

Initiation and termination of integration in a decision process

Tobias Larsen^{*,1}, Rafal Bogacz

Department of Computer Science, University of Bristol, United Kingdom

ARTICLE INFO

Keywords:

Decision making
Choice
Neural integrator
Optimality-sequential sampling

ABSTRACT

Experimental data indicate that simple motor decisions in vertebrates are preceded by integration of evidence in certain cortical areas, and that the competition between them is resolved in the basal ganglia. While the occurrence of cortical integration is well established, it is not yet clear exactly how the integration occurs. Several models have been proposed, including the race model, the feed forward inhibition (FFI) model and the leaky competing accumulator (LCA) model. In this paper we establish qualitative and quantitative differences between the above mentioned models, with respect to how they are able to initiate the integration process without integrating noise prior to stimulus onset, as well as the models' ability to terminate the integration after a decision has been made, to ensure the possibility of subsequent decisions. Our results show that the LCA model has advantages over the race model and the FFI model in both respects, leading to shorter decision times and an effective termination process.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

In simple decision tasks where a subject has to make a choice based on a visual stimulus, experimental data show that populations of sensory neurons represent evidence supporting alternative choices in their firing rate. For example, in tasks in which an animal has to decide the direction of prevailing motion in a stimulus, the activity of medial temporal (MT) neurons is higher if the motion in the stimulus moves in the neuron's preferred direction. Thus the firing rate of MT neurons provides information useful for making the decision, and we refer to MT firing rate in this task as the evidence.

However, this representation of sensory evidence is often noisy due to the noise present in the stimulus and the sensory pathways (Britten, Shadlen, Newsome, & Movshon, 1993). Consequently, in order to correctly evaluate the evidence, there is a need to accumulate the evidence over time, and the longer the time, the greater the chance of making the correct choice.² The data also indicate that such an evidence accumulation occurs in certain cortical neuron populations, and that these neurons representing alternative responses integrate the evidence supporting corresponding choices over time (Gold & Shadlen, 2001; Huk &

Shadlen, 2005; Kim & Shadlen, 1999; Mazurek, Roitman, Ditterich, & Shadlen, 2003; Roitman & Shadlen, 2002; Schall, 2001; Shadlen & Newsome, 2001). Furthermore, the data from these two-alternative choice tasks show that: (i) The activities of these integrator neurons are at a base level before stimulus onset. (ii) After stimulus onset, the activity of both integrators increase, but while the integrator neurons corresponding to the strongest alternative continue to increase, the activity of neurons corresponding to the weaker alternative drops to the base level. (iii) After a decision has been made, both activity levels settle back to the base level. The experimental findings (i) and (ii) are apparent in Fig. 1, but (iii) can only be seen in an experiment where measurements are continued after the response is given (e.g. Figure 2 in Shadlen and Newsome (1996)).

Several models have been proposed to describe the integration and decision making process, including the race model (Vickers, 1970), the Feed Forward Inhibition (FFI) model (Shadlen & Newsome, 2001) and the Leaky Competing Accumulator (LCA) model (Usher & McClelland, 2001). These models differ in their inhibitory connections, as shown schematically in Fig. 2 and described in detail in the next section. The race model, which lacks inhibitory connections altogether, predicts that neurons integrating evidence for the winning alternative, as well as those integrating for the losing alternative, should increase their firing rate throughout the duration of the stimulus presentation, and thus is inconsistent with the observation that neuronal activity supporting the weaker alternative decreases during the decision process. Since this makes the race model biologically implausible, its inclusion in this paper is for comparison and engineering reasons, as it might be favourable to find the simplest model that gives the same performance, whether it is implemented in the brain or not.

* Corresponding author.

E-mail address: larsent@tcd.ie (T. Larsen).

¹ New address: TCIN, Lloyd Institute, Trinity College, College Green, Dublin 2, Ireland.

² In these experiments, the reward given for each stimulus is the same, so that the evidence upon which the decision is made are purely sensory, and not affected by reward expectation.

