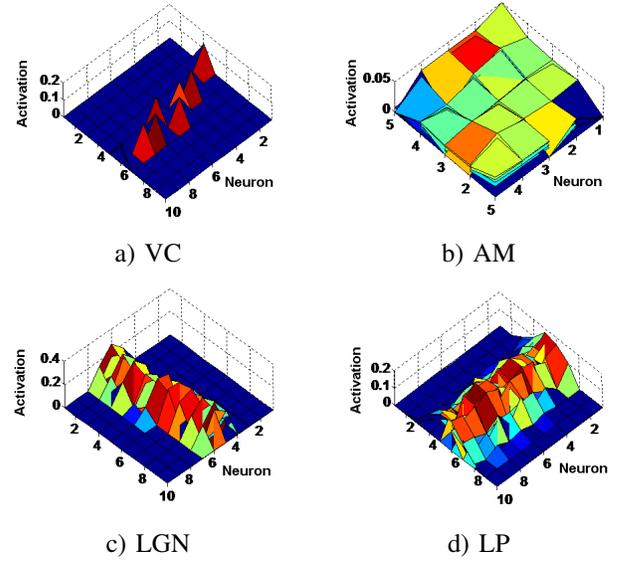


Fig. 4. Topographic representation of stimuli from (9, 1) to (9, 14) for the a) VC, b) AM, c) LGN and d) LP. Note the different orientations within the maps of the LGN and LP combined successfully in the VC.

organization did not occur, so that each neuron’s weights would be updated at each time step. This corresponds to the biological view that the amygdala itself is not topographically organized. Note also that the LGN and LP have developed different orientations for the same inputs since no such orientation preference is imposed by the learning algorithm (such as through seeding the weights). However, when these outputs are merged by the VC module, we notice that both representations have been successfully combined.

Having successfully trained the model to stability with the desired topographic organization, we now explore what effect conditioning has on learning in the maps. Starting from the pre-conditioning model, we train the model for a further 530 epochs until the pattern of activation was again stable. Here, training is performed on one input feeding through each module concurrently, rather than sequentially as before. All the parameters are the same as at the end of the pre-conditioning phase so that the radii of the winner areas have all reached their minimum values and the learning rate continues to drop with each epoch (continuing on from 700). For conditioning, we randomly chose one input as the conditioned stimulus (CS). For just this one CS input, the US was set to 1, but was 0 for all other inputs.

Once training was complete, we observed again the outputs from each map for each of the 266 inputs (Fig. 5). The CS resulted in increased activations from all three modules affected by conditioning (LP, VC and AM), but without effecting the topographic organization. This topographic representation therefore resulted in the activation of inputs close to the CS being increased. For example, Fig. 5 shows that the Gaussian activation pattern of the LP module is still preserved, while the VC activation for the conditioned stimulus (6, 7) is the highest, with similar inputs also highly activated: (5,6), (5,7), (7,7) and (7,8). The model therefore successfully

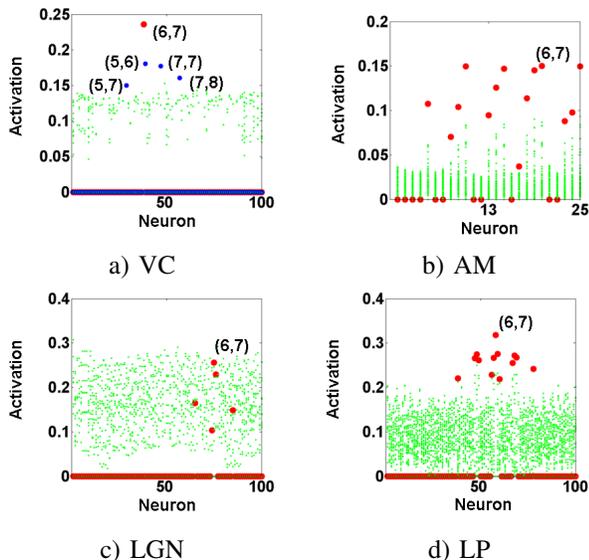


Fig. 5. Activation values from all 266 inputs for the a) VC, b) AM, c) LGN and d) LP. The red dots show the activation values of the selected CS; the blue dots show the neighboring stimuli activation values (for the VC only for comparison); the green dots show the activations of the remaining inputs. Note that for any given input only a few neurons are active per module (for example, the CS in the VC has 1 active neuron, the AM 14, the LGN 5 and the LP 13) while the rest have zero activation. The number of neurons active per input depends upon the minimum neighborhood radii selected for each module.

