A Biologically-Inspired Neural Implementation of Affect Control Theory

Aarti Malhotra (aarti.malhotra@uwaterloo.ca)
Cheriton School of Computer Science, University of Waterloo,
200 University Ave. W, Waterloo, ON, N2L 3G1, Canada

Terrence C. Stewart (terrence.stewart@nrc-cnrc.gc.ca)
National Research Council of Canada, University of Waterloo Collaboration Centre
200 University Ave. W, Waterloo, ON, N2L 3G1, Canada

Jesse Hoey (jhoey@cs.uwaterloo.ca)
Cheriton School of Computer Science, University of Waterloo,
200 University Ave. W, Waterloo, ON, N2L 3G1, Canada

Abstract

Social interactions are a part of day-to-day life of most human beings. Affect, decision-making and behavior are central to it. With increase in adaptation of technology in our society, interaction between humans and artificially intelligent agents is also increasing. Large-scale brain-inspired neural models have been equipped with capabilities to fulfil a variety of tasks, but there has been relatively limited focus on making them capable of handling social interaction. In this paper, NeuroACT, a neural computational model and implementation of a socio-psychological theory called Affect Control Theory (ACT) is presented. This is towards building an emotionally intelligent AI agent, that can handle interactions. It takes as input a continuous affective interpretation of a perceived event, consisting of an actor, behavior and an object and generates post-event predictions of the next optimal behavior to minimize deflection. The aim is to model the role of affect guiding decision-making in AI agents, resulting in interactions that are similar to human interactions, while inhibiting some behaviors based on the social context.

Keywords: Social interaction; Affect Control Theory; Brain-inspired simulation; Nengo

Introduction

Social interaction is central to the human experience; affect and emotion plays an important role in guiding human thinking and behavior. Affect is a property of consciousness (Barrett & Satpute, 2019) and a part of every psychological phenomenon, even those that are not explicitly emotional (Hutchinson & Barrett, 2019). Social neuroscience has primarily focused on sense of self identity and how a person's mind creates a perception of another person, whereas affective neuroscience concerns mainly with brain basis of emotions (Barrett & Satpute, 2013). We present a neural model and implementation of social interaction in an AI system using a socio-psychological theory called Affect Control Theory (ACT), and combine social and affective domain perspectives, thereby dissolving their artificial boundaries. This is towards gaining deeper understanding of the neural mechanisms of the role of affect guiding decision-making. The goal here is to show that the calculations required by ACT can be implemented by spiking neurons, using anatomical structure that fits the cortex, basal ganglia, and thalamus.

Affect Control Theory (ACT) is a comprehensive social psychological theory of human interaction that emphasises feelings (affect) as a key factor (Heise, 1987; N. J. MacKinnon, 1994; Heise, 2007; N. MacKinnon & Robinson, 2014). ACT proposes that social perceptions, actions, and emotional experiences of the people are governed by a psychological need to minimize deflections between culturally shared fundamental sentiments about social situations and transient impressions resulting from the dynamic behaviors of interactants in those situations. Sentiments are said to have three aspects in an affective space, forming a three-dimensional vector comprising of Evaluation, Potency, and Activity (EPA) (Osgood, Suci, & Tannenbaum, 1957; Heise, 2007). Evaluation concerns goodness versus badness, Potency concerns powerfulness versus powerlessness, and Activity concerns liveliness versus quietness. The value of each aspect can vary in degree, in the sense that it can be greater or lesser, in either a positive or negative direction. The dimensions of EPA vector is a common cross-cultural representation of social objects, such as interactants' concepts of situations, emotions, identities and behaviors, hypothesized to be an organizing principle of human socio-emotional experience (Osgood, May, & Miron, 1975).

EPA profiles of concepts can be measured with the semantic differential, a survey technique where respondents, both males and females rate affective meanings of concepts on numerical scales. These numerically-measured sentiments are useful for mathematical analysis. EPA measurements [as noted in (Heise, 2001)] are appropriate when one is interested in affective meanings rather than denotative or logical meanings. Affective meanings correspond to sentiments - that is, the general feelings that we have about something. The EPA system is notable for being a multivariate approach to measuring affect, as compared, say, to attitude measurement which deals only with the single dimension of evaluation.