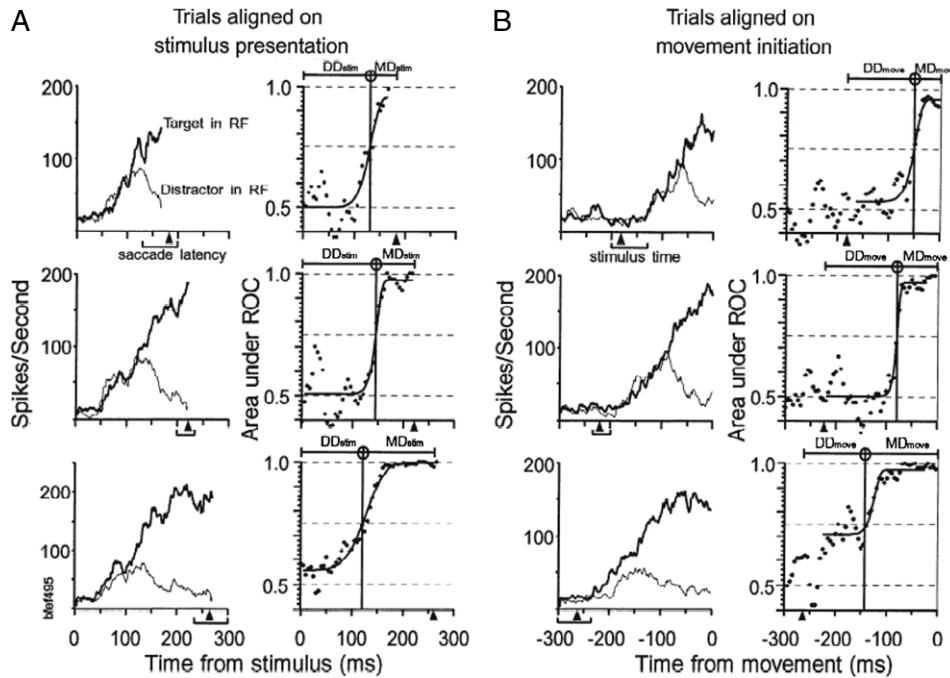


Fig. 1. Time course of target discrimination of a visuomovement cell in the Frontal Eye Field after trials were separated into short (top), middle (middle) and long (bottom) saccade latency groups. In A and B the data are aligned on stimulus presentation and movement initiation respectively. A left and B left show the neuron activity when either the target or a distractor was present in the measured cell's Response Field (RF). A right and B right are divided into Discrimination Duration (DD) and Motor Duration (MD) time.

Source: Reproduced with permission from Thompson et al. (1996).

In contrast, both FFI and LCA models (and their spiking neuron equivalents mentioned in the next section) have been shown to be consistent with both behavioural and neurophysiological data from decision making tasks (Mazurek et al., 2003; Shadlen & Newsome, 2001; Usher & McClelland, 2001; Wang, 2002), though the most successful model in describing human behavioural data from two alternative choice tasks is the diffusion model (Ratcliff, 1978; Ratcliff & McKoon, 2008). In the diffusion model a single integrator integrates the difference in evidence between the competing alternatives towards a decision threshold, either positive or negative. The diffusion model is not directly implementable, but since the integrators in the FFI model also integrate the differences between sensory evidence, the linear version of the FFI model is equivalent to the diffusion model (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006). The LCA model can also closely approximate the diffusion model for certain values of its parameters (Bogacz et al., 2006). The FFI and LCA models make different experimental predictions, as described in the Discussion, but the experiments testing them have not been performed yet. Therefore, on the basis of existing experimental data it is still unclear which of these two models is the more biologically plausible.

However, from an ecological point of view, a model allowing better performance would give an evolutionary advantage, and if such a model fulfils other requirements, such as energy efficiency and biological plausibility, it may be more likely to have been chosen in the course of evolution. Although it is known that animals and humans do not always employ optimal strategies (Kahneman & Tversky, 1984), such an ecological approach has been proved to be very successful in generating experimental predictions in psychology, biology and neuroscience (Anderson, 1990; Belovsky, 1984; Olshausen & Field, 1996; Salinas, 2006). Therefore, in this article we compare the performance of decision making models in which sensory evidence is cortically integrated according to FFI, LCA, a recently proposed improved version of the race model (Stafford, Humphries, & Chambers, 2005) and a novel version of the FFI model with an integration threshold, and describe

the predictions of this analysis. In particular, we compare how the above models are able to integrate sensory evidence without integrating noise before stimulus onset and after stimulus offset.

As a measure of performance, we use decision time for a given error rate, as the ability to make faster decisions for any given error rate has been shown mathematically to maximize reward rate (Bogacz et al., 2006), because the ability to make faster rewarded decisions leads to more reward per unit of time. This measure of performance has also been used in previous studies comparing decision making models (Bogacz & Gurney, 2007; Bogacz et al., 2006; Bogacz, Usher, Zhang, & McClelland, 2007; McMillen & Holmes, 2006).

For a choice between two alternatives based on noisy evidence, the optimal performance is achieved by the decision rule defined by the Sequential Probability Ratio Test (SPRT) (Wald, 1947). This test is optimal in that for any accuracy in a two alternative task, the SPRT gives the fastest decision time compared with any other statistical test (Wald & Wolfowitz, 1948). A decision maker employing SPRT will therefore maximise the rate of receiving rewards for correct choices in a wide range of tasks involving a choice between two alternatives (Bogacz et al., 2006).

There exists a generalization of SPRT for choices between multiple alternatives, called Multihypothesis SPRT (MSPRT) (Dragalin, Tartakovsky, & Veeravalli, 1999). In a choice between multiple alternatives, MSPRT was shown in McMillen and Holmes (2006) to achieve better performance than simple networks, including the race model and the LCA model. However, the decision processes of the brain do not necessarily rely solely on networks with simple architectures as in Fig. 2. It has been proposed that while the cortical regions integrate evidence supporting various alternatives, the basal ganglia resolves the competition between the alternatives (Frank, 2006; Gurney, Prescott, & Redgrave, 2001; Redgrave, Prescott, & Gurney, 1999). It has been shown recently that many aspects of the network involving the cortex and basal ganglia are exactly those required to implement MSPRT (Bogacz & Gurney, 2007).