demonstrates that it can represent our abstract visual stimuli appropriately, while demonstrating an increased pattern of activation for an arbitrary conditioned stimulus, which is our desired pattern of behavior from the neural structures that are being represented.

B. Simulation of Masking Experiments

Having established a model that appears to have the appropriate properties we need, we now use this platform to simulate the masking experiments. These experiments were conducted on humans to help provide evidence for the existence of the visual sub-cortical pathway leading to the amygdala, and in this paper we explore our proposed model as a computational platform for such experiments. This approach offers us the advantage that we can modify the connectivity of the neural structures being simulated to test which are necessary to reproduce the required behavior, albeit constrained by the abstractions we must make to establish the model. Here, we make the assumption that our platform sufficiently models the basic behavior of the structures having already observed the ability of the model to form topographic maps and appropriately respond to conditioning.

Our computational masking experiments focus on one aspect of the Morris et al experiments [9]. Here, we are interested particularly in determining whether the model can reproduce the increased pattern of activation seen in the amygdala, but not in the visual cortex, when a target and then a mask are presented to a subject. Using our model, we can test whether the required behavior occurs both with

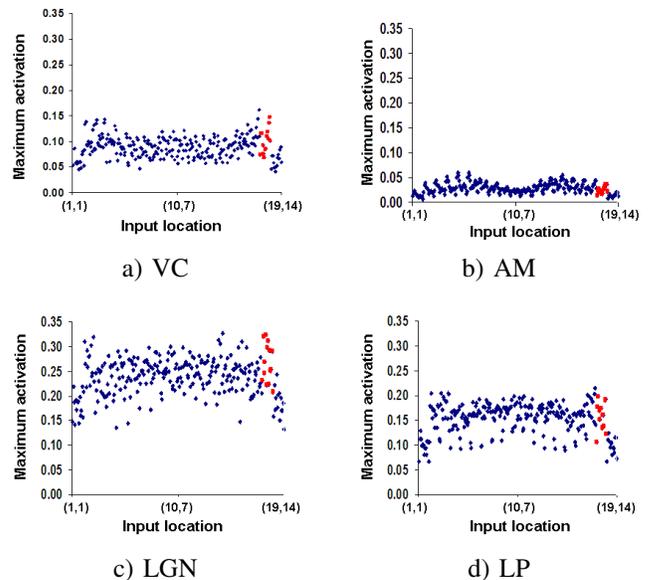


Fig. 6. The maximum activation values for each of the 266 input stimuli following habituation in the a) VC, b) AM, c) LGN and d) LP. The activations from stimuli that will be selected as CS are shown in red.

and without a direct connection from the LP to the AM – the sub-cortical pathway. To achieve this, we follow a similar pattern of training and testing conducted by Morris et al. First we go through a period of *habituation* in which the model is trained on all the inputs (in identical pairs). Second, we *condition* the model on a subset of the habituated inputs (CS) by continuing training but with the US turned on (1) for the selected CS and off (0) for all other inputs. Third, we test the model using the *masking* scenarios. Our experiments differ in scale to that of the human experiments, with an increased number of stimuli (4 for Morris et al versus 266) and an increased number of conditioned stimuli (1 versus 14). Apart from it being easier to use a larger number of stimuli with a computational model, these increases allow us to determine how robust the results are through exploring generalization over a wider range of responses.

