Parameter Values for FLIF Neurons

Christian Huyck
Middlesex University, UK
c.huyck@mdx.ac.uk

Abstract. A typical human brain consists of roughly 100 billion neurons, and one key aim of Biological Cybernetics is to simulate neural systems. A good model of a neuron accurately represents the behaviour of biological neurons, typically the spiking behaviour. For cybernetics systems that function in real time with thousands, millions, or even billions of simulated neurons, it is also important that the model is computationally efficient. Fatiguing Leaky Integrate and Fire neurons are models that have four free parameters per neuron. They have been used in cybernetic agents, but there have been few links to actual biological behaviour. A model of a rat neocortical neuron is developed with four specific parameter settings. This model is tuned to a particular input regime. When compared to a biological neuron it gets 90% of spikes roughly correct.

1 Introduction

One method for exploring Biological Cybernetics is to build simulated neural systems. At one extreme, this approach has led to a project hoping to simulate the entire human brain [23]. Another framework is to build increasingly sophisticated systems from simulated neurons [11, 15]; this framework builds cognitive models and agents that function in an environment.

One of the key questions for Biological Cybernetics is what neural model to use. There are a large number of neural models (see section 5), and many think that non-neural connectionist models are also a good basis (e.g. Multi Layer Perceptrons [25] or Self Organising Maps [20]).

One framework uses a simple point model of a neuron, the Fatiguing Leaky Integrate and Fire (FLIF) neuron (see section 2). There has been significant progress within this framework including simulated games agents [13] with vision, planning, action and language, machine learning [12], and cognitive models of natural language parsing [14] and task selection [3]. These have all used the same FLIF neural model. An individual neuron has several internal parameters including , firing threshold, leakage rates, and fatigue rates.

One key behaviour of neurons and neuron models is spiking. Neurons receive inputs from other neurons, and emit spikes. These spikes can be induced in biological neurons both *in vivo* and *in vitro*. For example, a neuron can be directly stimulated by electrical current, its internal electrical state can be measured, and spikes can be inferred (see section 3). Consequently, one way of calibrating a neural model, is to compare its spiking behaviour to biological data. In this paper, biological data is used to derive parameter settings for the FLIF neural model, and gauge its accuracy as a model.

2 Fatiguing Leaky Integrate and Fire Neurons

The FLIF model is an extension of the standard leaky integrate and fire (LIF) model, which is an extension of the Integrate and Fire model. A similar model [6] has been shown to account for inter-spike intervals under various input conditions better than the standard LIF model.

One variant of the Integrate and Fire (IF) model is the McCulloch Pitts neuron [24], which has a long standing history and is quite simple. Roughly, neurons are connected by uni-directional synapses. A neuron integrates activity from the synapses connected to it, and if the activity surpasses a threshold, the neuron fires, sending activity to the neurons it connects to. There are two possibilities regarding retaining activity between cycles. The McCulloch Pitts neuron merely throws it away; this prevents low amounts of input causing the neuron to fire. In a second model [1], all of the activity is retained; this allows small amounts of input over time to cause the neuron to fire.

In the LIF model, if a neuron does not fire, it retains a portion (but not all) of its activity making it easier to fire later [22]. Typically, both IF and LIF neurons lose all activity when they fire.

The LIF model is extended to FLIF by the addition of fatigue. When a neuron fires, it fatigues and becomes more difficult to fire.

One of the major components of the model is the firing threshold, θ. A neuron i fires if its activation A_i minus its fatigue F_i is above the threshold.

$$A_i - F_i >= \theta \tag{1}$$

If the neuron fires, it loses all its activation. If sufficient activation is provided from neurons sending spikes to it, it may fire in the next time step.

