Journal of Computational Neuroscience (2023) 51:463–474
https://doi.org/10.1007/s10827-023-00860-0

RESEARCH

A high-efficiency model indicating the role of inhibition in the resilience of neuronal networks to damage resulting from traumatic injury

Brian L. Frost[1] · **Stanislav M. Mintchev**[2]

Received: 4 November 2022 / Revised: 2 June 2023 / Accepted: 17 August 2023 / Published online: 26 August 2023

Abstract
Recent investigations of traumatic brain injuries have shown that these injuries can result in conformational changes at the level of individual neurons in the cerebral cortex. Focal axonal swelling is one consequence of such injuries and leads to a variable width along the cell axon. Simulations of the electrical properties of axons impacted in such a way show that this damage may have a nonlinear deleterious effect on spike-encoded signal transmission. The computational cost of these simulations complicates the investigation of the effects of such damage at a network level. We have developed an efficient algorithm that faithfully reproduces the spike train filtering properties seen in physical simulations. We use this algorithm to explore the impact of focal axonal swelling on small networks of integrate and fire neurons. We explore also the effects of architecture modifications to networks impacted in this manner. In all tested networks, our results indicate that the addition of presynaptic inhibitory neurons either increases or leaves unchanged the fidelity, in terms of bandwidth, of the network's processing properties with respect to this damage.

Keywords Focal axonal swelling · Integrate-and-fire network · Network adaptation

1 Introduction

Focal axonal swelling (FAS) is an abnormality that sometimes develops in neurons subjected to physical stress Maia and Kutz (2014). In recent years, as medical science has taken up the goal of understanding the various mechanisms by which concussions affect the brain, progressively more accurate microscope imaging technology has shown that traumatic brain injury can lead to structural deformations in neural axon morphology within regions of affected brain tissue (Johnson et al., 2013; Tagge et al., 2018). Specifically, impact trauma can cause the severance of microtubules involved in various cell transport mechanisms (Maxwell et al., 1997; Wu et al., 2021), resulting in material pile-ups at points along the length of the cell axon that cause widening of the cell body (Tang-Schomer et al., 2012); Wang et al., 2011). As a consequence the electrical impulse transmission capabilities of the neuron are affected significantly Maia and Kutz (2013).

From a biophysical perspective, the impact of axonal swelling on electrical conduction from soma to dendrite is likely due to the fact that the compromised cell body acts as a variable-(rather than constant-)width cable. Investigations into the phenomenology arising from such altered morphology have been carried out in, e.g., (Debanne et al., 2011; Maia et al., 2015; Maia & Kutz, 2013; Manor et al., 1991; Ramón et al., 1975). Specifically, Maia and Kutz (2013) presents a detailed mathematical study by way of numerical solution to variable-width cable PDEs, wherein the authors have observed a variety of effects on sequences of traveling wave packets traversing this sort of non-homogeneity while propagating diffusively. In particular, this study has indicated that such sequences are subject to a trichotomy of consequences including partial packet loss, total transmission blockage, and partial reflection; the specifics of each case depend on the parameter regime for the selected model.

Action Editor: David Golomb

✉ Brian L. Frost
b.frost@columbia.edu

Stanislav M. Mintchev
mintchev@cooper.edu

[1] Electrical Engineering, Columbia University, 500 W 120th St, New York, NY 10027, USA

[2] Mathematics, The Cooper Union, 41 Cooper Sq, New York, NY 10003, USA

Of the three categories of effects discussed in Maia and Kutz (2013), partial packet loss presents a compelling candidate for further study since it indicates that a swollen axon may behave as a non-linear frequency filter in the sense of signal processing. A systematic (numerical/computational) investigation of the nature of the packet loss using the aforementioned PDE model reveals a highly selective deletion scheme that is at least somewhat dependent on the time-sequencing of the packets. For example, there are scenarios in which a pair of closely spaced packets will transmit across the cable deformation reliably, but a triple with identical time spacing will see its third spike deleted Maia and Kutz (2017) ; these full-detail studies at the cellular level have led to the development of a comprehensive albeit artificial computational framework Maia et al. (2019) for mimicking a dilated axon's signal filtration effects on spiketrain sequences. The recorded phenomenology suggests a degree of complexity that merits investigation both from the perspective of rigorous analysis at the single-axon level and from the point of view of impact on network activity and information processing.