For our CS stimuli, we choose a strip of 14 inputs (18, 1) to (18, 14) near to the edge of the visual space. This selection is deliberately near the edge so that we can pair these targets with inputs that are distinct for masking (a strip over the center of the space). Our input to the model then comprises a pair of stimuli, as described in section II-B. All of model parameters remain the same as before, including map sizes, neighborhood and learning rate parameters. During habituation, the model was trained on all 266 stimuli presented as identical pairs of inputs, and reached stability after 700 epochs. For conditioning, all 266 stimuli were again presented in identical pairs, but with the conditioning signal turned on for the CS, reaching stability after 530 epochs. To test the model, we present a mask with each of the 14 conditioned targets to the model, where each mask follows consecutively. During testing, the conditioning signal is turned off.

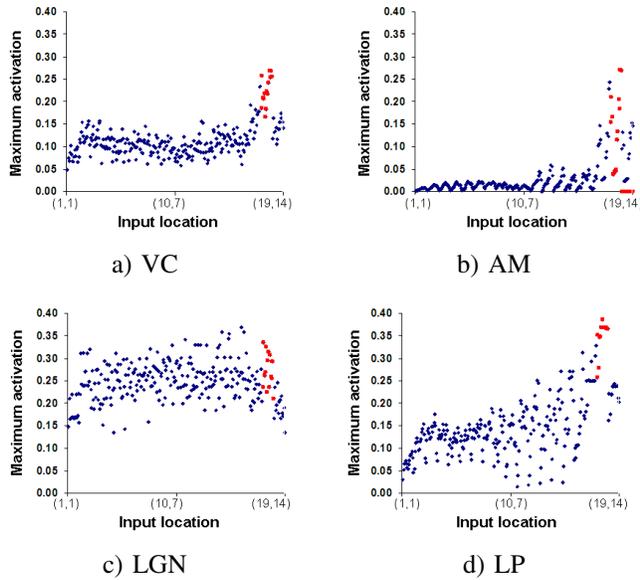


Fig. 7. The maximum activation values for each of the 266 input stimuli following conditioning on the CS in the a) VC, b) AM, c) LGN and d) LP. The activations from the CS stimuli are shown in red.

Fig. 6 shows the maximum activation value from each module for each of the inputs after habituation. These show a similar level of activation compared to the results obtained from just a single input (section III-A), despite the presentation of stimuli pairs. No significant differences in activation values are observed between the stimuli that are chosen to act as CS and the rest of the input examples in any of the four modules. In particular, the AM shows the lowest activations of all the modules with the CS stimuli showing activations within the range of 0.02 to 0.06.

Following conditioning on the selected CS stimuli, we observe that the LP, VC and AM show increased activations for the CS (Fig. 7), while the LGN activations remain at the same levels as before (no conditioning signal is applied to the LGN). Also, the breadth of the input activations affected by conditioning is much larger in the LP because of the different minimum radius applied compared to the LGN and VC in order to get a different level of map specificity. This time the CS activations of the AM have significantly increased compared to before, and now range from 0.05 to 0.29.

The results so far are consistent with the observations made when using single Gaussian activation patterns, and show that the model can adequately handle paired representations. Furthermore, we note that the model is also able to condition appropriately on multiple inputs, with all the 14 CS stimuli provoking a higher response in the LP, VC and AM.

Having established the trained model, we now simulate masking. Here we present a mask and a target as the pair of stimuli and observe the resulting maximum activation values in each module. We choose unique mask-target pairings, so that we present the pairs (10, 1) and (18, 1), through to (10, 14) and (18, 14).