If a neuron does not fire, some of its activation leaks away. This leak, or decay, is represented by a constant D where $D > 1$. Ignoring external input and assuming i did not fire at $t-1$, activation of neuron i at time t is

$$A_i^t = A_i^{t-1}/D \tag{2}$$

When neuron i fires, it sends activation (or inhibition) along its synapses to other neurons according to the strength of each synapse, so neuron j receives activation according to synaptic strength w_{ij}. The neuron is an integrator, so it accumulates activity from the synapses connected to it. So, given P_j, the prior activation of neuron j, either 0 or equation 2, the activation at time $t+1$ is

$$A_j^{t+1} = P_j + \sum_{i \in V_i} w_{ij} \tag{3}$$

where V_i is the set of all neurons that fired at time t.

These equations describe a LIF model [22], and fatigue is used to extend the model. Fatigue uses two constants; it is incremented by F_c in a cycle when the neuron fires, and is decremented by F_r in a

cycle when the neuron does not fire. $F_i >= 0$, so that firing always requires at least θ retained activation. This makes it more difficult for neurons to fire the longer they are firing.

The model has a loose link with time in biological neurons. The model does not incorporate conductance delays or refractory periods, and these behaviours all happen in under 10 ms., so each given cycle can be considered to be roughly 10 ms. Consequently, each neuron emits at most one spike per 10 ms. of simulated time, and the timing precision is at most 10 ms. This is a shortcoming of the model, but enables efficient simulation of hundreds of thousands of neurons in real time on a standard PC. It is consistent with the neural data (see section 3), as the neuron being modelled never spikes more than once in a 10 ms. interval.

3 Neural Data

The neural data was used for a neural modelling competition [4], and the data was from Challenge A of that competition. More details can be found there along with the data.

A neuron from the primary somatosensory neocortex of a rat was extracted. So, the datum that was collected was *in vitro*.

A probe was placed into the neuron. Current was injected directly to the neuron and cellular voltage was measured at .1 ms. intervals.

Input varied over 60 seconds, but only the first 39 seconds were available with the remaining 21 seconds used as the test for the competition. There was an initial input phase, followed by two seconds of 300mV input, then two seconds with no input, then two second of 600mV input, then two seconds with no input, then two seconds of 900mV input, then two seconds of no input, followed by 42.5 seconds of stochastic input, with 21.5 of that available as data. The same input regime was applied 13 times to the neuron.

The cellular voltage measurements were converted to spikes. With no input, the voltage hovered around -65mV. Under input this gradually increased. When it neared -40mV it then rapidly increased to positive values around 40mV. It then rapidly returned to a negative value, and then more gradually to a value around -55mV while under input. The spike was calculated from the first point the voltage crossed from negative to positive. It then was reset when it crossed back to negative. In the data, the shortest inter-spike interval was 13.5 ms.

4 Simulations and Parameter Settings

The task was to discover appropriate parameter settings for the FLIF model for input, the threshold θ, leak D, fatigue F_r, and fatigue recovery F_c. The parameters interact in relatively simple ways, so initially the goal was to find one parameter set, where the set led to model behaviour that was a relatively close fit to the biological data.

4.1 Clamped Input

The first question was how to scale the input. As the first input of interest was 300mV at a rate of .1 ms., the input was averaged over the full time of the step (10 ms. or 100 pieces of data), and converted to units in volts. For example, input at time steps 50,000 to 50,100 (5 to 5.1 seconds) were all 300mV, and this is converted to .3 for input.

Next some analysis of the first series of inputs was used to set initial parameters. This series of inputs was two seconds of 300mV. Analysis of the biological data of one neuron showed that there were 12 spikes at an average of 174.5 ms. apart. For the FLIF model, that was every 17 cycles. Using data derived from other studies [21], D was set at 1.25. Using these values to see what threshold was passed at 17 cycles, and this was $\theta = 1.46$. These parameter settings led to data that fit the 300mV input. However, the parameters also needed to work for the other inputs.