The present article aims to address how network-level properties (both cell and connection type, as well as connectivity and strength of connections) may endow a subpopulation of neurons with greater (resp., lesser) resilience in the presence of damage due to focal axonal swelling. Damage of the type described is modeled via decreased fidelity of the connections, which is exhibited in the form of transmuted temporal spiking sequences from presynaptic cells.

Various studies of the consequences of concussive damage at the network level have been carried out recently; see, e.g., Sharp et al. (2014) for a thorough overview of the seminal literature. As it relates to the focus of the present discussion, Rudy et al. (2016) used the FORCE algorithm of Sussillo and Abbott (2009) to establish a measure of tolerance to damage that shows a gradual decline in network plasticity-recovery capabilities as the proportion of damaged connections is increased. Notably, this analysis relied on a firing rate model (based on assumptions regarding the effects of axonal swelling damage on the spike-emission rates of affected cells) implemented for a spiking neuronal network tasked with image processing and decision making. Subsequent work in Lusch et al. (2018) uses a similar paradigm by leveraging a convolutional neural network approach. While illuminating, this line of work carries the drawback of potentially overlooking the specifics of neurobiology and detailed axon signal-transduction dynamics at the single-cell level. By appealing directly to and artificially imposing a value for the firing rate parameter; this *en gros* approach skips over the mechanisms at work within axons that have incurred physical defects, in favor of observables that measure a neuronal population's global network activity Ermentrout and Terman (2010). These remarks notwithstanding, a careful critical examination of the work to date on this topic reveals a significant computational challenge: to full detail, numerical investigation of this sort of network would involve intensive PDE simulations for each damaged axon in tandem with a resource-intensive neuronal network implementation. The complexity of such a study renders it prohibitive except in the case of very small networks over short time periods, a severe limitation if one has the objective of scaling to a size that would allow for modeling compromised neural tissue within, e.g., the mammalian nervous system Herculano-Houzel (2009).

Instead of working with an *a priori* firing rate model, which averages over the details of dynamics at the cellular/node level, our study extracts firing rate data as an observable *collected from* a specific output site within a spiking network of integrate-and-fire neurons Gerstner and Kistler (2002). The use of a simple one-dimensional integrate-and-fire model at the cell level allows for an efficient simulation of the spiking network that retains the basic features of voltage evolution and depolarization without slowing down the numerics Izhikevich (2004). We aim to investigate especially the manner in which the precise voltage dynamics of individual neurons influence the overall impact of traumatic stress, on the way to understanding whether specific detailed adaptations in network architecture can mitigate this impact. Attention has been given recently to network degeneracy and its potential to increase the robustness of network processing capabilities, with both positive and negative consequences of this effect (Stöber et al., 2023; Kamaleddin, 2021). We note in particular that brain trauma in the form of *Fluid Percussive Injuries (FPI)* has been associated with a *reduction* in inhibitory synaptic inputs Neuberger et al. (2017), a finding that is consistent with the results that we present in this paper.

To circumvent the simulation of the full axon PDE diffusion model proposed by Maia and Kutz, we implement a high-accuracy predictive machine learning algorithm that allows the recovery of a post-damage time series of spikes given precise knowledge of the input spike train. A statistical analysis shows that this algorithm is extremely accurate; in particular, sufficiently so to warrant its implementation *entirely instead of* the non-linear PDE model. This advance speeds up the numerics significantly and presents a data-driven alternative to the heuristic deletion formula posited in Maia et al. (2019); our spike deletion rule is trained directly from the continuous cable equation due to Hodgkin and Huxley; it performs favorably especially in light of its direct relation to the biological model -and- its low computational cost, and so we are able to leverage it in network-level simulations. While the networks we have investigated are only of modest size, the methodology proposed herein can be adapted strategically to scale for larger networks, more sophisticated connection topology, and features at the

cellular level that are ignored in simple integrate-and-fire schemes. We have proposed and partially investigated one such amenable extension in network size via the *layered network* formalism (Vogels, 2005; Vogels et al., 2005; Vogels & Abbott, 2007).