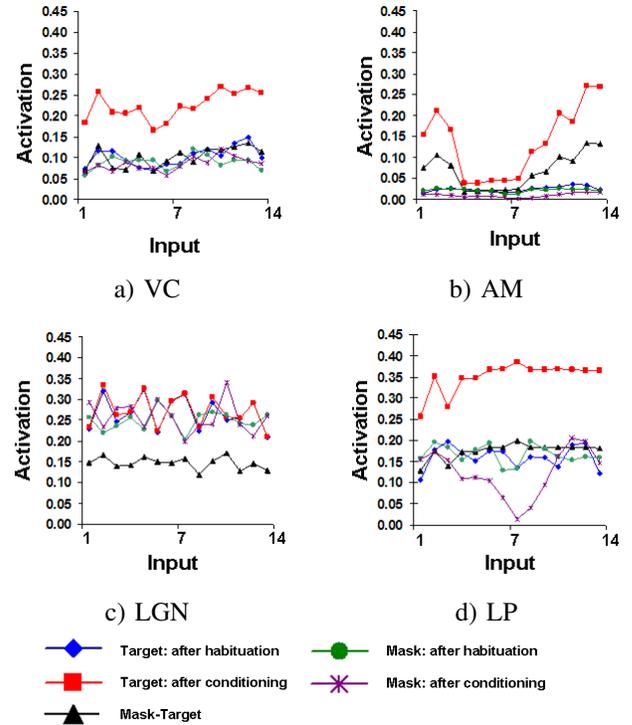


Fig. 8. The maximum activation values for the mask-target pairs compared to the habituation and conditioning maximum activations for the masks and the targets in the a) VC, b) AM, c) LGN and d) LP. The results are obtained from a model with a direct LP to AM connection.

First, we assume that there is a connection between the LP and the AM: that the sub-cortical pathway influences the amygdala. Fig. 8 shows the maximum activations in each module when presented with the mask-target pairs. Note that for comparison we show activations obtained after habituation and conditioning for both the 14 masks and the 14 targets on their own (identical pairs of inputs). In the LP, we observe that the activations for the presentation of the mask-target pair (black line) do not differ significantly from the post habituation responses to either the mask or the target, and hence are not as high as for conditioned responses to the target (red line). In the LGN, the mask-target responses are lower than all other responses; recall that the LGN does not receive the US input (is not conditioned). In the VC, the responses for the mask-target inputs are again not significantly different to the post habituation responses to either the mask or the target. Since these values are lower than the conditioned responses of the target in the VC, we note that the VC is therefore not responding as if these were conditioned responses, and hence is simulating the lack of a conscious perception of the targets. In contrast, the AM shows a mask-target response that is higher than the mask conditioned responses in the majority of cases, although lower than the activations for the conditioned target. An analysis of variance confirms these observations for the AM with the mask-target response (mean 0.0681) differing with statistical significance to the habituation mask (mean 0.0221,

$p < 0.0005$) and target (mean 0.0247, $p < 0.0009$), and conditioned mask (mean 0.0100, $p < 0.0001$), but not significantly different to the target (mean 0.0137, $p < 0.0112$). These results therefore show that the AM is responding to the target even though the VC is not.

Second, we assume that there is no connection between the LP and AM: that the sub-cortical pathway does not influence the amygdala. To achieve this we re-trained the model (habituation and conditioning) exactly as before, but without a connection from the LP to the AM. Fig. 9 shows the maximum activations in each module when presented with the mask-target pairs. Here we note no significant difference between the results for the LP, LGN and VC compared to Fig. 8, since these are connected and trained in the same way as before. However, the results for the AM differ in that they are far lower in value for all responses (with only one set of inputs from the VC, the cumulative input is lower), and crucially, do not show any significant difference to the conditioned mask response. An analysis of variance confirms these observations with the mask-target response in the AM (mean 0.0027) differing with statistical significance to the habituation mask (mean 0.0020, $p < 0.0096$) and conditioned target (mean 0.0061, $p < 0.0000002$), but not differing significantly to the habituation target (mean 0.0030, $p < 0.5569$) and conditioned mask (mean 0.0023, $p < 0.1957$). These results therefore show that the AM is not responding to the target.

With the AM responding as expected when the LP is connected to it, but not when this connection is absent, the link between the LP and AM therefore appears necessary for the required behavior to occur. This provides evidence of the influence of the sub-cortical pathway on the amygdala to match with the human observational data [9]. Although this evidence is derived from an abstract, computational model, such evidence is important in that it can be used to predict behavior and help design new psychophysics experiments, such as those conducted for auditory stimuli [20].

IV. CONCLUSION

The model presented in this paper supplements and extends the work of Armony et al [6] that computationally modeled auditory fear conditioning. It provides an abstract representation of the cortical and sub-cortical visual pathways leading to the amygdala. To achieve this we modified Armony et al's algorithm to incorporate topographic neighborhoods into the neural modules, while testing the model on its capability to condition on multiple two-dimensional stimuli.