Fig. 2. The architectures of the race (a), the feed forward inhibition (b) and the leaky competing accumulator (c) models for two alternatives. Arrows indicate excitatory connections, dots indicate inhibitory connections. In each panel, the left column of circles corresponds to populations of neurons providing sensory evidence for the two alternatives, while the right column of circles corresponds to populations of cortical integrator neurons.

Bogacz and Gurney (2007) showed that the circuit of basal ganglia and cortex can implement the MSPRT for different cortical integrative networks including the race model (Vickers, 1970), the FFI model (Shadlen & Newsome, 2001) and the LCA model (Usher & McClelland, 2001). These networks were originally proposed as complete models of a decision process, but with the MSPRT evaluating whether sufficient evidence has been integrated, they only handle the evidence integration. Since the MSPRT ensures optimal decisions, the model of the cortico-basal ganglia circuit achieves exactly the same decision time (for given error rate) no matter which of the models (race, FFI, LCA) is used to describe the cortical integration, and hence none of the integrative networks should have an evolutionary advantage over the other.

The problem with this comparison is that in any realistic system, there is a period before the integration where no evidence is presented to the integrators, and if nothing is done to prevent it, noise will be integrated in this period thereby altering the starting point of the actual evidence integration. This means that even though two models have the same performance when integrating evidence, they might behave differently before the integration, thereby having different overall performances.

A similar problem arises at the end of a decision. Once a decision has been made, the neural integrators need to reduce the firing rate to a base firing rate to allow for subsequent decisions.

In the following article, we investigate which of the proposed models of cortical integration provides the most efficient initiation and termination for the integration processes, thereby having an evolutionary advantage over the others. We will consider these cortical models as a part of a decision network also involving the basal ganglia.

Although the decision network involving cortical integrators and the basal ganglia can optimise the choice between multiple alternatives, for simplicity of explanation, we analyse the choice between two alternatives, because this is the only case in which the dynamics of cortical integration models can be visualised in two-dimensional state-space and can be easily understood. Nevertheless, our results naturally generalise to multiple alternatives.

2. Review of decision models

This section first describes the experimental paradigms that decision making models seek to characterise. Then it describes the models of the two parts of the brain's decision making network: cortical models of evidence integration, and the model of competition resolution in the basal ganglia.

2.1. Experimental paradigms

Many of the two alternative choice experiments are variations of the same basic setup: A visual stimulus is presented to a subject and after some time, a choice based on the stimulus must be

indicated. In these experiments the stimulus is designed in such a way that there is a continuous flow of evidence to be evaluated. A typical stimulus consists of a display of randomly moving dots with some dots coherently moving in one direction (Britten et al., 1993). The task is then to decide the direction of the prevalent motion.

There are two main variations of this paradigm. In the interrogation – or response signal paradigm – the subject has to give the response at a particular time (usually indicated to the subject by an additional stimulus), while in the free response paradigm they are free to respond at any time. In this article, we focus on the free response paradigm.

2.2. Models of evidence integration

The models reviewed in this subsection were originally proposed as models for the entire decision process, i.e., to describe both evidence integration and competition resolution. However, since in this article we assume that the competition resolution is implemented in the basal ganglia, we only review the evidence integration of the models.

In this subsection we focus on the population level models describing the average activity of neuronal populations rather than of individual neurons. Consequently, all these models include cortical *integrators* with activity denoted by y_i accumulating evidence supporting alternative i . In their original descriptions, these models assume that at the beginning of the integration process the starting point of integration is $y_1(0) = y_2(0) = 0$. The models can also be classified as Sequential Sampling Models (Ratcliff & Smith, 2004) as they assume that new evidence is added to the evidence accumulated so far as it is presented.

In all models described in this subsection we additionally assume that the activity levels of the integrators cannot be negative, so that if $y_i + dy_i < 0$, then y_i is set to 0.

2.2.1. Race model

Although earlier Sequential Sampling Models had been proposed (e.g. LaBerge, 1962), the Accumulator Model (Vickers, 1970) and the Poisson Counter Model (Pike, 1966) were the first models that were able to correctly predict the skewed shape of reaction time (RT) distributions (Ratcliff & Smith, 2004). In both of these models (and in LaBerge, 1962, as well), two separate integrators y_i are accumulating evidence supporting competing responses (see Fig. 2(a)). This can be described in the continuous time limit as

$$\begin{aligned} dy_1 &= x_1 dt + c dW_1 \\ dy_2 &= x_2 dt + c dW_2 \end{aligned} \quad (1)$$

where dy_i is the change in accumulated evidence in integrator i during time dt , x_i is the mean level of evidence supporting alternative i , dW_i is an independent noise process with mean 0 and variance dt , and c describes the amount of noise. We refer to the model described in the above equation as the race model.