Two other sets of input were of considered next; after rest periods of two seconds, there were two seconds of inputs at 600mV, then two seconds rest, and two seconds of 900mV. Some analysis of this data showed that the neurons spiked on average every 60.38 ms. for 600mV and every 40.51 ms. for 900mV, or every six and four FLIF cycles respectively. With $\theta = 1.46$ and $D = 1.25$ the model led to firing every three steps for 600mV and every two steps for 900mV. Reducing D, retaining more activation per step, moved things in the right direction for 600mV and 900mV.

Setting $D = 1.1$ led to values that worked. Using the process described above for 300mV, θ was calculated at 2.6. This led to the desired behaviour with spikes every 17 cycles for 300mv, every 6 cycles for 600mV, and every four cycles for 900mV.

These parameters determine a LIF model. Perhaps fatigue was unnecessary. Some further analysis of the data showed that fatigue would improve fit. Figure 1 shows the latencies, time between spikes, of the neuron under the three input regimes. After initial rapid firing, all three spike latencies are relatively stable. However, there is a gradual increase at 900mV, increasing on average of .329 ms. between each spike after the third spike; that is, the spikes are coming more slowly, implying that at this elevated firing rate fatigue has an effect. At the lower firing rates, the slope is virtual flat, so the rate remains roughly constant.

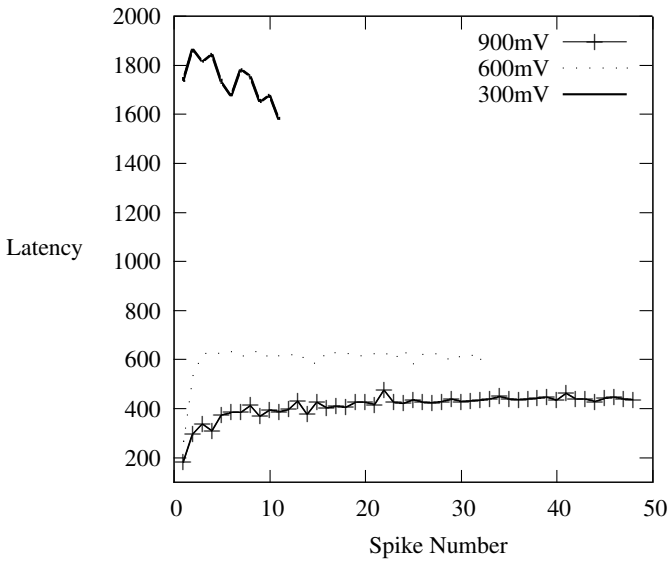

Figure 1. Latencies between spikes at different input values over two seconds

To calculate the fatigue F_c and fatigue recovery F_r values, this behaviour imposes some constraints. Rates at less than once per six cycles should not accumulate fatigue, and rates at once per four cycles should. So $F_c > 3F_r$, and $F_c < 5F_r$.

For the 900mV case, firing rates are initially just over every 35 ms., and they pass every .45 ms. around spike 35. Expanding the formulas with $\theta = 2.6$ and $D = 1.1$ shows a neuron having 3.138 units of

activation at any fourth cycle. It needs 2.6 to fire, so it has 0.538 surplus activity. For accumulated fatigue to cause it pause for another cycle, it must be greater than .538. As this should not accumulate for 35 cycles, $F_c - 3F_r \sim .538 \div 35$. So, $F_c - 3F_r \sim 0.015$.

F_r was selected as 0.1, leaving F_c as 0.315. Running simulations on this showed that indeed spike rates at this the 35th cycle were every five cycles, but they returned to every four cycles there after. Reducing F_r to 0.01 left F_c at 0.045. The model latencies increased to five cycles after 35 spikes and continued to increase thereafter. The model predicts that after 200 spikes the latency would increase to six cycles, a testable hypothesis though not in the data.

This leaves all parameter values determined. Threshold is $\theta = 2.6$, leak is $D = 1.1$, fatigue is $F_c = 0.045$ and fatigue recovery is $F_r = 0.01$.