Having defined and evaluated this model, we use it as a computational platform to simulate masking experiments. Here we include a time-delay like second input to the model to simulate the presentation of a mask immediately after the presentation of a target. This simple extension allows us to reproduce behavior observed in human studies, suggesting that this approach provides a suitable platform upon which further experiments can be conducted.

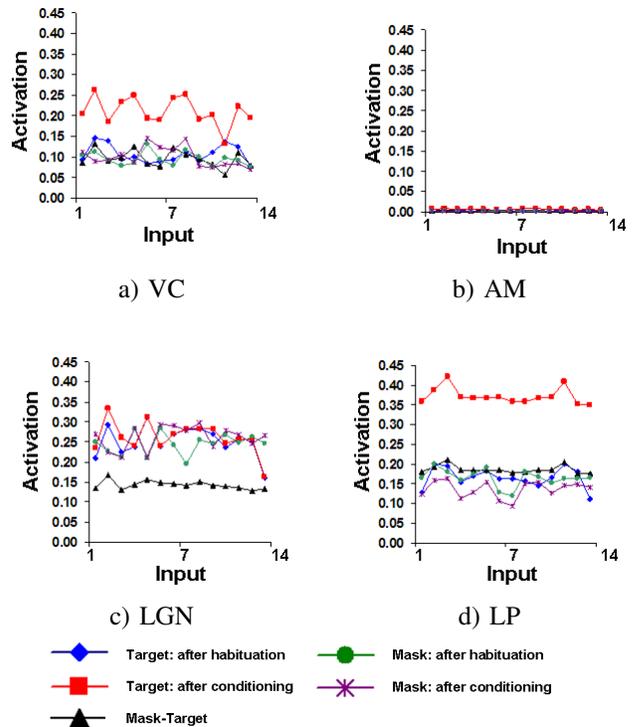


Fig. 9. The maximum activation values for the mask-target pairs compared to the habituation and conditioning maximum activations for the masks and the targets in the a) VC, b) AM, c) LGN and d) LP. The results are obtained from a model that has no direct LP to AM connection.

As a concrete example of how the platform can be used, we tested the hypothesis that the sub-cortical visual pathway is crucial in provoking a higher level of activation from the amygdala when a particular target image is processed, but not consciously perceived in the visual cortex. We evaluated the behavior of the model both with and without this sub-cortical connection and compared the results. When either the sub-cortical pathway is connected or not, the LP, VC and AM show significantly higher levels of activation in response to the CS compared to the non-conditioned stimuli, when each stimulus is presented as an identical pair of images to the model (not masking); the LGN in contrast shows no such difference because it is not conditioned. This means that the model is correctly responding to conditioning. When simulating masking, both whether the sub-cortical pathway influences the AM or not, the VC shows levels of activation that are commensurate with the target not being consciously perceived (lower levels of activation compared to the conditioned responses) demonstrating that the platform is simulating the required masking effects. However, only when the sub-cortical pathway is connected does the AM have the higher levels of activation associated with the conditioned responses when masking. Without this connection, the AM responds with low activation levels, despite being capable of responding correctly when only the target is presented without the mask. This example provides computational evidence of the effect that the sub-cortical pathway has on

the behavior of the amygdala, and therefore provides an indication that the pathway exists.

As an abstraction, our platform has two key limitations that need further work. First, the model uses a simple rate coding scheme that assumes the synchronous presentation of signals. Although our masking experiments overcome this through a time delay, a pulse coded model may offer a better way of simulating masking, provided an appropriate adaptive modular architecture could be developed that can simulate conditioning. Second, for this evaluation we have used abstract input representations that have enabled us to test the capabilities of the model to simulate conditioning and masking. Preliminary work has shown that the model can condition and classify upon presentation of more complex inputs, such as emotional face images [12], but we now need to explore more realistic data with masking. To achieve this a better understanding of the processing in the retina and thalamus is required to develop better approaches to low level visual processing. This may need us to consider feedback connections between and within modules beyond that simulated with neighborhoods.

ACKNOWLEDGMENT

We would like to thank the four anonymous reviewers for their helpful comments and suggestions.

