Computational Role of Astrocytes as Bayesian Inference Agents in Shaping Neural Networks

Martin Dimkovski and Aijun An

Technical Report EECS-2015-01

February 27 2015

Department of Electrical Engineering and Computer Science
4700 Keele Street, Toronto, Ontario M3J 1P3 Canada

Computational Role of Astrocytes as Bayesian Inference Agents in Shaping Neural Networks

Martin Dimkovski[1], **Aijun An**[1]

[1]York University, Toronto, Canada.

Keywords: Astrocytes, glia, Bayesian neural networks

Abstract

Glia cells are increasingly suspected of having an information processing role in the nervous system, however, it is not clear what their precise role could be. Based on the intracellular Ca^{2+} wave mechanics of astrocytes, we derive a capability of astrocytes to encode information probabilistically, and we present their effect on neural networks as Bayesian inference over synapse parametrization, analogous to Markov Chain Monte Carlo (MCMC) sampling. The proposed framework suggests a Bayesian nature at the cellular level in the neocortex. We also make an argument that astrocytes have a central role in learning for a new behavior, and shaping neural networks for this new behavior.

1 Introduction

Glia cells are increasingly suspected of having an information processing role in the nervous system (Perea et al., 2014; Clarke and Barres, 2013). However, it is not clear what their precise computational role could be. There have been some particularly telling findings in the past several years in terms of behavior and astrocytes, which are the principal type of glia in the neocortex. Han et al. (2013) demonstrated that engrafting human astrocyte progenitor cells in neonatal mice, improves learning and memory. Lee et al. (2014) showed that selectively disabling astrocyte communication without affecting neurons, prevents mice from paying particular attention to only novel object in their environment. In terms of biology, findings have accumulated in the last decade about the bi-directional relationship of astrocytes with neural synapses, prompting the development of the *tripartite synapse* perspective (Volterra et al., 2002; Araque et al., 1999). Clarke and Barres (2013) explain how astrocytes "powerfully control every stage of synapse formation, maturation and elimination and that they "can no longer be thought of as passive support cells".

The change in perspective on astrocytes can be intuitively shared when one considers the updated physical image of astrocytes. Old imaging techniques did not capture the finer arborization of astrocytes and revealed only about 15% of their cell volume, in star-like shapes which gave them their name (Haber and Murai, 2005). New imaging techniques show that what was thought to be a star-like cell, is actually a much larger meshed ball, engulfing millions of synapses in discreet micro-domains, each under the control of a single astrocyte.

Perea et al. (2014) define three aspects required for an information processing role,

and elaborate how astrocytes satisfy all of them: they receive incoming information, integrate and code information in a way that is not a side-effect of neural processing, and transmit the results to other cells. The fact that astrocytes form a bi-directional feedback loop with neural networks, and that neural networks have a known information processing function, also suggests that astrocyte participate in that function.

Based on the intracellular Ca^{2+} wave mechanics of astrocytes, we derive a capability of astrocytes to encode information probabilistically, and we present their effect on neural networks as Bayesian inference over synapse parametrization, analogous to Markov Chain Monte Carlo (MCMC) sampling. The proposed framework suggests a Bayesian nature at the cellular level in the neocortex. We also make an argument that astrocytes have a central role in learning for a new behavior, and shaping neural networks for this new behavior.

Our framework is conjectural because it is based on previous findings and theoretical derivations. However, we believe that there is enough support by now to raise attention, and that the possible implications warrant further verification and improvement by others. By following this framework we developed a model for Bayesian neural networks which has unique advantages for incremental learning and scalability (Dimkovski and An, 2015).

In section 2 we discuss an idea that astrocytes provide ongoing brain rebuilding based on experience. In section 3 we discuss how wave dynamics can encode memory and probability density functions, and present the idea of astrocytes as Bayesian inference agents. Section 4 reviews some related models, and section 5 ends with a conclusion.

2 Shaping Neural Networks For New Experience

In order for astrocytes to have a central role in learning and to shape neural networks based on this role, the following relationships should exist: impairing astrocytes has to impair learning without affecting recognition of existing knowledge; developmental and evolutionary rise of astrocytes has to translate into better learning capacity which cannot be attributed to other improvements such as better neurons; and astrocytes need to be able to affect neural networks and the biochemical and genetic processes which are know to be related to learning.