A fundamental challenge in the present framework is the determination of a measure that serves as a reliable indicator of a network's task performance capabilities. For neurons subjected to stochastically generated spike trains, the effects of swelling-induced traumatic stress on spike-arrival-time based encoding are direct but not immediately illuminating. One assumes that a time sequence of spikes – modeled by a binary string of 1's and 0's with each bit representing a fixed small time interval – is observed after having traversed an axonal deformation. It features fewer spikes because of spike deletion, and certain select spike arrival times may also undergo augmentation (Delahunt et al., 2021; Maia & Kutz, 2014).[1] This results in a discretized neural code exhibiting greater numbers of consecutive 0's, less frequently interspersed with 1's; equivalently, a specific spike train will see on average an increase in the lengths of inter-spike intervals and an effective decrease in spike arrival frequency. While these bottom-line effects are clear, the stochastic variability typical in biological neurons obscures the interpretation that precise bit-string encoding is the most appropriate conduit of information on damage within the neural code (Abeles, 1994; Bialek et al., 1991; Haslinger et al., 2010; Lestienne, 1996).

We focus instead on stimulus and response firing rate measurements averaged over many realizations. A spiking network representing a neural tissue with a specific neural processing task is bathed in a global stimulus modeled by a Poisson spike train with a tunable frequency. In the spirit of electrode recordings from a specific location *in vivo*, a target neuron is selected as an output site and its spiking statistics are stored and analyzed for various frequencies of the global stimulus. This leads to a natural input-output relation between the frequency of the injected stimulus and the mean inter-spike interval at the recording site that typically features a specific monotone curve, the details of which are dependent on the characteristics of the particular network under investigation. In turn, changes to this graph indicate the effect that damage has, both at the single-cell level and on the network's at-large signal processing capabilities.

Proceeding in this manner, we are able to analyze successively more complex layered feedforward networks for input-output frequency response. We determine a measurement that assesses the difference between response curves of a damaged network and its undamaged analog (identical cell type of inhibitory or excitatory variety, and network architecture), and this allows us a direct view into characteristics that make a network more or less susceptible to damage from this type of traumatic injury. We show evidence indicating that network plasticity - specifically, a strategic addition of inhibitory cells or undamaged axonal connections - serves to attenuate the effects of damage, altogether pointing to mechanisms that may be at play in neural tissues as they recover from concussive trauma.

The paper is organized as follows. In Section 2 we present the details of the model investigated in the study, specifically, its properties as a network of interacting dynamical systems. We also motivate and define the measure leveraged in deducing the effects of focal axonal swelling on this type of network model. In Section 3 we provide a summary of results for the exhaustive study carried out for basic two-(modified to three-)element networks, as well as results from the exploratory analysis of layered networks by way of the principles developed in the small-network study. We also provide detailed discussion of the observed phenomenology. In Section 4 we summarize the main takeaways from the results presented in Section 3 and provide a look ahead into possible future work in this setting.

2 Methods

This investigation focuses on spiking networks of integrate-and-fire neurons employing axonal connections that can be either healthy or damaged by swelling resulting from traumatic stress. The present setting constitutes a specific example of a *network of dynamical systems*, featuring a complex interplay between local intracellular dynamics and interactions between cells Young (2022). A mathematical formalism for studying how the system changes over time can be achieved via a triple consisting of (1) a directed graph (the vertices and edges of which indicate the network structure), (2) the local dynamical rules governing intracellular processes at each vertex, and (3) the precise communication protocol in place between vertices adjacent due to the presence of an edge, the direction of which serves to distinguish between pre- and post-synaptic cells in the sense of neuroscience. We proceed to specify here parts (2) and (3) in this formalism precisely for our model, thereby drawing attention to the lexicon that we will use to describe and interpret the results of the paper. Part (1) will be detailed after this, in preparation for presentation and discussion of the results.

