

RFR: An Actor-Critic Decision-Making Model with the Frontal-Cortex-Basal-Ganglia Loop

Proposed by [Project AGI](#) and [the Whole Brain Architecture Initiative](#)

Summary

Background and motivation

As intelligent agents make decisions, any project aiming to realize human-like AGI should model decision-making. As we have been pursuing [the WBA approach](#) to create AGI by learning from the architecture of the entire brain, we request you to model the decision-making of the mammalian brain. While a number of models have been proposed, we refer to O'Reilly's model on [his textbook for computational cognitive neuroscience](#) (CCNBook hereafter) as the standard, where decisions are supposed to be made with [the loop consisting of the frontal cortex, basal ganglia, and related areas](#), which reinforces decisions in an [actor-critic](#) way.

Objective

You are requested to implement a biologically-plausible yet computationally effective model for decision-making and action selection. The model should serve as a reference model for other brain-inspired models of intelligence. Thus, its implementation should be as **simple** as possible for being used and maintained in the community. We outline the structure of the model to be implemented in [the Detailed Project Description section](#) below.

Success criteria

The implementation will be judged with the following criteria:

- Biological plausibility
The implementation should be 'compatible' with the structure and function of the mammalian brain.
- Usability
The implementation should be easily used and maintained together with documentation.
- **Specifications**
See [the Detailed Project Description section](#) below.
- Performance
Use one or more tasks from the [Dataset and Test](#) section below.

Detailed Project Description

The request is to implement a decision-making model consisting of the following modules: (They are [refactored](#) from the model in the CCNBook.)

- FC Module
 - provides with options based on input from outside
 - may or may not implement a winner-take-all logic for decision-making
 - may or may not implement an accumulator logic to accumulate scores for options
 - may use recurrent networks
- Actor module
 - corresponds to part of BG and the thalamus
 - modulates the strength of each option
 - receives the [TD signal](#) from the Critic to learn
 - receives the *state* input from outside
 - may or may not implement a winner-take-all logic for decision-making
- Critic module
 - corresponds to part of BG and the amygdala
 - creates the TD signal based on the exterior reward
 - receives the *state* input from outside

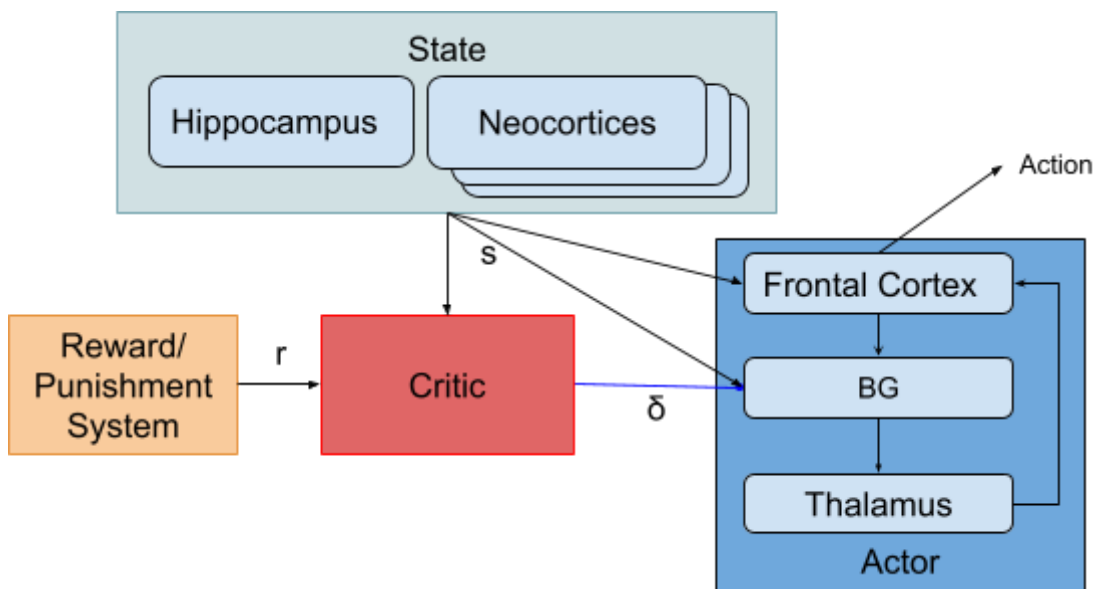


Figure 1

Both Actor and Critic contain parts of BG.

Model characteristics:

The model should have the following performance characteristics (see [here](#) for a discussion of these constraints):

- Few hyperparameters: These should either adapt automatically to model conditions, or not vary over time.
- Possibility to select 1 or more actions simultaneously, if necessary
- Conflict resolution: A method of excluding incompatible action selections
- Clean switching: Marginally better actions should be selected quickly and definitively without vacillation or dithering.
- Full selection: To avoid pathological interference, unselected actions should be fully inactive.

Computational Tools for Implementation:

You may use any well supported machine learning algorithms or open source frameworks for implementation. You may prefer more biologically plausible algorithms such as those provided by [emergent](#) from O'Reilly's group (i.e., Hebbian learning). Note, however, your new implementation should have more usability than the original implementation in *emergent*.

We recommend Python for programming.

We request you to make modules **well 'encapsulated' so that modules could be reused** without large modifications and used in a hybrid framework environment (e.g., mixing implementation with TensorFlow and Caffe) for better interoperability among developer teams. For that matter, WBAI has been developing our own framework for brain-inspired computing ([BriCA](#)).

We invite you to discuss the tool issue with us before you start working on the request.

Dataset and Tests

While currently we do not offer our own dataset or test batteries, the [Executive Function](#) chapter of CCNBook refers to tests for decision-making models. You can choose one or more tasks here to test your implementation. Note that most of them require working memory so that they test working memory as well.

- Task switching tasks (psytoolkit.org):
 - [Wisconsin card sorting task](#) (WCST)
 - [Dimensional change card sorting task](#) (DCCS)

- [N-Back task](#) ([Wikipedia](#))
- Stroop task ([CCNBook](#)) / ([Psytoolkit.org](#)) ([Wikipedia](#))
- [Store-Ignore-Recall \(SIR\) task](#) (CCNBook)
- Reproducing the [A Not B](#) error (CCNBook) ([Wikipedia](#))

Background Information

The following chapters of CCNBook are relevant to this RFR.

- [Motor Control and Reinforcement Learning](#) in CCNBook ([our summary](#))
- [Executive Function](#) in CCNBook ([our summary](#))

Note that O'Reilly's model is also a model of working memory ([PBWM: the Prefrontal cortex Basal ganglia Working Memory model](#)). Also note that the model presented here is a quite simplistic one for a starter. The accumulator/ramping aspect of decision making would have to be also taken into consideration as in [[Simen 2012](#)].

See also [the Action Selection article on Scholoarpedia](#).