4.2 Stochastic Input

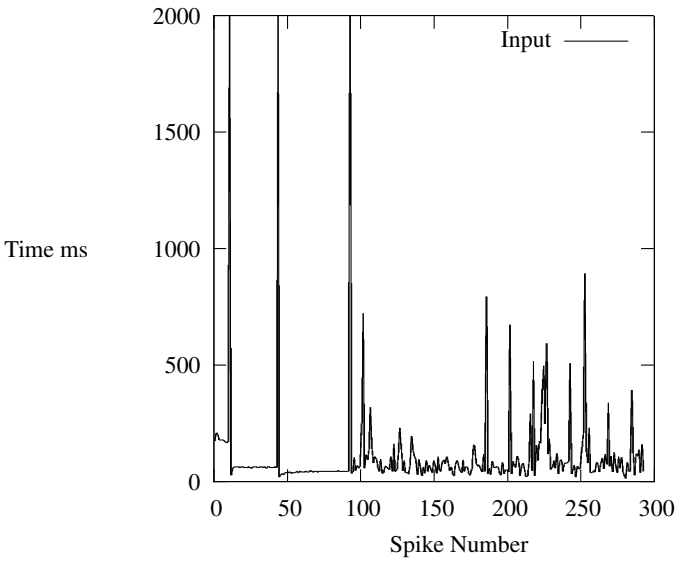

Figure 2. Interspike Latency of Biological Neurons

The first 17.5 seconds of input was clamped with either -300mV, 0mV, 300mV, 600mV or 900 mV, and the stimuli persisting for two seconds, with intervals of no input in between. After this, the remaining 21.5 seconds of input went through a rapidly varying stochastic input. Figure 2 shows the response of the biological neuron to these inputs in the form of inter-spike intervals. The first 11 spikes come around 170 ms. apart, followed by two seconds with no input, and thus no spikes. After spike 93 and the third two second delay, varying stochastic input began. Note that many periods during stochastic input experience low input and thus high spike latency.

Using the same parameters, the model was run to account for this behaviour. For the biological neuron, input varied every .1 ms.; as the simulation steps accounted for 10 ms. of time, the inputs were binned into groups of 100 and averaged over that time.

During the clamped period, the model produced 92 spikes, and the biological data 94. The missing spikes were during the 900mV input.

During the stochastic period, the model produced 151 spikes and the biological data produced 200. The model spikes and real spikes were aligned, with the model spikes being placed adjacent to the nearest real spike. Analysis of these aligned spikes showed three categories of problems.

The first problem was evident even under clamped input. The first biological spikes after input resumed came earlier than the model predicted. This is evident in figure 1 when the initial spikes are quite rapid, but the first spikes are even more rapid; for example at 300mV, the first spike occurs 36 ms. after input resumes. This is also evident in the stochastic period when spikes are missed after relatively long periods of low or negative input are followed by moderate or high input.

The second problem was that periods of rapid biological spiking led to missing modelled spikes. This implied that the threshold was too high under the parameter settings.

The third problem was that spikes were missed over periods of low input. This implied that there was too much leakage in the parameter settings.

A further search of the parameter space ensued. Reducing the threshold and decay separately or together led to improved behaviour under stochastic input, but worse behaviour under clamped input. In particular, reducing decay requires increasing threshold, which had an adverse effect when there is high input. On the other hand, moderate increases of decay (e.g. $D \sim 1.2$) made it difficult to spike under low input (e.g. 300mV). Reducing threshold to allow this, made it spike too much under higher input.

Setting $D = 1.12$ and $\theta = 2.2$ left a good compromise. There were too many spikes at low clamped input (300mV had 14 model and 11 real, and 600mV had 40 model and 32 real), but most of the stochastic spikes (183 of 200) were present.

Fatigue had a minor effect on the results. Removing fatigue did increase the false positives with 900mV input. The original parameters correctly produced all 49 spikes, but added two incorrectly. Removing fatigue added another seven incorrect spikes. It also added an extra eight spikes to the stochastic input. The 300mV and 600mV clamped input remained unchanged. The effect of fatigue was negligible because, outside the clamped 900mV period, there was not a sustained period of high input to cause fatigue to accumulate.