Affect control theorists have compiled datasets of a few thousand words along with average EPA ratings obtained from survey participants who are knowledgeable about their culture (Heise, 2010) [1]. For example, most English speakers agree that professors are about as nice as students (E), however more powerful (P) and less active (A). The corre-

[1] The EPA profiles in this paper are from 'Indiana 2002-4' dataset

sponding EPA profiles are {1.61, 1.58, 0.35} for professor and {1.49, 0.31, 0.75} for student. The values range by convention from -4.3 to +4.3 (Heise, 2010). In general, within-cultural agreement about EPA meanings of social concepts is high even across subgroups of society, and cultural-average EPA ratings from as little as a few dozen survey participants have shown to be extremely stable over extended periods of time (Heise, 2010).

Large-scale brain-inspired neural models like Spaun (Eliasmith et al., 2012) implement some mechanisms of cognitive processing to perform variety of tasks. A neural model called POEM (POinter EMotions) (Kajic, Schröder, Stewart, & Thagard, 2019) provides a detailed account of neurocomputational mechanisms responsible for psychological functions required for emotions. Implementing social interaction mechanisms in an AI system still remains a challenge.

In this paper, NeuroACT, a neural model of social interaction based on ACT is presented along with its implementation using Nengo, a python library for building and simulating large-scale neural models, showing the possibility of developing a human-like AI agent, which can interact with human or other agents. In this paper, the mathematical model of ACT and its non-neural implementation is first outlined. Then details about NeuroACT model, its implementation using spiking neurons and its simulation in a doctor-patient scenario and prisoner's dilemma game play are provided. Lastly, conclusion and some directions for future work are presented.

Mathematical model of ACT

In an interaction, actor is the agent who behaves (or acts) towards a target object (who can be self or some other person). An actor or an object can be specified by identity-nouns, whereas behaviors are specified by verbs.

Interaction Event Description

ACT models the formation of transient impressions from events with a minimalist grammar of the form Actor Identity-Behavior-Object Identity (*ABO*). Each of these have an EPA profile. A person's basic identity can be particularized with specifications of emotion, traits, moods, biological characteristics, statuses, or moral dispositions. In our current version of neural model, emotion is considered as the modifier. The state of being has generally more impact on the outcome impression than the identity does. [Modifier(*Mod*)-Identity(*I*) combinations were analyzed in (Smith-Lovin & Heise, 1988)]. The interaction can be written as

$$Event = [Mod_{actor}][I_{actor}][B_{actor}][Mod_{object}][I_{object}] \quad (1)$$

Note that each of these elements is an EPA profile.

Modifier-Identity Combination

The addition of attributes or adjectives that modify the identities (e.g., "good friend" or "abusive father") is calculated from the EPA values of both the identity and the modifiers. The EPA profiles of particular modifiers are symbolized as $P = \{Pe, Pp, Pa\}$, identities as $R = \{Re, Rp, Ra\}$, and the profile for identity-modifiers amalgamation as $C = \{Ce, Cp, Ca\}$. The modifier-identity profile is computed by the equation given below:

$$C = pP + rR + a \quad (2)$$

where p and r are coefficients estimated from empirical studies of the modifiers and identity, respectively and a is a constant. For example, using the affective dictionary, the EPA profile of a "stranger" is {0.02 -0.09 -0.23}, the EPA profile of "happy" is {2.92 2.43 1.96}, and that of "happy stranger" is {0.6 0.5 0.5}.

The interaction in (1) can be re-written as follows:

$$Event = [C_{actor}][B_{actor}][C_{object}] \quad (3)$$

For the rest of the paper, actor A means C_{actor} and object O means C_{object}

Deflection

According to ACT grammar, the fundamental sentiment f (represented by over-bar) is represented as follows:

$$f = \{\bar{A_e}\ \bar{A_p}\ \bar{A_a}\ \bar{B_e}\ \bar{B_p}\ \bar{B_a}\ \bar{O_e}\ \bar{O_p}\ \bar{O_a}\} \quad (4)$$

and the transient impression τ (represented by caret) evoked by an event is given by:

$$\tau = \{\hat{A_e}\ \hat{A_p}\ \hat{A_a}\ \hat{B_e}\ \hat{B_p}\ \hat{B_a}\ \hat{O_e}\ \hat{O_p}\ \hat{O_a}\} \quad (5)$$