2.2.2. Race model with threshold

The race model suffers from the problem that it integrates noise even in the absence of evidence because it has no method for initiating the integration. This happens even though the noise is Gaussian with mean 0, because the activity levels of integrators are not allowed to be negative, and the integration therefore has a boundary at 0 which results in an average increase in activity. This results in a variance of the starting point of integration ($y_1(0), y_2(0)$) that is proportional to the time allowed to elapse before the stimulus onset. As a solution to this problem, a decay of activity (described by parameter k) can be introduced to the race model when the activity level of integrators is below a certain integration threshold, θ_{int} , thereby creating a race model with threshold (RwT). This model is a simplified version of work presented in Stafford et al. (2005).

$$\begin{aligned} dy_i &= \begin{cases} (x_1 - ky_1)dt + cdW_1 & \text{if } y_1 < \theta_{int}, \\ x_1dt + cdW_1 & \text{if } y_1 \geq \theta_{int} \end{cases} \\ dy_2 &= \begin{cases} (x_2 - ky_2)dt + cdW_2 & \text{if } y_2 < \theta_{int}, \\ x_2dt + cdW_2 & \text{if } y_2 \geq \theta_{int}. \end{cases} \end{aligned} \quad (2)$$

This mechanism ensures that the model integrates evidence only if the input is sufficiently high to push the activity level of an integrator above the integration threshold θ_{int} . The introduction of a leak only for weakly active neurons is not biologically plausible, but the model was originally formulated by Stafford et al. (2005) in a biologically plausible way. In particular, in their original model, each integrator has a constant leak, and sends excitatory feedback to itself which balances the leak and thus allows perfect integration. However, the feedback is provided via another neuronal population which transmits the feedback only when the activity of integrator $y_i \geq \theta_{int}$, thus it results in a leaky integration when $y_i < \theta_{int}$.

2.2.3. Feed forward inhibition model

In the FFI model, each integrator, in addition to accumulating evidence from the appropriate input, receives inhibition from the competing input channels (see Fig. 2(b))

$$\begin{aligned} dy_1 &= (x_1dt + cdW_1) - v(x_2dt + cdW_2) \\ dy_2 &= (x_2dt + cdW_2) - v(x_1dt + cdW_1) \end{aligned} \quad (3)$$

where v is the weight of the inhibitory connection, set to $v = 1$ in the original description of the model (Shadlen & Newsome, 2001).

A spiking neuron version of this model has also been developed (Mazurek et al., 2003), and has been shown to successfully describe both the responses of integrator neurons in the lateral intraparietal (LIP) area and the reaction times in a motion discrimination task (see Section 2.1) (Mazurek et al., 2003).

2.2.4. Feed forward inhibition model with threshold

The FFI model has the same problem with integration initiation as the race model, and in a similar attempt to resolve it, a leak can be introduced whenever the activity of the integrators are below a threshold.

$$\begin{aligned} dy_1 &= \begin{cases} ((x_1 - ky_1)dt + cdW_1) - v(x_2dt + cdW_2) & \text{if } y_1 < \theta_{int}, \\ (x_1dt + cdW_1) - v(x_2dt + cdW_2) & \text{if } y_1 \geq \theta_{int} \end{cases} \\ dy_2 &= \begin{cases} ((x_2 - ky_2)dt + cdW_2) - v(x_1dt + cdW_1) & \text{if } y_2 < \theta_{int}, \\ (x_2dt + cdW_2) - v(x_1dt + cdW_1) & \text{if } y_2 \geq \theta_{int}. \end{cases} \end{aligned} \quad (4)$$

2.2.5. Leaky competing accumulator model

In the LCA model two leaky integrators accumulate evidence in support of the alternatives, and they inhibit each other (see Fig. 2(c)). A two alternative LCA is described by

$$\begin{aligned} dy_1 &= (x_1 - wy_2 - ky_1)dt + cdW_1 \\ dy_2 &= (x_2 - wy_1 - ky_2)dt + cdW_2 \end{aligned} \quad (5)$$

where k is the leak from the integrators and w is the weight of the inhibitory connection between the two integrators.

Due to the leak from integrators, the LCA model has an inherent initiation process that ensures that noise is not integrated significantly in the absence of stimuli (Brown et al., 2005), as will be described in detail in Section 4.2.

A spiking neuron model of integration has been proposed by Wang (2002), in which the integrator neurons project to a population of inhibitory neurons which inhibit all integrators. This model has also been demonstrated to successfully describe the responses of integrator neurons in the area LIP. It has been shown that a population-level model with the architecture of the Wang model predicts approximately the same dynamics of integration process as the LCA model (Bogacz et al., 2006). Due to the close relation of the Wang model to the LCA model we do not analyse the Wang model in this paper separately, but we note that we consider the Wang model as a possible biologically realistic implementation of the LCA model.