Of the 288 spikes emitted by the model, 26 alignments had two model spikes aligned with a biological spike. The first of these was taken for a timing comparison. Of the 260 directly aligned spikes, the average variance between the model and biological time was 16.3 ms.

The simulation is open to the criticsm of testing on the training set. While there are only four parameters, this is still a valid criticism. However, when the model was compared to a second run of the same input on the biological data, the model's fit improved the a small amount. In the second run, the biological data produced an extra spike for each of the three clamped input regimes, and one less spike for the stochastic input. Of the 288 spikes emitted, 261 aligned with an average variance of 15.3 seconds.

5 Other Neural Models

The FLIF model presented and compared in this paper is a relatively simple point model, where the model does not consider any spatial components of the neuron. Compartmental models [9, 7] are appreciably more complex, mapping the structure of the entire neuron body. These models can be further refined to include, synaptic delays, ion transfer, and so forth. The FLIF model was compared to the real spike data. Compartmental models can compare, relatively

accurately, at the actual voltage level. While they are more accurate, compartmental models are appreciably more expensive computationally to simulate.

A primary motive of the FLIF model is computational efficiency, so that thousands, millions and even billions of neurons can be simulated in real time to support the cognitive aspects of behaving agents. The trade off between biological accuracy and computational efficiency is skewed more toward efficiency in the FLIF (and other point) model and toward accuracy in compartmental models.

Simple IF models [24, 10] would not have faired well simulating this data. If no activity were retained between steps, it would not have been able to spike with small inputs, or it would have emitted thousands of spikes if the threshold was set low enough to enable spikes with small inputs. If it had retained all activity between inputs, it would have spiked too frequently with small inputs.

The LIF models [2] would have faired much better. As noted in section 4.2, fatigue only had a significant impact when there was a sustained rate of high input (i.e. two seconds at 900mV).

Another issue that is relevant to both compartmental and firing models is time. It would be relatively simple to modify the FLIF model to have a finer time grain, e.g. .1 ms. per cycle, but that would increase computational speed.

The Spike Response Model [8, 16] is another model combining thresholding, refractory periods and randomness. As the spike data accounted for seems, at best, weakly effected by refractory periods, it is likely that this model would not perform particularly well on this data set.

There are of course other higher level models of neurons. For example, a model of cell assembly behaviour [17] models the behaviour of sets of neurons. Similarly, there is a theoretical mapping between adaptive resonance theory [5] and group of neuron behaviour. Clearly, these models could not account for the spike data.

A final property should be considered, spontaneous neural activation. In many cases, without input, neurons spike spontaneously [18]. Even during two seconds of inactivity, the biological neurons do not spike in this data. However, this may be part of the reason for the rapid occurrence of the first spike after a period of inactivity. One possibility to account for this, and improve the model is to modify the model so that fatigue could reduce the firing threshold ($F_i < 0$). If no input continued for many seconds, this could induce spontaneous activation.

6 Conclusion

This paper has described a FLIF neural model, discussed some biological neural data, and derived parameters for the FLIF model from that neural data. Thus, the FLIF model with parameters set to $\theta = 2.2$, $D = 1.12$, $F_c = 0.045$, and $F_r = 0.01$ is a relatively faithful model of this particular rat neuron.

The paper has indicated that fatigue can play a useful component of the model. This is the case when there is a significant amount of input for a reasonably long duration. Fatigue may also be used in learning to account for particular psychological phenomena [19].

It should be emphasised that this is a model of a particular neuron. It is almost certain that, even when applicable, the FLIF model of different neurons would require different parameter settings. This paper has only considered one model neuron, though it was the first the author actually tried to model.

None the less, the four free parameters of the model have been set so that they account for 90% of the spikes relatively well. Consequently, it seems that this model is of reasonable biological fidelity.

REFERENCES

[1] L. Abbott, 'Lapicque's introduction of the integrate-and-fire model neuron (1907)', *Brain Research*, **50**, 303–304, (1999).