[1] Of note, the computational studies carried out herein are related to "Rule (v)" (a.k.a. *Increasing refractoriness*) in Maia et al. (2019) and examine in detail spike deletion in the absence of temporal augmentation scenarios; the methodology and algorithm outlined in the Appendix can be applied in the more general context as well.

2.1 Spiking integrate-and-fire neurons

The spiking network simulation method used here is a modified version of the synchronous or clock-driven algorithm presented by Brette et al. (2007), to which we have introduced a refractory period for each neuron. The latter amounts to a dynamical rule ensuring that sufficient time has passed since a neuron has last spiked before it is able to spike again. This mandatory quiescent phase places a *de facto* bound on individual neural firing rates; its incorporation into our model is motivated by the behavior of physical neurons Dayan and Abbott (2001). The communication between cells in these networks is achieved via the processing of incoming spikes, which serve to affect the local dynamics by resulting in contributions to the receiving cell's voltage. An extensive treatise comparing clock-driven algorithms to their event-driven counterparts is given in Rudolph and Destexhe (2006); the latter more biologically accurate alternative is not feasible in the present context.

Each neuron processes spikes according to an integrate-and-fire model much like that presented in Gerstner and Kistler (2002). A neuron can be *excitatory* or *inhibitory*, which means either a positive or negative contribution to the rate of change of voltage, V, of any neuron that is post-synaptic to it. Each neuron has three non-negative state variables that are updated at each timestep in the simulated dynamics – a voltage V, and excitatory/inhibitory potentials J_E and J_I which determine the impact of incoming excitatory and inhibitory spikes on the voltage evolution. Assuming there are no spikes at the input of a given cell, its state variables evolve as

$$\tau_V \frac{dV}{dt} = -V + J_E - J_I \tag{1}$$

$$\tau_J \frac{dJ_E}{dt} = -J_E \tag{2}$$

$$\tau_J \frac{dJ_I}{dt} = -J_I. \tag{3}$$

The time constants τ_V and τ_J determine the rate of these evolutions. Considering Eq. (1), we can see that in the absence of input, the voltage decays exponentially. Equations (2) and (3) provide that J_E and J_I also decay exponentially in the absence of spikes. Excitatory and inhibitory potentials lead to increases and decreases, respectively, to the rate of change of V, allowing the voltage to increase when the J_E term dominates.

The input of a spike causes an instantaneous jump by w in potential J_E (resp., J_I) if the spike is produced by a presynaptic excitatory (resp., inhibitory) cell. Cells have a threshold voltage V_T as well as a refractory period T_R which determine when a cell is able to spike. A neuron will *spike* if V rises above threshold voltage once a refractory period has passed.

Table 1 Parameters used in simulating networks of neurons, as described in Algorithm 1. Values were informed by (Brette et al., 2007; Gerstner & Kistler, 2002; Hansel et al., 1998)

Parameter	Units	Value	Significance
Δt	ms	0.1	Discrete timestep
τ_V	ms	18	Voltage time constant
τ_J	ms	5	J_E and J_I time constant
w	mV	0.5	Discrete potential step from incoming spike
V_T	mV	0.2	Threshold voltage
T_R	ms	1	Refractory period

The routine followed at each time step, separated by Δt, of the simulation is presented in Algorithm 1. In particular, note that voltage reset and spike emission occur only if voltage is above threshold *after* a refractory period has elapsed, a provision that in theory allows the voltage to fluctuate above and below threshold on the refractory timescale. The numerical values of the parameters in Algorithm 1 used in our simulation can be found in Table 1.