To begin with, glia create the neurons in all stages of life. During prenatal neurogenesis, radial glia create the scaffolding over which neurons are positioned into columns and layers (Ma et al., 2005). Radial glia are also the progenitor cells for cortical neurons and astrocytes during this period. In adulthood, radial glia transform into astrocyte-like neural stem cells that provide adult neurogenesis (Corty and Freeman, 2013). Seri et al. (2001) identifies the neural stem cells more closely to astrocytes adult forms, and suggest that astrocytes could also provide adult neurogenesis in the neocortex.

In addition to creating the neurons, glia also induce the creation of synapses and their removal from the network (Allen et al., 2012; Corty and Freeman, 2013; Haydon, 2006). In perinatal development, radial glia initiate synapses by forming transient neuron-glia synapses which result in neuron-neuron synapses later. Adult astrocytes can remodel the neuropil in the hippocampus in a matter of minutes (Haber and Murai, 2005). As mentioned earlier, astrocytes partition the neuropil into functional islands, called micro-domains, with an overlap of less than 5% (Nedergaard et al., 2003a). Nakae et al. (2014) found through visual inspection that neurons also had a tendency to

be under the domain of a unique glia cell. In summary, neural networks in the neocortex can be seen as a three-dimensional array of discrete domains (Haber and Murai, 2005) which are shaped and reshaped by astrocytes.

Astrocytes are not just passive neural network builders, because they form a feedback loop with the tripartite synapses under their micro-domains. For example, the glutamate released by the presynaptic action potential also affects metabotropic glutamate receptors on the astrocyte membrane which wraps the synapse, and through a chain chemical reaction causes puffs of Ca^{2+} released from internal cellular stores. These puffs can accumulate into Ca^{2+} waves across the whole astrocytes and even into neighboring astrocytes through gap junctions. The Ca^{2+} waves, in turn, can trigger release of neurotransmitters from inside the astrocyte into the synapse, affecting both pre- and post-synaptic receptors (Wade et al., 2011). Other feedback of Ca^{2+} waves include effects on spike-timing-dependent plasticity (ref), regulation of gene expression over months or years (Thul et al., 2009), and secretion of thrombospondins which induce neurons to form synapses (Barres, 2008). Therefore, through the astrocyte, activity of one synapse can considerably affect many remote synapses of the same or other neurons.

The communication that happens across the tripartite synapse is not merely a side effect of information processing in neural networks, because the information flow inside the astrocyte has considerably different spatial and temporal characteristics. While neurons communicate in a point-to-point manner, with a specific peer, astrocyte communication is of a broadcast nature. In addition, neural communication is on a scale of milliseconds, where as intracellular astrocyte communication is on a scale of minutes,

in some cases in excess of 15 minutes (Araque et al., 2001). The effect of astrocytes in the tripartite synapse is also of a longer temporal scale compared to neural synapse effects. For example, astrocytes cause "a slow inward current (SIC), which has a rise time of 60 ms and a decay time of 600 ms, and is thus very different from the classical excitatory postsynaptic current (EPSC) (6.4 ms and 10 ms rise and decay time respectively) (Wade et al., 2011).

Various findings suggest that the unique information processing of astrocytes is related specifically to learning. Lee et al. (2014) found that using a toxin, which selectively disables astrocyte communication without affecting neurons, prevents mice from spending extra time with a novel object in its environment as they would normally, without affecting their behavior with familiar objects. Miranda et al. (2011) found that experimental hampering of the calcium cellular mechanisms, crucial to glia, does not affect memory retrieval, only memory formation. Another revealing finding is that the effect of marijuana on memory is through astrocytes, not neurons (Han et al., 2012). Mller and Best (1989) showed that injection of immature astrocytes into the visual cortex of adult cats in vivo reopens the window of ocular dominance plasticity. Markham and Greenough (2005) found that "increased astrocytic volume can be inferred to arise in association with learning-specific synaptogenesis, and not merely constitute a response to a general increase in neural activity. A suggestion that glia's role is not in recognition of known knowledge is the fact that such recognition happens on the order of hundreds of milliseconds (Thorpe et al., 1996), which is too fast for the glia dynamics. Han et al. (2013) engrafted human glia progenitor cells in neonatal mice, which later developed in hominid glia in the mature mouse. A hominid astrocyte is consid-

erably larger and more complex than murine astrocyte, and covers 2 millions synapses in its territorial domain, compared to the 20,000 to 120,000 synapses of a murine astrocyte. In addition hominid astrocytes communicate faster. Mice with hominid glia showed improved learning and memory, assessed by Barnes maze navigation, object-location memory, and both contextual and tone fear conditioning.