2.3. Model of competition resolution in basal ganglia

This subsection briefly reviews the involvement of the basal ganglia in decision making and how the MSPRT is related to the disinhibition of motor actions. The basal ganglia is a set of subcortical nuclei connected with one another and with the cortex. The output nuclei in their default state send inhibition to the thalamus and the brainstem thus blocking execution of motor actions, and the actions are selected when the activity of corresponding neurons are decreased (Chevalier, Vacher, Deniau, & Desban, 1985; Deniau & Chevalier, 1985). On the basis of these properties, Redgrave et al. (1999) have suggested that the basal ganglia acts as a ‘‘central switch’’ resolving the competition between cortical regions which vie for behavioural expression.

Gold and Shadlen (2001) formulated the goal of the decision process as choosing the alternative with the highest mean evidence (Gold & Shadlen, 2001, 2002). Hence the problem of decision making between N alternatives can be formulated in statistical terms as a choice between N hypotheses H_i . Each hypothesis H_i states that of all the alternatives, the sensory input supporting alternative i has the highest mean (Bogacz & Gurney, 2007; Gold & Shadlen, 2001, 2002). In the original formulation of the MSPRT (Baum & Veeravalli, 1994), at each moment of time and for each alternative, the probability P_i of the hypothesis H_i given the sensory evidence observed so far is evaluated, and the decision is made as soon as any P_i exceeds a decision threshold.

It has been proposed that the basal ganglia is organised into channels defined as groups of neurons representing individual motor actions, and these channels traverse all basal nuclei (Alexander, Crutcher, & DeLong, 1990). Bogacz and Gurney (2007) proposed that in the case of motor decisions, the firing rate of channel i of the output nuclei of the basal ganglia is proportional to $OUT_i = -\log P_i$. Hence the decision needs to be made as soon as any OUT_i decreases below a decision threshold, which is consistent with the selection by disinhibition mentioned above. Computing $-\log P_i$ from Bayes theorem gives (Bogacz & Gurney, 2007)

$$OUT_i = -y_i + \ln \sum_{k=1}^N \exp y_k \quad (6)$$

where y_i is the integrated evidence supporting alternative i , which can be computed using the various models presented in Section 2.2. Eq. (6) includes two terms, the first ($-y_i$) describing the amount of integrated evidence supporting alternative i , and

the second describing the amount of conflict between alternatives (as it involves the summation over all alternatives). Thus the more conflict between alternatives, the higher y_i needs to be for OUT_i to decrease below the decision threshold. The details of how the computation of Eq. (6) could be performed in the cortico-basal ganglia circuit can be found in Bogacz and Gurney (2007).

Note that in this model the cortex and basal ganglia do not operate sequentially, but rather in parallel: the cortical neurons integrate sensory evidence, while the basal ganglia continuously monitor whether sufficient information has been accumulated. Such a parallel involvement of the cortex and basal ganglia is supported by experimental data showing that basal ganglia neurons gradually change their firing rate during the decision process long before the response is initiated (see Samejima, Ueda, Doya, and Kimura (2005) or Williams, Neimat, Cosgrove, and Eskander (2005), Figure 5).

3. Methods of simulation

To assess the integration of evidence in RWT, FFI, FFIwT and LCA models we compare the performance of decision networks in which evidence is integrated according to one of these models and the competition is resolved by the MSPRT model of the basal ganglia.

The models were simulated using the Euler method with an integration constant of 0.001s. The performance of these decision networks is described by error rate (ER) – defined as a proportion of trials in which the alternative with lower x_i is chosen, and decision time (DT) – defined as the interval from stimulus onset until decrease of any OUT_i below the decision threshold. In the models simulated, there always exists a speed-accuracy trade-off controlled by the decision threshold. Hence, to be able to compare the DT of different networks, it is necessary to ensure that they produce the same ER. This is done by altering the decision threshold until the desired ER is met with the desired precision (Bogacz et al., 2006; McMillen & Holmes, 2006). In the simulations described in this article, DT is compared for ER = 1%. For a given tested value of the decision threshold, the simulated decision process is repeated until it is established whether this decision threshold results in ER = 1% with confidence intervals $\pm 0.2\%$.

Due to the noisy nature of the simulations, there is an inherent possibility that the decision threshold will not be crossed in finite time. This possibility is usually marginal, but with decreasing decision threshold the probability increases. For this reason the simulations are implemented with a time limit of 14 s that if exceeded will terminate that particular simulated decision, and the trial is discarded.

In all our simulations, the number of alternatives (and thus the number of basal ganglia channels, N) is 2, and with $x_1 > x_2$ we assume that the first alternative is the correct choice. In all the simulations, each trial is divided into an initial time and a decision time as shown in Fig. 3. At the beginning of the initial time, the integrators are set to $y_1(0) = y_2(0) = 0$. In the initial time the input parameters are: $x_1 = 0, x_2 = 0, c = 0.33$, and in the decision time the input parameters are: $x_1 = 4.5, x_2 = 3, c = 0.33$. These values are estimated from a representative subject of an experiment described by Bogacz et al. (2006), and kept fixed for all simulations.