Algorithm 1 Network loop followed at each discrete timestep

```
for cell in network do
    cell.time+ = Δt
    for input in cell.inputs do
        if input == 1 then
            J_E = J_E + w
        else if input == −1 then
            J_I = J_I + w
        end if
    end for
    update cell via the solution to the ODEs
    if cell.V ≥ V_T and cell.time ≥ T_R then
        cell.V = 0
        cell.time = 0
        if cell.type = excitatory then
            cell.outputs = 1
        else if cell.type = inhibitory then
            cell.outputs = −1
        end if
    end if
end for
```

2.2 Axonal transmission and damage

In the case of an undamaged connection, the spikes from one neuron are incorporated into the dynamics of connected neurons with no intermediate processing. However, if a connection is damaged, the spike train is first subjected to a filter that integrates the effects of the damage before the processing of the connected neurons. This filter can either delete spikes or allow them to be transmitted across the connection faithfully.

A number of detailed models of FAS-related complications to signal transmission have been developed in recent years. The literature surveyed in Wu et al. (2021) presents

a diverse picture of considerations that are important in the study of *traumatic brain injuries (TBI)*. Details of axonal geometry and its impact on electrical transmission are recognized as critical to modeling these phenomena faithfully, see, e.g. Ofer and Shefi (2016) for an overview of relevant considerations. In particular, an important category of models for phenomenology specific to *spherical axonal swelling* is supported by physiological data Wu et al. (2020) and provides a two-dimensional generalization of the model we have selected for the present investigation; it should be noted that, although we have chosen not to focus on this aspect, such further detail in the modeling allows for the incorporation of surface area effects around FAS deformations that likely play a crucial role in altering the local action-potential firing threshhold along the extent of the damaged axon.

The damage filter implemented in this study is inspired by the model of axonal swelling presented and analyzed by Maia and Kutz (2013). This model was derived by introducing a continuous deviation of cable radius in the active cable equation given by

$$\frac{\partial V}{\partial x} = \frac{D}{d(x)}\frac{\partial}{\partial x}\left(\frac{d^2(x)}{r_L(x)}\frac{\partial V}{\partial x}\right) + V(V-a)(V-1) - R \tag{4}$$

$$\frac{\partial R}{\partial t} = bV - cR \tag{5}$$

where V is the voltage, x is distance along the axon, $d(x)$ is the radius of the axon at position x, $r_L(x)$ is the resistance of the axon at position x, R is a *gating variable* (which models the probability of ion channels being open, as in the Hodgkin-Huxley model Hodgkin and Huxley (1952)), and a, b, c and D are physical constants. The damage is implemented by letting $d(x)$ grow from width δ_B to δ_A (that is, to *swell*) over a distance δ_T, according to

$$d(x) = \begin{cases} \delta_B, & \text{for } x \le 0 \\ \tilde{d}(x), & \text{for } 0 \le x \le \delta_T \\ \delta_A, & \text{for } x \ge \delta_T \end{cases} \tag{6}$$

where

$$\tilde{d}(x) = \delta_B + (\delta_A - \delta_B)\left(10\frac{x^3}{\delta_T^3} - 15\frac{x^4}{\delta_T^4} + 6\frac{x^5}{\delta_T^5}\right). \tag{7}$$

The resistance r_L in the present setting is assumed to be a constant 1 Ω. The constants' physical interpretations, units, and values used to calibrate the subsequent studies can be seen in Table 2.

A numerical solution of this PDE model can be carried out by an adaptation of the spectral method Thapa and Gudejko (2014) to nonlinear equations Fornberg and Sloan (1994). The resulting *pseudo*-spectral method employed by Maia and Kutz is detailed in the Appendix. These numerical solutions are computationally expensive even on the single-cell scale, and a faster simulation method is required so that extensive simulations can be performed at the network level.