In ACT, the weighted sum of squared Euclidean distances between fundamental sentiments and transient impressions is called total deflection D:

$$D = (f - \tau)^2 \quad (6)$$

Calculation of τ will be discussed in the next subsection. Deflection arises when impressions produced by an event differ from sentiments. Deflection that cannot be resolved produces psychological stress, which is a serious condition that can undermine one's health. Deflection is related to unlikelihood: the more deflection an event produces, the more that event seems stranger, more surprising, more unique and even inconceivable.

Consider for example, a professor who yells at a student. Most observers would agree that this professor appears considerably less nice (E), a bit less potent (P), and certainly more active (A) than the cultural average fundamentals of a professor. ACT treats the dynamics of emotional states and behaviors as continuous trajectories in affective space. Deflection minimisation is the only prescribed mechanism.

Transient impression formation The transients existing after an event can be predicted from the transients that precede the event by the equation given below:

$$\tau = Mt \quad (7)$$

M is the matrix of prediction coefficients estimated in impression-formation research, with one column for each post-event transient being predicted. For example, Matrix M is 20 x 9, consisting of coefficients estimated from U.S male data on ABO. Vector t contains pre-event transients along with interaction terms that have been found to have predictive value in empirical analyses. Vector t given below is 1 x 20, hence τ is 1 x 9. [Refer to (Heise, 2007)]

$$
\begin{aligned}
t = \{1\ \bar{A_e}\ \bar{A_p}\ \bar{A_a}\ \bar{B_e}\ \bar{B_p}\ \bar{B_a}\ \bar{O_e}\ \bar{O_p}\ \bar{O_a} \\
\bar{A_e}\bar{B_e}\ \bar{A_e}\bar{O_p}\ \bar{A_p}\bar{B_p}\ \bar{A_a}\bar{B_a}\ \bar{B_e}\bar{O_e} \\
\bar{B_e}\bar{O_p}\ \bar{B_p}\bar{O_e}\ \bar{B_p}\bar{O_p}\ \bar{A_e}\bar{B_e}\bar{O_e}\ \bar{A_e}\bar{B_e}\bar{O_p}\}
\end{aligned}
\tag{8}
$$

To show an example of how M and t affects the calculation of τ, the following equation shows the post-event Actor's evaluation dimension estimated using the impression equations (considering non-zero values of first column of M which related to $\hat{A_e}$):

$$
\begin{aligned}
\hat{A_e} = -0.26 + 0.41\bar{A_e} + 0.42\bar{B_e} - 0.02\bar{B_p} - 0.10\bar{B_a} + 0.03\bar{O_e} \\
+0.06\bar{O_p} + 0.05\bar{A_e}\bar{B_e} + 0.03\bar{A_e}\bar{O_p} + 0.12\bar{B_e}\bar{O_e} \\
-0.05\bar{B_e}\bar{O_p} - 0.05\bar{B_p}\bar{O_e} + 0.03\bar{A_e}\bar{B_e}\bar{O_e} - 0.02\bar{A_e}\bar{B_e}\bar{O_p}
\end{aligned}
\tag{9}
$$

The coefficients in the above equation indicate the factors and the degree to which they contribute towards the post-event evaluation of the actor. For example, the positive coefficient on pre-event evaluation of actor $\bar{A_e}$, means that the good actors are evaluated more positively (in E) and bad actors are evaluated more negatively (in E), with a factor of 0.41. The positive coefficient on combination terms like pre-event behavior and object evaluation $\bar{B_e}\bar{O_e}$ means that the actors are evaluated more positively (in E) if they are observed doing good things to good people, or bad things to bad people, but more negatively (in E) if they are observed doing bad things to good people or good things to bad people, with a factor of 0.12. Similarly, the other dimensions can be calculated for $\hat{A}$, $\hat{B}$ and $\hat{O}$ giving us the value of τ as mentioned in eq. (5).