The learning function of astrocytes has been related specifically to long-term memory. Orr et al. (2015) found that tampering with a certain receptor on astrocytes affected only long-term memory formation, while other functions remained unaffected. Chen et al. (2012) also show that disabling astrocytes prevents only long-term memory formation. As discussed above, it is astrocyte-like cells, if not astrocytes themselves, that enable adult neurogenesis. The two places that are widely acknowledged to host adult neurogenesis are the subgranular zone and the subventricular zone (Grabel, 2012), and in both the newly produced neurons are implicated in long-term memory formation (Sultan et al., 2011).

Glia have been associated with less stereotypical and more challenging behavior. In their study of C. Elegans, Oikonomou and Shaham (2011) find it "...striking that only C. Elegans head muscles form partnerships with glia. These muscles mediate fine motor behaviors that are less stereotypical than the undulations produced by body wall muscles, perhaps explaining the need for glial companionship." Hartline (2011) notes that between the evolutionary branches of some of the earliest taxa to develop glia, or even between stages of their metamorphosis, loss of glia is related to less active lifestyle. For example, urochordates have glia in the roaming (larva) stage but lose it in the sedentary adult phase. The author does point that the evidence is not equivocal. Bacaj et al.

(2008) found that glia-ablated C.Elegans fails to adjust its behavior and migrate to its cultivation temperature, and it does not do chemotaxis or long-range avoidance, even though its sensory and basic movement capability remains normal. Chemotaxis and long-range avoidance require the ability to adjust behavior depending on whether the situation is improving. Markham and Greenough (2005) note that astrocytes ensheath synapses more in response to a complex environment. Evolutionary, the ratio of astrocytes to neurons increases progressively with a perceived capacity for intelligence. The exact ratios vary from brain region to region and are difficult to estimate, but some example figures show it increases from about 0.05 in leeches, 0.2 in frogs, 0.4 in rats, 1.2 in cats, to 1.5 in humans (Nedergaard et al., 2003b).

The above findings support the requisite relationship we identified at the beginning of the section: astrocyte create and change neural networks, impairing astrocytes impairs learning for new experience without affecting behavior based on previous experience or genetic inheritance, and addition of glia (evolutionary, engrafting, or developmental) results in improved adaptability. Astrocytes appear to provide ongoing brain rebuilding based on experience, continuing the work of their perinatal progenitors (such as radial glia) which initially build the neocortex. In other words, astrocytes appear to continuously shape neural networks for newly learned behavior. A similar proposal was also made by Markham and Greenough (2005), where based on histological studies they ”speculate that astrocytic changes might be necessary to induce, but not to maintain, adaptive changes in the brains wiring diagram in response to experience.

3 Ca^{2+} Waves, Synapses, and Bayesian inference

In order for astrocytes to learn within their micro-domains, the crucial Ca^{2+} wave dynamics inside the astrocytes has to be related to ability to encode and process information. De Pitt et al. (2009) and Wade et al. (2011) present studies on such ability by combining various dynamical system models of neurons, Ca^{2+} waves, and synapses. Using their composite models, they describe information encoding capability in the Ca^{2+} wave mechanics in terms of bifurcations, amplitude and frequency modulation, and oscillatory properties. For example, Wade et al. (2011) show that activity from one synapse can induce plasticity at other remote synapses on the astrocyte. These studies, however, simulate a few units and few synapses due to their complexity, and are not very informative about macro computational properties.

There are some intuitive examples in physics which link wave mechanics to information encoding. Eddi et al. (2011) show one such example with a bouncing droplet coupled to a vibrating fluid surface. When the vibrating amplitude becomes comparable to gravity, the droplet can bounce indefinitely. If the vibrating amplitude approaches the Faraday instability threshold the droplet couples to a pilot" wave and starts moving with it. With every bounce, the droplet causes ripples, which interfere with ripples from previous activity and create a path memory. At the point of next contact, the path memory in turn reads out" the droplets next movement. The authors point out the interpretation in terms of information encoding: "The dual nature ... is contained in the path memory dynamics: the wave nature lies in the coding while the particle nature lies in the reading.

Such a memory functionality, based on wave mechanics with a coupled point process, has been found also in other domains. In a study of crack propagation in physical

medium, Goldman et al. (2010) note how the crack tip progression can be dominated by memory from superimposed elastic waves it caused in the past, reflected from boundaries and inhomogeneous zones.