REFERENCES

- [1] I. P. Pavlov, *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*. London: Oxford University Press, 1927.
- [2] J. E. LeDoux, "Emotion circuits in the brain," *Annual Reviews Neuroscience*, vol. 23, pp. 155–184, 2000.
- [3] —, "Emotion, memory and the brain," *Scientific American Special Edition*, vol. 12, no. 1, pp. 62–71, 2002.
- [4] K. S. Labar, J. C. Gatenby, J. C. Gore, J. E. LeDoux, and E. A. Phelps, "Human amygdala activation during conditioned fear acquisition and extinction: a mixed-trial fmri study," *Neuron*, vol. 20, pp. 937–945, 1998.
- [5] C. Shi and M. Davis, "Visual pathways involved in fear conditioning measured with fear-potentiated startle: Behavioral and anatomic studies," *The Journal of Neuroscience*, vol. 21, no. 24, pp. 9844–9855, 2001.
- [6] J. L. Armony, D. Servan-Schreiber, J. D. Cohen, and J. E. LeDoux, "An anatomically constrained neural network model of fear conditioning," *Behavioral Neuroscience*, vol. 109, no. 2, pp. 246–257, 1995.
- [7] A. Öhman, D. Lundqvist, and F. Esteves, "The face in the crowd revisited: A threat advantage with schematic stimuli," *Journal of Personality and Social Psychology*, vol. 80, no. 3, pp. 381–396, 2001.
- [8] A. Öhman, A. Flykt, and F. Esteves, "Emotion drives attention: Detecting the snake in the grass," *Journal of Experimental Psychology: General*, vol. 130, no. 3, pp. 466–478, 2001.
- [9] J. S. Morris, A. Öhman, and R. J. Dolan, "A subcortical pathway to the right amygdala mediating "unseen" fear," *Proceedings of the National Academy of Sciences*, vol. 96, pp. 1680–1685, 1999.
- [10] N. N. Doron and J. E. LeDoux, "Organization of projections to the lateral amygdala from auditory and visual areas of the thalamus in the rat," *The Journal of Comparative Neurology*, vol. 412, pp. 383–409, 1999.
- [11] L. Pessoa, "To what extent are emotional visual stimuli processed without attention and awareness," *Current Opinion in Neurobiology*, vol. 15, pp. 188–196, 2005.
- [12] A. Pavlou and M. C. Casey, "Identifying emotions using topographic conditioning maps," in *Proceedings of the International Neural Network Society - New Directions in Neural Networks Symposia (INNS-NNN 2008)*. Springer-Verlag, in press.
- [13] R. Linsker, "Self-organization in a perceptual network," *Computer*, vol. 21, no. 3, pp. 105–117, 1988.
- [14] R. Miikkulainen, J. A. Bednar, Y. Choe, and J. Sirosh, *Computational Maps in the Visual Cortex*. New York: Springer Science+Business Media, 2005.
- [15] D. O. Hebb, *The Organization of Behavior: A Neuropsychological Theory*. New York: John Wiley & Sons, 1949.
- [16] D. E. Rumelhart and D. Zipser, "Feature discovery by competitive learning," in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, D. E. Rumelhart and J. L. McClelland, Eds. MIT Press, 1986, vol. Volume 1: Foundations, book chapter 5, pp. 151–193.
- [17] C. von der Malsburg, "Self-organization of orientation sensitive cells in the striate cortex," *Kybernetik*, vol. 14, no. 2, pp. 85–100, 1973.
- [18] T. Kohonen, "Self-organized formation of topologically correct feature maps," *Biological Cybernetics*, vol. 43, pp. 59–69, 1982.
- [19] —, *Self-Organizing Maps*, 2nd ed. Berlin, Heidelberg, New York: Springer-Verlag, 1997.
- [20] J. L. Armony, D. Servan-Schreiber, L. M. Romanski, J. D. Cohen, and J. E. LeDoux, "Stimulus generalization of fear responses: Effects of auditory cortex lesions in a computational model and in rats," *Cerebral Cortex*, vol. 7, no. 2, pp. 157–165, 1997.