Optimal Behavior

Action selection in an interaction would be based on any institutionally-appropriate, feasible, and sentiment-affirming behavior. For example, in a medical setting, there would be a doctor-patient interaction, where doctor's identity is generally considered as quite good and potent and somewhat active with an EPA profile as {1.90,0.69,0.05}, whereas a patient identity is considered a bit good, less powerful and quite weak with an EPA profile as {0.90,-0.69,-1.05}. The sentiment-affirming behavior for a doctor would be to treat, instruct etc to the patient, so that his impression is maintained as good. If he does acts of yelling, cruelty etc, his impression will be bad and will cause deflection and conflict.

An event seems more unlikely, uncanny, or unique as deflections (D) are larger. In ACT, the EPA profile for the optimal behavior is regarded as the one that minimizes the unlikeliness of an event, that is defined as below.

$$k + \sum_{i=A_e}^{O_e} w_i D_i \tag{10}$$

From eq. (10) and (6), we have

$$k + \sum_{i=A_e}^{O_e} w_i (f_i - \tau_i)^2 \tag{11}$$

where k is a constant and w stands for summation weights. Minimizing unlikeliness or maximizing normality is obtained by setting partial derivatives of the right side of the above equation to zero and solving for behavior terms, giving us the suggested optimal behavior [for details refer (Heise, 2007)].

Predicted Emotion and Identity

ACT predicts the emotion and identity of both the actor and the object post interaction, which can also affect the dynamics. In this paper only optimal behavior will be focused upon.

Non-Neural implementations

Interact A computer software tool named Interact, implements ACT's mathematical model in Java. It provides a user interface to setup the interactions and analyze the results. It has a dictionary of various datasets across six nations, ranging from 1977 to 2007, and consists of EPA profile ratings for identities, behaviors, modifiers rated by male and female raters, which is useful in cross-cultural and historical analysis. [New datasets can be found at https://research.franklin.uga.edu/act].

BayesAct BayesAct (Hoey, Schröder, & Alhothali, 2016; Schröder, Hoey, & Rogers, 2016) generalises ACT by maintaining multiple hypothesis of behaviors and identities simultaneously as a probability distribution. It uses partially observable Markov decision process (POMDP). [Some applications include (Lin et al., 2014; Jung, Hoey, Morgan, Schröder, & Wolf, 2016)].
BayesAct and Interact can be accessed at http://bayesact.ca

Neural Model

The novel contribution of this paper is to take the underlying mathematics of ACT and implement them using the spiking neurons. In particular, it is striking that the overall form of the theory maps very well onto a neural model of the cortex/basal ganglia/thalamus loop that has been previously used to model a variety of tasks (Eliasmith et al., 2012).

The core part of the algorithm that is modelled here and its relation to the neural model of the brain is shown in Figure 1. In this work, the mechanisms for maintaining and tracking the EPA values of the current situation is not modelled; rather, focus is on the calculation of deflection and hence unlikeliness, given the event perception from an object's (AI agent) perspective and time t. That is, given the EPA values of the current situation, the question is: what action should be performed by the object of the event?

This maps well onto the traditional roles of the cortex, basal ganglia, and thalamus. Neurons in the cortex (1 in Figure 1) will represent the EPA values, the connections between cortical neurons and basal ganglia neurons (2) will compute eq. (11), the basal ganglia (3) will find the action with the largest deflection minimizing utility value, and the thalamus (4) will activate that particular action.