If a parallel is drawn with astrocytes, then a tripartite synapse can be described as coupled point process to the Ca^{2+} wave dynamics inside the astrocyte. In such a perspective, the Ca^{2+} waves can be seen as coding information based on perturbations from connected synapses, while Ca^{2+} wave effect on the synapse can be seen as a read-out of the coded information. The longer time frame and cumulative nature of Ca^{2+} waves allows for many synapse interaction to accumulate in the memory it represents.

The above perspectives for information coding and decoding with wave mechanics, provide valuable intuition. However, they are based fully on deterministic wave modeling, which is done through simulation of differential equations for minute physical properties. Seeking a more practical approach to modeling computation, we turn to a probabilistic interpretation of wave dynamics. A well-studied example of that nature is the Copenhagen interpretation of quantum mechanics, where Schroedingers wave equation is interpreted probabilistically, by relating it to the concept of a probability current, derived from fluid mechanics.

In fluid mechanics the concept of *current* **j** is described by the continuity equation:

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot \mathbf{j}, \tag{1}$$

which is characteristic when there is some physical quantity Q which moves continuously and is conserved (Kroemer, 1994). Known types of Q include dissolved ions, which is the case of Ca^{2+} ions inside the astrocyte. ρ is the density of Q, i.e. the quantity in a unit volume, t is the time of change, **j** is the current (sometime called *flux*)

which a vector field which tells how much Q passed through a unit area in the cross-section perpendicular to $\mathbf{j}$ in a unit time, and ∇ is the divergence of $\mathbf{j}$. In summary, $\mathbf{j}$ describes how density changes in a unit volume.

The probability of finding a Ca^{2+} ion in a specific location is proportional to its density ρ in that location. In addition, probability is also conserved because its integral is always 1. Therefore, a *probability current* can be defined in the same form as Eq.(1), as follows:

$$\frac{\partial P}{\partial t} = \frac{\alpha(\partial \rho)}{\partial t} = -\alpha(\nabla \cdot \mathbf{j}_\rho) = -\nabla \cdot \mathbf{j}_P, \tag{2}$$

since divergence is a linear operator. α is the proportionality constant, and $\mathbf{j}_P$ is the probability current.

Given the probability current we can get the probability by integrating across a cross-section and time (Kroemer, 1994). Therefore, the probability P_A that a Ca^{2+} ion inside the astrocyte affects a specific surface on the astrocyte membrane ΔS during Δt is:

$$P_A = \int_{\Delta t} \left[\int_{\Delta S} \mathbf{j}_P \cdot dS \right] dt. \tag{3}$$

Eq.(3) above explains how changes in density of Ca^{2+} ions, caused by activity of the tripartite synapses, can translate into a probability density function P_A. P_A represents wave memory encoded by the astrocyte due to historical activity in its associated neural network. Because a tripartite synapse interface within ΔS enacts the information reading of wave memory, a synapse is in effect sampled from P_A. Therefore, if the neural network in the astrocyte micro-domain is parametrized by a set of synapses $\mathbf{s} = \{s_1, s_2, \dots\}$, and the data processed by this neural network is $\mathbf{D}$, then $P(\mathbf{s}|\mathbf{D}) \equiv P_A$. In summary, the processing of data $\mathbf{D}$ builds P_A over time, which then parametrizes the

current neural network as $P(\mathbf{s}|\mathbf{D})$.

In order to use Eq.(3) in that form we need to be able to get the probability current $\mathbf{j}_P$ of the Ca^{2+} waves. This can theoretically be accomplished with any mathematical model for $\partial Ca^{2+}/\partial t$, because Ca^{2+} is proportional to its density ρ, and using a derivation similar to Eq.(2) we can get, with some other proportionality constant β:

$$\frac{\partial Ca^{2+}}{\partial t} = -\beta(\nabla \cdot \mathbf{j}_P). \tag{4}$$

Doing inverse divergence to get j_P from Eq.(4) is not trivial. We use the $\partial Ca^{2+}/\partial t$ models only borrow support for conditional independence of the tripartite synapses, which then lets us frame a more practical framework. In particular, Wade et al. (2011) use a $\partial Ca^{2+}/\partial t$ model which describes the change of Ca^{2+} levels inside an astrocyte and around the interface to a single tripartite synapse s_i, and describes how individual puffs of Ca^{2+}, around the synapse interface, remain mostly self-contained, and how their intra-cellular propagation effect is caused by cascade activation of puffs in neighboring areas. Therefore, we use Eq.(3) and Eq.(4) as the foundation for a conditional probability distribution $P(s_i|s_{-i}, \mathbf{D})$ of a single synapse given the state of all other synapses and the processed data.