4. Results

This section starts with an analysis of dynamics of the integration models which helps with the understanding of the comparison of performance presented in Section 4.2. Then the ability of the models to terminate the decision process is compared.




Fig. 3. Changes of input parameters x_1 , x_2 and c with stimulus onset dividing a simulated trial into an initial time without stimulus, and a decision time with stimulus.

4.1. Dynamics of integration process

In order to get an understanding of the models' behaviour, it is useful to analyse their potential fields (vector fields) in the different stages of decision making. Fig. 4 compares the potential fields of the RWT ((a), (e)), the FFI ((b), (f)), the FFIwT ((c), (g)) and the LCA ((d), (h)) models before and after stimulus onset. The axes describe the levels of activity in each of the two integrators, y_1 and y_2 . The vector field in each panel shows how the integration process would behave on average (i.e., without noise) in a current state. Each arrow shows how the integrators would, on average, change their activity from that particular state. The lengths of the arrows denote the relative strength of the activity change but cannot be compared between panels.

The black line in each of the panels shows the development of the activity in the integrators during a single simulation of an integration process.

4.1.1. Before stimulus onset

The upper row of panels in Fig. 4 illustrates the dynamics of the models when there is no stimulus input ($x_1 = x_2 = 0$).

In this case the integrators in the race model will have no preferred direction of change (as according to Eq. (1), $E(dy_1) = E(dy_2) = 0$), and will exhibit random walk behaviour based on the noise levels. This is also the case in the RWT model when both integrators are above the integration threshold (there are no arrows in top right part of Fig. 4(a)). However, in the RWT model each integrator will have a preferred direction of change towards 0 for $y_i < \theta_{int}$ as shown in Fig. 4(a). This means that even though noise might raise the activity of an integrator by chance, it will, provided the activity has not crossed θ_{int} , drift back towards 0, thereby limiting the possible starting points of integration.

The FFI model does not have an integration threshold, and the integrators have no preferred direction of change. Fig. 4(b) shows that even though the integrators on average have no preferred direction, they only change along diagonal lines with the same slope. This is due to the fact that the slope of each step of the integration path is $\frac{dy_2}{dy_1}$ and from Eq. (3) it is equal to $\frac{dy_2}{dy_1} = -1$ for $v = 1$. This is the case for all inputs and is therefore independent of the noise. The sum of activity of integrators only increases when the integrator that receives the least input is at 0 and therefore unable to drop any further.

The integration path in the FFIwT follows roughly the same diagonal lines as FFI without a threshold, even though due to the leak, the slope will not always be -1 . When comparing Fig. 4(b) and (c) it can be seen that the leak ensures that the activity is kept closer to (0, 0) thereby constraining the potential starting points.

In the LCA model without a stimulus, each of the integrators has an inherent drift towards 0 (Brown et al., 2005) due to leak of activity, as can be seen in Fig. 4(d). This is effectively ensuring that the starting point of integration will be kept close to (0, 0).