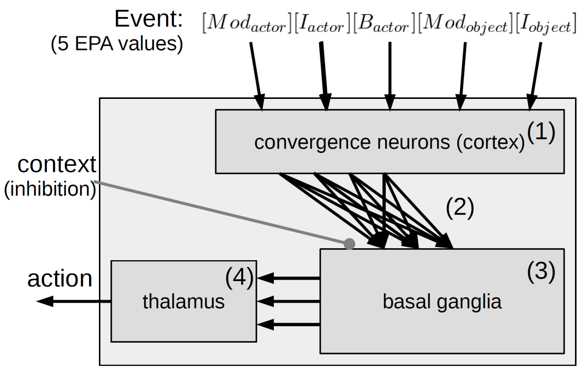

Figure 1: Neural implementation of ACT

While the overall mathematical function of this system is easy to describe and implement, it will be shown how spiking neurons can perform these operations. In particular, here Neural Engineering Framework (NEF; Eliasmith & Anderson, 2004) is used, which is a general method for finding how to connect simulated neurons so as to get the best approximation of any given algorithm. In general, the idea here is that the activity of groups of neurons can be thought of as representing vectors, and the connections between groups of neurons can be thought of as computing functions on those vectors. If we know the set of functions that we want to compute then we can perform a sequence of local optimizations (one for each set of connection weights) that will find the best approximation of the algorithm, given whatever type of neurons we want to use (including spiking and non-spiking neuron models).

For the basal ganglia and thalamus, we can make use of already-existing models of how to use the NEF to implement exactly the function that is desired here: a system that takes in a set of values from eq. (11) and determines which one is the largest utility, say U, outputting that information to the thalamus. This has been previously shown to both map on well to the anatomy of the basal ganglia and to exhibit realistic reaction times (Stewart, Choo, & Eliasmith, 2010). This system has been used in many previous models, including models of the bandit task (Stewart, Bekolay, & Eliasmith, 2012) and the large functional brain model Spaun (Eliasmith et al., 2012). The same is used here without adjusting any parameters. Also, an inhibitory "context" input that provides a large negative value for any actions that should not currently be considered.

While the basal ganglia and thalamus model take care of computing which of the action values has the largest deflection minimizing utility U (i.e. which action should be taken), this still leaves the question of how to have neurons calculate the eq. (11) values for each action, given the basic EPA values that constitute t.

Since this is simply a function, it is possible to train a neural network to approximate that function. However, the general challenge of neural networks is that if the function being approximated is too complicated, we will need a very large neural network to do this (either very deep or very broad, or both). Importantly, the networks generated using the Neural Engineering Framework have been analyzed in terms of the class of functions that they are good at approximating when using a Leaky Integrate-and-Fire neuron model with the default distribution of tuning curves (Eliasmith & Anderson, 2004). This analysis indicates that these neurons are best at approximating functions that consist of linear combinations of low-degree polynomials. Crucially, this is *exactly* the form of the calculation being done here (see eq. (9)). This means that we can use small numbers of neurons (here we use 1500) with the same parameter settings as has been used in the other biological models to approximate this function.

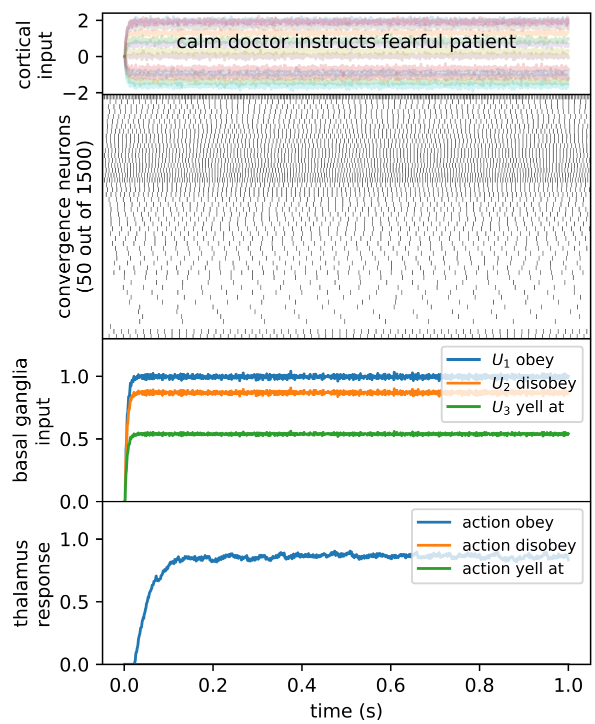

Figure 2: Example of behavior of NeuroACT

An example of the overall behavior of the resulting model is shown in Figure 2. The input is the EPA values for each of the 5 relevant terms. In this case, the situation is