The order in which synapses are updated is determined by how the Ca^{2+} wave interacts with the astrocyte membrane on the inside. The update order can be described as stochastic, since at worst it is pseudorandom due to the chaotic interaction patterns of the wave and the membrane, or it could be random due to randomness in the biological processes that underpin it. For simplicity, we can assume a single synapse updating at one time, by considering an arbitrary level of precision in measuring the update time. This view can be extended to blocks of synapses updating simultaneously.

The effect of the astrocyte on its tripartite synapses can now be described as continually updating them one by one in a stochastic order, by sampling the parametrization of each synapse s_i from $P(s_i|s_{-i}, \mathbf{D})$ during each update. Since each update conditions on the current state of all other synapses, the states of all synapses are continually interpolated across the astrocyte. This update process is equivalent to Gibbs sampling (Gelman, 2014), which is a Bayesian inference method which guarantees that the sampled synapse states, each from its conditional distribution $P(s_i|s_{-i}, \mathbf{D})$, will all together asymptotically converge to the joint distribution $P(\mathbf{s}|\mathbf{D})$, regardless of the initial states of the synapses. At the core of this Bayesian inference is the relationship:

$$P(\mathbf{s}|\mathbf{D}) \propto P(\mathbf{s})P(\mathbf{D}|\mathbf{s}), \tag{5}$$

where the prior and posterior are based to the astrocyte and the likelihood in the neural network.

A biologically plausible framework must describe how continually streaming data from the environment are incrementally processed, i.e. it must not require the explicit storage of all the data seen. The wave memory encoding discussed earlier integrates previous data into the Ca^{2+} current $\mathbf{j}$, i.e. at time $t-1$, $\mathbf{j}^{(t-1)}$represents the effect of all $\mathbf{D}^{(t-1)} = \{D_1, D_2, \ldots, D_{t-1}\}$, and can be thought of a surrogate sufficient statistics for them. The difference between $P_A^{(t)}$ and $P_A^{(t-1)}$ as per Eq.(3) should be describable using only D_t. Therefore, any model implementing the proposed glia framework needs a Bayesian inference setup where only D_t is used for the likelihood. In the case of Gibbs sampling this means we would need a conditional probability distribution of the form $P(s_i|s_{-i}, D_t)$ instead of $P(s_i|s_{-i}, \mathbf{D})$. Guhaniyogi et al. (2014) show how standard Gibbs sampling, which requires all the previous data, can be modified for

incremental use by integrating sufficient statistics $\xi^{(t-1)}$ from previous data $\mathbf{D}^{(t-1)}$ into $P(s_i|s_{-i}, D_t, \xi^{(t-1)})$. In such cases, an empirical Bayesian relationship can then be expressed as:

$$P(s_i|s_{-i}, D_t, \xi^{(t-1)}) \propto P(s_i|s_{-i}, \xi^{(t-1)})P(D_t|\mathbf{s}). \tag{6}$$

This equation implies several points: that the astrocyte encodes internally an empirical prior for the neural network parametrization, learned from previous data; that the likelihood needed for inference is contributed by the neural network activity $P(D_t|\mathbf{s})$ over the tripartite synapses into the astrocyte; and that the result of the Bayesian inference acts back on the neural network over the tripartite synapse as parametrization samples from the posterior.

In Gibbs sampling implementations, the prior and the likelihood are not explicit because they are normally subsumed through conjugate relationships into the posterior which is directly sampled. In our implementation of this glia framework we used feed-forward neural networks with sigmoidal McCullochPitts neurons, whose likelihood function does not offer obvious conjugate setup needed for Gibbs sampling. Instead, we extended Metropolis-Hastings (Gelman, 2014), which is a more general form of Gibbs sampling, to an incremental form.

Our Incremental Metropolis-Hastings (IMH) gives an incremental MCMC which recursively uses the previous posterior as a new prior through kernel density estimations, and only the last data D_t. IMH is a standalone Bayesian inference method, published in a separate paper (Dimkovski and An, 2015) where we do not talk about glia, but presents only computational properties and performance results on machine learning datasets. Similarly, the glia framework we propose here does not depend on IMH as

it could be implemented in other ways, and in various proportions between practicality and biological fidelity. For example, using more realistic spiking neurons might extend work such as (Rao, 2004; Deneve, 2005) to further develop the framework.