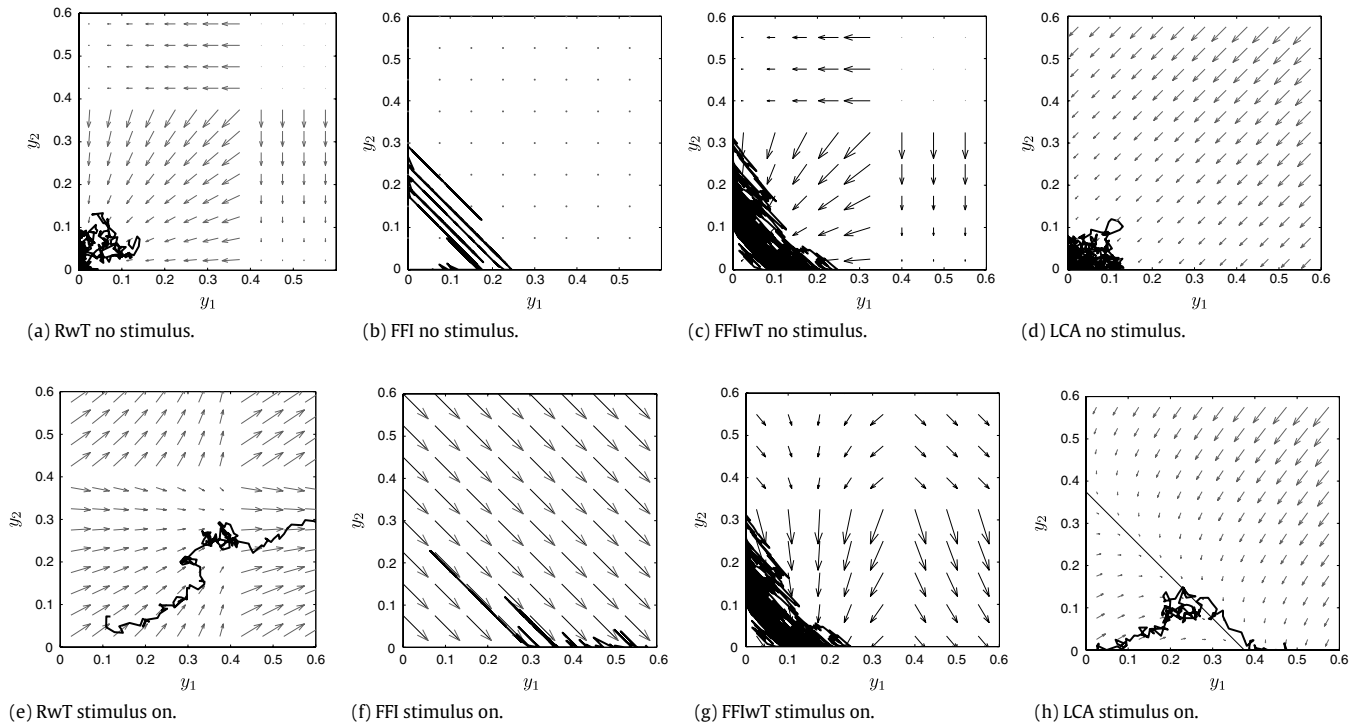


Fig. 4. The potential fields before stimulus onset and with stimulus for the race with threshold (RwT) ((a), (e)), the feed forward inhibition (FFI) ((b), (f)), the feed forward inhibition with threshold (FFIwT) ((c), (g)) and leaky competing accumulator (LCA) ((d), (h)) models respectively. The paths show the evolution of the models in a single simulated trial. The following parameters of the models were used: RwT: $\theta_{int} = 0.33$, $k = 10$. FFI: $v = 1$. FFIwT: $v = 1$, $\theta_{int} = 0.33$. LCA: $k = 10$, $w = 10$.

4.1.2. With stimulus

The lower row of panels in Fig. 4 illustrates the dynamics of the models when the stimulus is present. Recall from Section 3 that we choose $x_1 > x_2$ thus we assume that the first alternative is the correct choice.

Once the stimulus is presented, each of the integrators in the RwT model increases its activity. As the activity increases, change in activity decreases until the leak and the evidence input evens out (i.e., $x_1 - ky_1 = 0$). If the integration threshold θ_{int} is situated before this point (where $y_1 = \frac{x_1}{k} = 0.3$ for the losing alternative in our simulations), the integration will continue without boundaries, as seen in integrator y_1 in Fig. 4(e), but if the integration threshold is situated after this point, the integration may saturate, as seen in integrator y_2 in Fig. 4(e). Once θ_{int} has been crossed, the change in activation is constant, only depending on the evidence input.

The FFI model's crossed inhibitory connections from inputs to the integrator units will result in the level of activity in the losing integrator decreasing towards 0. This is the desired behaviour, as it mimics the behaviour of integrating neurons observed in the brain. To ensure that the activity of the losing integrator decreases towards 0, the weight of the inhibition (v) must obey $v \geq x_2/x_1$, where x_2 is the mean input to the incorrect alternative. For this to be true for all $x_1 > x_2$, the inhibition weight must be 1. As in the case before stimulus, the model changes its state only along a line with a slope of -1 , or along one of the axes.

The FFIwT has the same behaviour as described for the FFI and must also have $v = 1$ in order for the losing alternative to decrease its activity towards 0. In addition to this, the integration threshold must be positioned low enough to allow the integrator for the winning alternative to drift across the threshold. Otherwise, the model would be equivalent to adding a leak to the integrators without the threshold. To allow for this, the threshold must be set less than or equal to $\frac{x_1 - x_2}{k} = 0.15$. In Fig. 4(g) the threshold is set higher to illustrate the problem. Note that although the RwT model is simulated for the same parameters as the FFIwT model,

the former is able to cross θ_{int} (Fig. 4(e)) while the latter is not (Fig. 4(g)).

When the stimulus is presented to the LCA model, an attracting line (the diagonal line in Fig. 4(h)) with a slope of -1 appears in the state space and it is situated a distance $\frac{x_1 + x_2}{\sqrt{2 \cdot (k + w)}}$ from $(0, 0)$ (Bogacz et al., 2006). The intuition for the existence of this line is the following. When the stimulus is turned on, both integrators increase their activity levels, which corresponds to movement towards the attracting line. However, when they are sufficiently active, the inhibition starts to play a bigger role, so that if one increases, the other needs to decrease, which corresponds to the movements along the attracting line. The average direction of evolution along the attracting line depends on which of the alternatives has higher mean input, resulting in integration saturation at the point in the state space where the losing integrator is 0 and the value of the winning integrator is defined by the attracting line, see Fig. 4(h). Thus in the LCA model the integration reaches a natural saturation point, which does not happen in the other models.