[2] D. Amit, *Modelling Brain Function: The world of attractor neural networks*, Cambridge University Press, 1989.

[3] R. Belavkin and C. Huyck, 'Conflict resolution and learning probability matching in a neural cell-assembly architecture', *Journal of Cognitive Systems Research*, (accepted).

[4] T. Berger and R. Naud. Quantitative single-neuron modelling, 2009.

[5] G. Carpenter and S. Grossberg, 'The art of adaptive pattern recognition by a self-organizing neural network', *IEEE Computer*, **21**, 77–88, (1988).

[6] M. Chacron, K. Pakdaman, and A. Longtin, 'Interspike interval correlations, memory, adaptation, and refractoriness in a leaky integrate-and-fire model with threshold fatigue', *Neural Computation*, **15**, 253–278, (2003).

[7] P. Dayan and L. Abbott, *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*, MIT Press, 2005.

[8] W. Gerstner and J. vanHemmen, 'Associative memory in a network of spiking neurons', *Network*, **3**, 139–164, (1992).

[9] A. Hodgkin and A. Huxley, 'A quantitative description of membrane current and its application to conduction and excitation in nerve', *J. of Physiology*, **117**, 500–544, (1952).

[10] J. Hopfield, 'Neural nets and physical systems with emergent collective computational abilities', *Proc. of the Nat. Academy of Sciences USA*, **79**, 2554–8, (1982).

[11] C. Huyck, 'Cell assemblies as an intermediate level model of cognition', in *Emerging Neural Architectures based on Neuroscience*, eds., S. Wermter, J. Austin, and D. Willshaw, 383–397, Springer, (2001).

[12] C. Huyck, 'Creating hierarchical categories using cell assemblies', *Connection Science*, **19:1**, 1–24, (2007).

[13] C. Huyck, 'CABot1: a videogame agent implemented in flif neurons', in *IEEE Systems, Man and Cybernetics*, (2008).

[14] C. Huyck, 'A psycholinguistic model of natural language parsing implemented in simulated neurons', *Cognitive Neurodynamics*, (2009).

[15] C. Huyck and H. Ghalib, 'A neuropsychological framework for advancing artificial intelligence', in *AAAI fall symposium on Biologically Inspired Cognitive Architectures*, p. 87, (2008).

[16] R. Jolivet, T. Lewis, and W. Gerstner, 'The spike response model: A framework to predict neuronal spike trains', *Springer Lecture Notes in Computer Science*, **2714**, 846–853, (2003).

[17] S. Kaplan, M. Sontag, and E. Chown, 'Tracing recurrent activity in cognitive elements(trace): A model of temporal dynamics in a cell assembly', *Connection Science*, **3**, 179–206, (1991).

[18] L. Katz and C. Shatz, 'Synaptic activity and the construction of cortical circuits', *Nature*, **274**, 1133–1138, (1996).

[19] L. Kleinsmith and S. Kaplan, 'Paired-associate learning as a function of arousal and interpolated interval', *Journal of Experimental Psychology*, **65:2**, 190–193, (1963).

[20] T. Kohonen, *Self-Organizing Maps*, Springer, 1997.

[21] P. Lansky, P. Sanda, and J. He, 'The parameters of the stochastic leaky integrate-and-fire neuronal model', *Journal of Computational Neuroscience*, **21**, 211–223, (2006).

[22] W. Maas and C. Bishop, *Pulsed Neural Networks*, MIT Press Cambridge, MA, 2001.

[23] H. Markram, 'The blue brain project', *Nature Reviews Neuroscience*, **7**, 153–160, (2006).

[24] W. McCulloch and W. Pitts, 'A logical calculus of ideas immanent in nervous activity', *Bulletin of Mathematical Biophysics*, **5**, 115–133, (1943).

[25] D. Rumelhart and J. McClelland, *Parallel Distributed Processing*, MIT Press, 1986.