4 Related Work

All the models we are aware of are deterministic and without Bayesian aspects. Some models do not present astrocytes as encoding information. Porto-Pazos et al. (2011) add to each neuron one astrocyte unit which activates when the neurons connection are highly active and then causes additional and gradual increase of the neurons connection weights over a longer period (4-8 iterations), or if the astrocyte is not active it decreases the weights gradually. In effect, this model equates astrocytes with long-term potentiation (LTP) and depression (LTD) of a transient type (without permanent changes). Ikuta et al. (2011) similarly add to each neuron one astrocyte unit which activates when the neuron is highly active, and spreads an impulse throughout the network stimulating neurons additionally, in a form of broadcast LTP/LTD, also transient, acting on remote neurons in the network.

We believe that a simple and transient LTP/LTD role for astrocytes is overly simplified, and that any LTP/LTD related to astrocytes, especially late phase type (permanent), is an outward, local, and partial manifestation of the long-term memory encoding done by astrocytes. LTP/LTD was defined from a neuron-only perspective, before the role of astrocytes was being considered. Wallace and Bluff (1995) discusses shortcomings of a neuron-centric LTP, and discusses, in terms of biochemistry, how astrocytes can

broaden the idea of LTP for multiple and longer time scales, and selective action on different neurons.

Models which consider astrocytes as encoding information, approach the problem through the internal wave dynamics of astrocytes. Reid and Barrett-Baxendale (2008) propose that information is encoded through Ca^{2+} wave propagation and interference patterns triggered by neural activity, and that it could act in return on the neural network to change the post-synaptic potential and refractory period, and add new synapses. However a formal complete model is left as future work.

5 Conclusion

Tenenbaum et al. (2011) note how Bayesian inference quickly becomes a suspect for explaining how the brain works, when we consider how it learns and generalizes too fast and beyond what only the latest relevant data allow for, and they point to the need for some "abstract background knowledge" which must be present at all time. The challenges for Bayesian brain theories are in explaining how the involved probability distributions are learned, stored, and utilized. Existing theories often use single-point approximations to probability distributions (such as MAP), which is a problem when learning with recent data only. The principal challenge however is explaining the low-level physical foundation. Existing theories mostly address this question on higher conceptual levels, such as the *computational* and *algorithmic* Marr's (1982) levels. Few theories tackle the question on the *physical* level in terms of cellular processes. Tenenbaum et al. (2011) note that uncovering a physical basis of Bayesian inference in the

brain remains an open challenge.

Our Bayesian brain framework is rooted as low as the intracellular Ca^{2+} waves in astrocytes and the biochemistry of tripartite synapses, and it models probability distributions holistically through MCMC samples, not though single-point estimates. The framework proposes how Bayesian priors can be encoded at the cellular level, and serve as the inductive bias which continually biases neural networks towards a background knowledge based on long-term memory. Such biasing can facilitate detection of novelty, when the immediate stimulus deviates from the background knowledge. A key point is that the background knowledge can be updated through Bayesian inference, which suggests that astrocytes enable an organism to prosper in challenging and changing environments, by providing brain rebuilding into adulthood based on experience. We also show that the proposed framework can be abstracted in forms which are practical for machine learning, which do not require complicated biophysical modeling.

References

Allen, N. J., Bennett, M. L., Foo, L. C., Wang, G. X., Chakraborty, C., Smith, S. J., and Barres, B. A. (2012). Astrocyte glypicans 4 and 6 promote formation of excitatory synapses via GluA1 AMPA receptors. *Nature*.

Araque, A., Carmignoto, G., and Haydon, P. G. (2001). Dynamic signaling between astrocytes and neurons. *Annual Review of Physiology*, 63(1):795–813.

Araque, A., Parpura, V., Sanzgiri, R. P., and Haydon, P. G. (1999). Tripartite synapses: glia, the unacknowledged partner. *Trends in Neurosciences*, 22(5):208–215.

Bacaj, T., Tevlin, M., Lu, Y., and Shaham, S. (2008). Glia Are Essential for Sensory Organ Function in C. elegans. *Science*, 322(5902):744–747.

Barres, B. (2008). The Mystery and Magic of Glia: A Perspective on Their Roles in Health and Disease. *Neuron*, 60(3):430–440.