4.2. Performance with initial time

In this subsection we compare the performance of decision making models in which evidence is integrated according to the RwT, FFI, FFIwT or LCA model, and the competition is resolved in the MSPRT model of basal ganglia. By introducing an initial time, we also introduce the possibility that on stimulus onset the starting point of integration is not at 0. Such a distribution of the starting points will increase ER for a given decision threshold. Since we compare the performance for the same ER, the decision threshold will need to be adjusted, which will give longer DTs. If the distributions of the starting points are different for the four models, they will have different DTs for the same ER.

The performance of decision networks also depends on the parameters of the integration models (i.e., k , θ_{int} , v , w). In order to make a fair comparison between the models of integration,

Fig. 5. Decision times (DTs) for decision making networks in which the competition resolution is implemented in MSPRT model of basal ganglia and the integration is performed by one of the following models: (a) race with threshold, (b) feed forward inhibition and (c) leaky competing accumulator. DTs are shown for initial times between 0 and 1500 ms, and different values of the parameters of integration models.

Fig. 6. Comparison of decision times for varying leak and integration threshold in the Feed Forward Inhibition with Threshold model with $v = 1$, initial time 500 ms and error rate 1%.

we need to compare the models for their best parameters; hence before we perform the comparison we first find the optimal parameters for each model of integration.

4.2.1. Optimal parameters for the RwT model

When simulating the RwT model, there are two parameters to optimise; the leak (k) and the integration threshold (θ_{int}). However simulations (not shown here) show that as long as the θ_{int} is set accordingly, there is no difference in performance for different values of k , as long as it is of a magnitude comparable to the inputs. Thus in simulations of decisions with integration by the RwT model in Fig. 5(a) the leak is fixed at 10.

Fig. 5(a) show that the optimal value of the integration threshold is around 0.3–0.35. When the integration threshold is set at this level, the initial time before stimulus onset does not influence the DT, which can be seen as the “valley” in Fig. 5(a). The intuition for why the integration threshold should be around 0.3 can be provided by an analysis of the potential field for the RwT model (Fig. 4(e)). It can be seen that when $y_2 = 0.3$, the potential field is perpendicular to y_2 . This happens because at this point $x_2 = y_2 \cdot k$ and the leak therefore balances the input, hindering any further integration. The higher above x_2/k the threshold is set, the lesser is the chance of y_2 crossing the threshold, but this also slows down the integration of the winning channel.

If the integration threshold is set too low, the losing integrator will drift across the integration threshold. If the integration threshold is set too high, none of the integrators will cross the integration threshold. This suggests that the optimal integration threshold is between x_1/k and x_2/k . From a biological point of view, it does not seem likely for a system to use such a specific, input dependent parameter unless there is some automatic setting procedure.

It is also worth noting in Fig. 5(a) that if there is no initial time, increasing the integration threshold increases DT. This follows from the fact that the leak in the RwT model introduces forgetting of information, hence the higher the integration threshold, the more early evidence is discarded by the model, necessitating longer DT to gain required accuracy.

4.2.2. Optimal parameters for the FFI model

There is only one parameter to set when simulating the FFI model, and that is the weight of the inhibitory connections. However, we have argued that to get the desired behaviour of the integrator neurons, the inhibition weight has to be 1. This is however not the optimal parameter when there is an initial time and the output from FFI projects to the MSPRT model of the basal ganglia. Simulations for input parameters given in Section 3 shows that the inhibition weight that gives the shortest DTs for the FFI model is approximately 0.6 (see Fig. 5(b)).

4.2.3. Optimal parameters for the FFIwT model

The motivation for introducing the threshold to the FFI model was to decrease the variation in starting points while maintaining near perfect integration, thereby improving on the performance with initial time. Because there are two parameters to set in the FFIwT model (θ_{int} and k), we find the optimal parameter setting by fixing the initial time to 500 ms and, through an exhaustive parameter search, compare the decision times for all values of θ_{int} and k in a range tested. Note that the points where either $\theta_{int} = 0$ or $leak = 0$ corresponds to the FFI model without threshold.

However, as can be seen in Fig. 6, the optimal performance is achieved for the parameters that reduce the FFIwT model to the FFI model. This shows that there is no performance gain by introducing the leak. The reason for this is that the integration threshold (θ_{int}) needs to be set low to allow for the winning alternative to drift across it, but this low threshold is not enough to constrain the starting point sufficiently. Recall from Section 4.1 that the highest θ_{int} that can be exceeded by integrators after stimulus onset without the help of noise is given by $\frac{x_1}{k}$ in the RwT model, and by $\frac{x_1 - x_2}{k}$ in the FFIwT model, so to allow effective evidence integration after stimulus onset, θ_{int} needs to be set to a lower value in the FFIwT model than in the RwT model. But setting θ_{int} to a lower value impairs noise suppression in the initial time, which

