



Leendert A Remmelzwaal¹ (leenremm@gmail.com), Amit K Mishra¹ and George F R Ellis¹.

¹University of Cape Town, Rondebosch, Cape Town, South Africa 7700.



We present a brain-inspired cognitive architecture that incorporates sensory processing, classification, contextual prediction, and emotional tagging.

The cognitive architecture is implemented as three modular web-servers, meaning that it can be deployed centrally or across a network for servers.


Our approach is a radically different approach to taking existing software (e.g. OpenCV) and adapting it to be more brain-like - it is fundamentally brain-inspired.

It is common knowledge that bio-realistic artificial intelligence is still a distant dream to many researchers.

We embarked on this work because we believe that a modular bio-inspired cognitive architecture could be a key step in designing a bio-realistic artificially intelligent system.


Our bio-inspired cognitive architecture is modelled closely after key structures play in the cortex that have important roles in sensory processing, prediction, emotional tagging, attention direction and consciousness:

- the thalamus
- the cortex
- the arousal system, and
- the amygdala



*Key autonomic and voluntary pathways in the cortex
which we model in this paper*


The thalamus, arousal system and memory structures were each modelled as their own independent Python Flask web-servers. Each web-server could be called by any other web-server, to request information (e.g. current state of the cortex) or to request an action (e.g. classifying a new image).



Overview of the distributed system architecture

This paper confirm neuroscience understanding: that the thalamus and arousal system indeed play a crucial roles in processing incoming data, and directing cortical attention.






Without the thalamus constantly comparing incoming signals to a contextual prediction, our brains would be overwhelmed with a non-stop overload of incoming information to process.



Without salience tagging of memories and event enabled by the arousal system, the cortex wouldn't receive the guidance it requires to direct its attention to the most salient issue

Our model reveal two distinct operations of behaviour, which closely model attention in the brain.

The first mode of operation occurs when the system is presented with an input image that has a low-salience association. The cortex is engaged for one computational iteration, and returns a contextual prediction to the thalamus. For subsequent viewings with the same image the thalamus does not engage the cortex.



Two modes of operation. Low-salience (A) and High-salience (B) inputs

The second mode of operation occurs when the system experiences an input image with a high associated salience.

Similar to the first mode of operation, the thalamus notices a significant change in input, and engages the cortex to classify the image. Again, the cortex is engaged for one computational iteration, and returns a contextual prediction to the thalamus.

However this time, the cortex identifies a high salience association with the high level concept, and two things happen: firstly, the cortex triggers a release of noradrenaline from the arousal system, and secondly, the release of noradrenaline lowers the internal threshold of the thalamus system. This has a direct impact on how the system processes subsequent images.

With a lower internal threshold, the thalamus System sends images to the cortex more frequently, demanding the computational power and attention of the cortex system

The experiments reveal two distinct operations of behaviour, which closely model attention in the brain:

1. Low-salience mode (single engagement of the cortex)
2. High-salience mode (ongoing engagement of the cortex)

In this paper we have demonstrated with the aid of a fully-connected simulation of the thalamus, arousal system and cortex, how these three key structures rely on each other to allow for:

1. Attention direction
2. Saliency detection
3. A simple form of prediction

We present a cognitive architecture specifically with future research in mind. The distributed modular nature of the architecture we present here allows research groups to independently enhance and extend a single module in the system.

In this paper we have presented an open-source brain-inspired distributed cognitive architecture.

The model is:

1. Extensible
2. Optimizing performance of complex tasks
3. Confirming neuroscience understanding of the cortex

The source code as well as records of the tests conducted in this paper are publicly available online.

- https://bitbucket.org/leenremm/affect_in_the_cortex

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.