Chen, N., Sugihara, H., Sharma, J., Perea, G., Petravicz, J., Le, C., and Sur, M. (2012). Nucleus basalis-enabled stimulus-specific plasticity in the visual cortex is mediated by astrocytes. *Proceedings of the National Academy of Sciences*, 109(41):E2832–E2841.

Clarke, L. E. and Barres, B. A. (2013). Emerging roles of astrocytes in neural circuit development. *Nature Reviews Neuroscience*, 14(5):311–321.

Corty, M. M. and Freeman, M. R. (2013). Cell biology in neuroscience: Architects in neural circuit design: Glia control neuron numbers and connectivity. *The Journal of Cell Biology*, 203(3):395–405.

De Pitt, M., Goldberg, M., Volman, V., Berry, H., and Ben-Jacob, E. (2009). Glutamate regulation of calcium and IP3 oscillating and pulsating dynamics in astrocytes. *Journal of Biological Physics*, 35(4):383–411.

Deneve, S. (2005). Bayesian inference in spiking neurons. *Advances in neural information processing systems*, 17:353–360.

Dimkovski, M. and An, A. (2015). A Bayesian model for canonical circuits in the neocortex for parallelized and incremental learning of symbol representations. *Neurocomputing*, 149:1270–1279.

Eddi, A., Sultan, E., Moukhtar, J., Fort, E., Rossi, M., and Couder, Y. (2011). Information stored in Faraday waves: the origin of a path memory. *Journal of Fluid Mechanics*, 674:433–463.

Gelman, A. (2014). *Bayesian data analysis*. Chapman & Hall/CRC texts in statistical science. CRC Press, Boca Raton, third edition edition.

Goldman, T., Livne, A., and Fineberg, J. (2010). Acquisition of Inertia by a Moving Crack. *Physical Review Letters*, 104(11).

Grabel, L. (2012). Developmental Origin of Neural Stem Cells: The Glial Cell That Could. *Stem Cell Reviews and Reports*, 8(2):577–585.

Guhaniyogi, R., Qamar, S., and Dunson, D. B. (2014). Bayesian Conditional Density Filtering. *arXiv preprint arXiv:1401.3632*.

Haber, M. and Murai, K. K. (2005). Reshaping neuronglial communication at hippocampal synapses. *Neuron Glia Biology*, 2(01):59.

Han, J., Kesner, P., Metna-Laurent, M., Duan, T., Xu, L., Georges, F., Koehl, M., Abrous, D. N., Mendizabal-Zubiaga, J., and Grandes, P. (2012). Acute cannabinoids impair working memory through astroglial CB 1 receptor modulation of hippocampal LTD. *Cell*, 148(5):1039–1050.

Han, X., Chen, M., Wang, F., Windrem, M., Wang, S., Shanz, S., Xu, Q., Oberheim, N., Bekar, L., Betstadt, S., Silva, A., Takano, T., Goldman, S., and Nedergaard, M. (2013). Forebrain Engraftment by Human Glial Progenitor Cells Enhances Synaptic Plasticity and Learning in Adult Mice. *Cell Stem Cell*, 12(3):342–353.

Hartline, D. K. (2011). The evolutionary origins of glia. *Glia*, 59(9):1215–1236.

Haydon, P. G. (2006). Astrocyte Control of Synaptic Transmission and Neurovascular Coupling. *Physiological Reviews*, 86(3):1009–1031.

Ikuta, C., Uwate, Y., and Nishio, Y. (2011). Performance and features of multi-layer perceptron with impulse glial network. In *Neural Networks (IJCNN), The 2011 International Joint Conference on*, pages 2536–2541. IEEE.

Kroemer, H. (1994). *Quantum Mechanics: For Engineering, Materials Science, and Applied Physics*. Prentice Hall.

Lee, H. S., Ghetti, A., Pinto-Duarte, A., Wang, X., Dziewczapolski, G., Galimi, F., Huitron-Resendiz, S., Pina-Crespo, J. C., Roberts, A. J., Verma, I. M., Sejnowski, T. J., and Heinemann, S. F. (2014). Astrocytes contribute to gamma oscillations and recognition memory. *Proceedings of the National Academy of Sciences*, 111(32):E3343–E3352.

Ma, D. K., Ming, G.-l., and Song, H. (2005). Glial influences on neural stem cell development: cellular niches for adult neurogenesis. *Current Opinion in Neurobiology*, 15(5):514–520.