$$[calm][doctor][instructs][fearful][patient]$$

and the corresponding input EPA values are [1.97 1.32 -1.4][1.9 0.69 0.05][1.85 1.65 0.3][-1.64 -0.94 -1.15][0.9 -0.69 -1.05]. These values are fed into the convergence neurons. These connections are completely random, meaning that any particular input will produce some random pattern of neural activity that is unique to that input. From that activity, the connection weights from the convergence neurons to the basal ganglia compute the eq. (11) function for all of the different actions in parallel. For simplicity, here we only plot three of those actions: 'obey', 'disobey' and 'yell at'. Finally, the basal ganglia model finds the largest of these activity values (i.e. 'obey') and directs that result to the thalamus, so the object of the event, which is the patient in this case, can perform for better interaction. This is also the optimal behavior according to the mathematical model.

Simulation

To simulate NeuroACT model for social interaction involving affect, decision-making and behavior, a single play of prisoners dilemma game scenario was used. Of the two players involved; one represents a simulated human player agent (Actor) and the other represents NeuroACT AI agent (Object). In the play round, each player can decide to either give two coins to the other player (cooperation strategy) or take one coin (defection strategy) from a common pile. Players can maximize their returns by defecting while their partner cooperated, and although the Nash equilibrium is mutual defection, the players can jointly maximize their scores through mutual cooperation.

In the simulation scenarios, the AI agent perceives the emotional state, identity and behavior of the human agent, and outputs the optimal behavior it would choose ('give' or 'take') based on the ACT prescription of deflection-minimization. The decision-making dynamics over the time scale are tested, such that if the perceived emotion of the human agent changes during the play round, the AI agent changes its strategy as well. For simplicity, the identity of both the players was kept as 'stranger'. EPA profiles used for identity, modifiers and behaviors are as below:

$$[happy] : [2.92, 2.43, 1.96]$$

$$[angry] : [-1.45, -0.30, 1.13]$$

$$[stranger] : [0.02, -0.09, -0.23]$$

$$[gives\ to] \ : [1.60, 1.47, 1.55]$$

$$[takes\ from] \ : [-1.40, 1.62, 1.50]$$

Inhibition: The dictionary consists of 500 behaviors, out of which 498 are inhibited in this case due to the game context. If there was no mechanism of inhibitory neurons, AI agent would have selected a deflection-minimizing behavior out of 500 options, but in our case, it selects between 'give' or 'take' behavior only and others get inhibited.

To demonstrate the behavior of the model and show its ability to use neurons to perform similar calculations as found in the standard Affect Control Theory, we provide cortical input of 5 sets of EPA values representing a particular situation. Since neurons require time to respond, we hold this input constant for 0.5 seconds and then present a new situation. In particular, we manually adjust the recognized emotion from 'happy' to 'angry', as this causes ACT to produce a different action. It should be noted that, in this example, the 'object' is meant to correspond to the NeuroACT AI agent itself.

Scenario 1: Human agent cooperates with AI agent

Perception at time $t \leq 0.5$:

$$[happy][stranger][\textbf{gives to}][happy][stranger]$$

Perception at time $t > 0.5$:

$$[\textbf{angry}][stranger][\textbf{gives to}][happy][stranger]$$

Scenario 2: Human agent defects with AI agent

Perception at time $t \leq 0.5$:

$$[happy][stranger][\textbf{takes from}][happy][stranger]$$

Perception at time $t > 0.5$:

$$[\textbf{angry}][stranger][\textbf{takes from}][happy][stranger]$$

Results

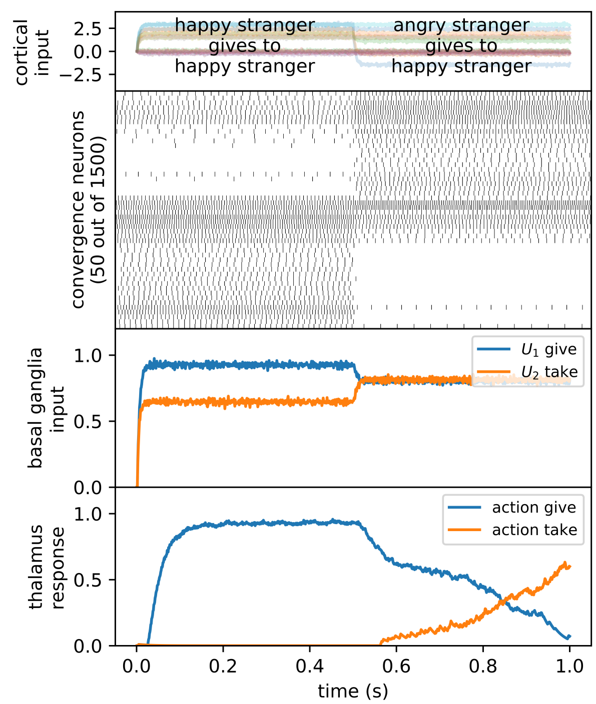

Figure 3: Human agent cooperates with AI agent

Results for the simulation runs for Scenarios 1 and 2 are shown in Fig 3 and 4 respectively. In both scenarios, the resultant behavior changes from 'give' to 'take' on perceiving the emotion of the human agent that changes from 'happy' to 'angry', given the affective dynamics. In scenario 1 (Fig 3), the change in behavior seems slower and more deliberate than

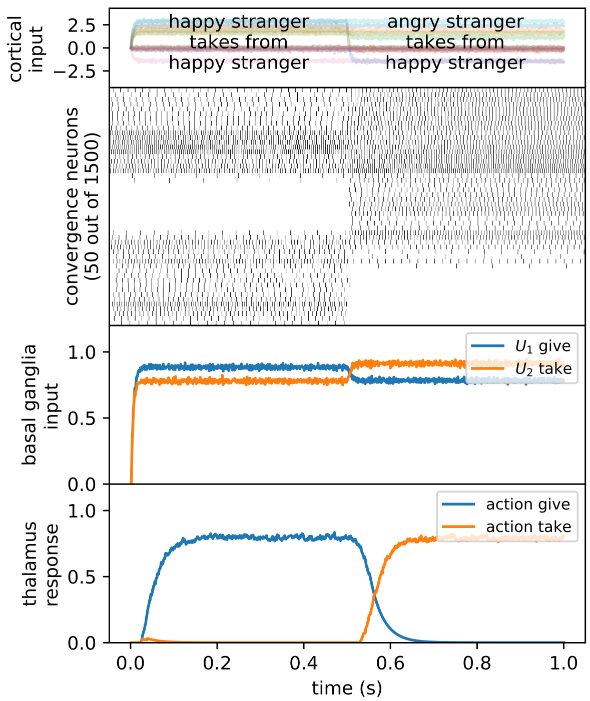

Figure 4: Human agent defects with AI agent

in scenario 2 (Fig 4), where the change is faster and somewhat automatic. This may be due to actor's behavior being more positive in scenario 1 as compared to scenario 2.

NeuroACT shows how affect influences decision-making and behavior. The behavior chosen by the model matches with its non-neural counterpart in choosing the optimal behavior as prescribed by ACT. The ability of the neural model to handle time dimension is important for the temporal order of information processing similar to human brain circuit (Gupta & Merchant, 2017).

Conclusion and Future Work

Social interaction is a challenging area to replicate in brain simulations. NeuroACT is a novel contribution implementing affective social interaction in spiking neurons. It is a generalized and extensible neural model of ACT, capable of providing an AI agent with the ability to interact with the other AI agents or humans. The input is an interaction perception and output is an optimal behavior selection. This is a step towards making emotionally intelligent AI agents.

A specific doctor-patient interaction is tested for the model. Simulation of a single play in Prisoner's dilemma game is provided. This can be iterated as well, taking into account that in the next round of play, the actor and the object change. NeuroACT can be used to model any other interaction. Future enhancement can include settings for additional context, such as location. The model can be expanded using similar methods to predict the emotion and generate re-identification of the actor and the object post-interaction. This system can be enhanced by incorporating some sensorimotor signals to integrate with physical robots.

Some other improvements could be considered involving a working memory component for the agent to utilize experience from the previous interactions. The input to the model is a generic input, that can incorporate visual, textual, or auditory forms, as all would eventually translate into verbal concepts. Advances in neuroimaging techniques like hyperscanning to study the inter-brain synchronisation (Liu et al., 2018) in social interaction may give more insight into the neural mechanisms at play.