Markham, J. A. and Greenough, W. T. (2005). Experience-driven brain plasticity: beyond the synapse. *Neuron Glia Biology*, 1(04):351.

Miranda, M. I., Gonzlez-Cedillo, F. J., and Daz-Muoz, M. (2011). Intracellular calcium chelation and pharmacological SERCA inhibition of Ca2+ pump in the insular cortex

differentially affect taste aversive memory formation and retrieval. *Neurobiology of Learning and Memory*, 96(2):192–198.

Mller, C. M. and Best, J. (1989). Ocular dominance plasticity in adult cat visual cortex after transplantation of cultured astrocytes.

Nakae, K., Ikegaya, Y., Ishikawa, T., Oba, S., Urakubo, H., Koyama, M., and Ishii, S. (2014). A Statistical Method of Identifying Interactions in NeuronGlia Systems Based on Functional Multicell Ca2+ Imaging. *PLoS Computational Biology*, 10(11):e1003949.

Nedergaard, M., Ransom, B., and Goldman, S. A. (2003a). New roles for astrocytes: Redefining the functional architecture of the brain. *Trends in Neurosciences*, 26(10):523–530.

Nedergaard, M., Ransom, B., and Goldman, S. A. (2003b). New roles for astrocytes: Redefining the functional architecture of the brain. *Trends in Neurosciences*, 26(10):523–530.

Oikonomou, G. and Shaham, S. (2011). The Glia of Caenorhabditis elegans. *Glia*, 59(9):1253–1263.

Orr, A. G., Hsiao, E. C., Wang, M. M., Ho, K., Kim, D. H., Wang, X., Guo, W., Kang, J., Yu, G.-Q., Adame, A., Devidze, N., Dubal, D. B., Masliah, E., Conklin, B. R., and Mucke, L. (2015). Astrocytic adenosine receptor A2a and Gs-coupled signaling regulate memory. *Nature Neuroscience*.

Perea, G., Sur, M., and Araque, A. (2014). Neuron-glia networks: integral gear of brain function. *Frontiers in Cellular Neuroscience*, 8.

Porto-Pazos, A. B., Veiguela, N., Mesejo, P., Navarrete, M., Alvarellos, A., Ibez, O., Pazos, A., and Araque, A. (2011). Artificial Astrocytes Improve Neural Network Performance. *PLoS ONE*, 6(4):e19109.

Rao, R. P. (2004). Hierarchical Bayesian inference in networks of spiking neurons. In *Advances in neural information processing systems*, pages 1113–1120.

Reid, D. and Barrett-Baxendale, M. (2008). Glial Reservoir Computing. pages 81–86. IEEE.

Seri, B., Garca-Verdugo, J. M., McEwen, B. S., and Alvarez-Buylla, A. (2001). Astrocytes give rise to new neurons in the adult mammalian hippocampus. *The Journal of Neuroscience*, 21(18):7153–7160.

Sultan, S., Rey, N., Sacquet, J., Mandairon, N., and Didier, A. (2011). Newborn Neurons in the Olfactory Bulb Selected for Long-Term Survival through Olfactory Learning Are Prematurely Suppressed When the Olfactory Memory Is Erased. *Journal of Neuroscience*, 31(42):14893–14898.

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., and Goodman, N. D. (2011). How to Grow a Mind: Statistics, Structure, and Abstraction. *Science*, 331(6022):1279–1285.

Thorpe, S., Fize, D., and Marlot, C. (1996). Speed of processing in the human visual system. *nature*, 381(6582):520–522.

Thul, R., Bellamy, T. C., Roderick, H. L., Bootman, M. D., and Coombes, S. (2009). Calcium oscillations. In *Cellular Oscillatory Mechanisms*, pages 1–27. Springer.

Volterra, A., Magistretti, P. J., and Haydon, P. G., editors (2002). *The tripartite synapse: glia in synaptic transmission*. Oxford University Press, Oxford ; New York.

Wade, J. J., McDaid, L. J., Harkin, J., Crunelli, V., and Kelso, J. A. S. (2011). Bidirectional Coupling between Astrocytes and Neurons Mediates Learning and Dynamic Coordination in the Brain: A Multiple Modeling Approach. *PLoS ONE*, 6(12):e29445.

Wallace, J. G. and Bluff, K. (1995). Should ANN be ANGN? In *From Natural to Artificial Neural Computation*, pages 53–60. Springer.