References

Barrett, L. F., & Satpute, A. B. (2013). Large-scale brain networks in affective and social neuroscience: towards an integrative functional architecture of the brain. *Current Opinion in Neurobiology*, *23*, 361-372.

Barrett, L. F., & Satpute, A. B. (2019). Historical pitfalls and new directions in the neuroscience of emotion. *Neuroscience Letters*, *693*, 9-18.

Eliasmith, C., & Anderson, C. H. (2004). Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems. *IEEE Transactions on Neural Networks*, *15*, 528-529.

Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., & Rasmussen, D. (2012). A Large-Scale Model of the Functioning Brain. *Science*, *338*(6111), 1202–1205.

Gupta, D. S., & Merchant, H. (2017). Understanding the Role of the Time Dimension in the Brain Information Processing. *Frontiers in Psychology*, *8*, 240.

Heise, D. R. (1987). Affect control theory: Concepts and model. *The Journal of Mathematical Sociology*, *13*(1-2), 1–33.

Heise, D. R. (2001). Project Magellan: Collecting Cross-cultural Affective Meanings Via The Internet..

Heise, D. R. (2007). *Expressive order: Confirming sentiments in social actions*. Springer US.

Heise, D. R. (2010). *Surveying Cultures: Discovering Shared Conceptions and Sentiments*. John Wiley & Sons, Inc.

Hoey, J., Schröder, T., & Alhothali, A. (2016). Affect Control Processes: Intelligent Affective Interaction using a Partially Observable Markov Decision Process. *Artificial Intelligence*, *230*, 134–172.

Hutchinson, J. B., & Barrett, L. F. (2019). The Power of Predictions: An Emerging Paradigm for Psychological Research. *Current Directions in Psychological Science*, *28*, 280 - 291.

Jung, J. D. A., Hoey, J., Morgan, J. H., Schröder, T., & Wolf, I. (2016). Grounding Social Interaction with Affective Intelligence. In *Canadian Conference on AI.*

Kajic, I., Schröder, T., Stewart, T. C., & Thagard, P. (2019). The semantic pointer theory of emotion: Integrating physiology, appraisal, and construction. *Cognitive Systems Research*, *58*, 35-53.

Lin, L., Czarnuch, S., Malhotra, A., Yu, L., Schröder, T., & Hoey, J. (2014). Affectively Aligned Cognitive Assistance Using Bayesian Affect Control Theory. In *IWAAL*.

Liu, D., Liu, S. Z., Liu, X., Zhang, C. Q., Li, A., Jin, C., … Zhang, X. (2018). Interactive Brain Activity: Review and Progress on EEG-Based Hyperscanning in Social Interactions. *Frontiers in Psychology*, *9*, 1862.

MacKinnon, N., & Robinson, D. (2014). Back to the Future: 25 Years of Research in Affect Control Theory. *Advances in Group Processes*, *31*, 139–173.

MacKinnon, N. J. (1994). *Symbolic interactionism as affect control.* Suny Press. (Pages: xvi, 245)

Osgood, C. E., May, W. H., & Miron, M. S. (1975). *Cross-cultural universals of affective meaning.* University of Illinois Press. (Pages: xix, 486)

Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1957). *The measurement of meaning*. University of Illinois press.

Schröder, T., Hoey, J., & Rogers, K. B. (2016). Modeling dynamic identities and uncertainty in social interactions: Bayesian affect control theory. *American Sociological Review*, *81*(4), 828–855.

Smith-Lovin, L., & Heise, D. R. (1988). *Analyzing social interaction: Advances in affect control theory*. Taylor & Francis.

Stewart, T. C., Bekolay, T., & Eliasmith, C. (2012). Learning to select actions with spiking neurons in the basal ganglia. *Frontiers in Decision Neuroscience*, *6*, 2.

Stewart, T. C., Choo, X., & Eliasmith, C. (2010). Dynamic Behaviour of a Spiking Model of Action Selection in the Basal Ganglia. In *Proceedings of the 10th International Conference on Cognitive Modeling* (pp. 235–40).