Table 2 Constants used in our simulations of the active cable model, along with their units and biological significance. Values are informed by those used by Maia and Kutz (2013)

Constant	Units	Value	Significance
D	$cm^2/\mu F$	0.02	Inverse of specific membrane capacitance
a	mV	0.1	Fitzhugh-Nagumo offset voltage
b	$\mu A{\cdot}\Omega/cm^2$	0.01	Voltage contribution to R evolution
c	Hz	0.05	R contribution to R evolution
d_A	cm	4	Final diameter of axon
d_B	cm	2	Initial diameter of axon
d_T	cm	0.25	Distance across which swelling occurs

This has been previously considered by Maia et al. at the single-cell level, in which the authors present formulaic representations of "rules" that qualitatively reproduce the character of spiketrain transmission along damaged axons Maia et al. (2019). Their fifth rule, *Increasing refractoriness*, is most applicable to the present study as it outlines spike deletions without temporal shifts. While this method does very well at representing certain parameter regimes for damaged axons, it cannot reproduce more complex deletion patterns that do not simply introduce an effective increase in the refractory period.

Our approach does not assume any behavior *a priori*, and amounts to a statistical algorithm trained from the behavior of the damaged active cable model. We note that increased refractoriness of spike trains through FAS deformations has been observed to be linked to the frequency of incoming spikes, with more pronounced effects on higher frequencies Crodelle and Maia (2021). In this regard, the literature supports the notion that the selective spike-deletion effect of FAS-traversal is a fixed-length time frame effect that can be modeled appropriately with a finite-memory kernel reconstruction such as that proposed herein. The details of this *damaged axon prediction algorithm* (DAPA) are presented in the Appendix, with a side-by-side comparison with the output of the pseudo-spectral numerical method and the increasing refractoriness rule Maia et al. (2019). This comparison indicates that the shortcut produces results statistically equivalent to the pseudo-spectral method. It is more general than the increasing refractoriness rule of Maia et al., as it can learn more complex deletion patterns.

2.3 Quantifying the effects of damage

In assessing the effects of damage on a network's capability for rate-encoding, we propose a measure based on the mean

inter-spike interval I_{av} of a network's output over an interval of frequencies. The nonlinearity of the network model precludes studying the frequency response via tonal analysis (i.e. stimuli with one frequency component). Networks are instead stimulated by a Poisson spiketrain Heeger et al. (2000) presented as a uniform background excitatory stimulus to all neurons. For each input frequency λ (Hz), at each time step, we simulate a single sample from the Bernoulli distribution with "success" probability $p = \lambda \Delta t$. A success leads to a spike being processed as an input at every neuron. This spike is handled in the same way as a spike from a presynaptic (excitatory) neuron.

For each input stimulus frequency λ, many realizations of the input are simulated to obtain an associated $I_{av}(\lambda)$ for a given network architecture and damage paradigm. Reported $I_{av}(\lambda)$ values are averaged across all realizations; in this study, the presented average interspike intervals are averaged over 10,000 realizations.

The prototypical frequency response curve is shown in Fig. 1. Such a curve will monotonically decrease towards a horizontal asymptote at T_R, which is the theoretical lower bound of the output spiking period. Thus, the network's capacity to discriminate between different inputs diminishes as the input frequency increases.

We define a *cutoff frequency* above which we say the network poorly discriminates between input stimulus frequencies. Every network comes with an associated cutoff frequency. When a network is damaged, it stands to reason that the cutoff frequency will change. This change in cutoff frequency depends on both the damage paradigm and the original architecture to which damage is applied. We are interested in how network architecture itself, as well as modification to it via addition of connections and cells, result in robustness to the effect of damage in this sense.

For an undamaged network of some architecture, we define the *undamaged network cutoff frequency* λ_u by

$$\lambda_u = \inf\{\lambda \mid I_{av}(f) \leq (1.1)T_R, \forall f > \lambda\}. \tag{8}$$

Analogously, let the *damaged network cutoff frequency* λ_d be such a frequency for the associated damaged network of the same architecture under a given damage paradigm. Figure 1 also shows prototypical $I_{av}(\lambda)$ curves for associated undamaged and damaged networks with λ_u and λ_d.

We define a network's *bandwidth* as the length of the frequency interval over which the network discriminates faithfully. We define the *bandwidth damage metric*, d_{BW}, as

$$d_{BW} = \left| \frac{\lambda_d - \lambda_u}{\lambda_u} \right| \times 100. \tag{9}$$

This is the absolute percent deviation of the damaged bandwidth from the undamaged bandwidth incurred by

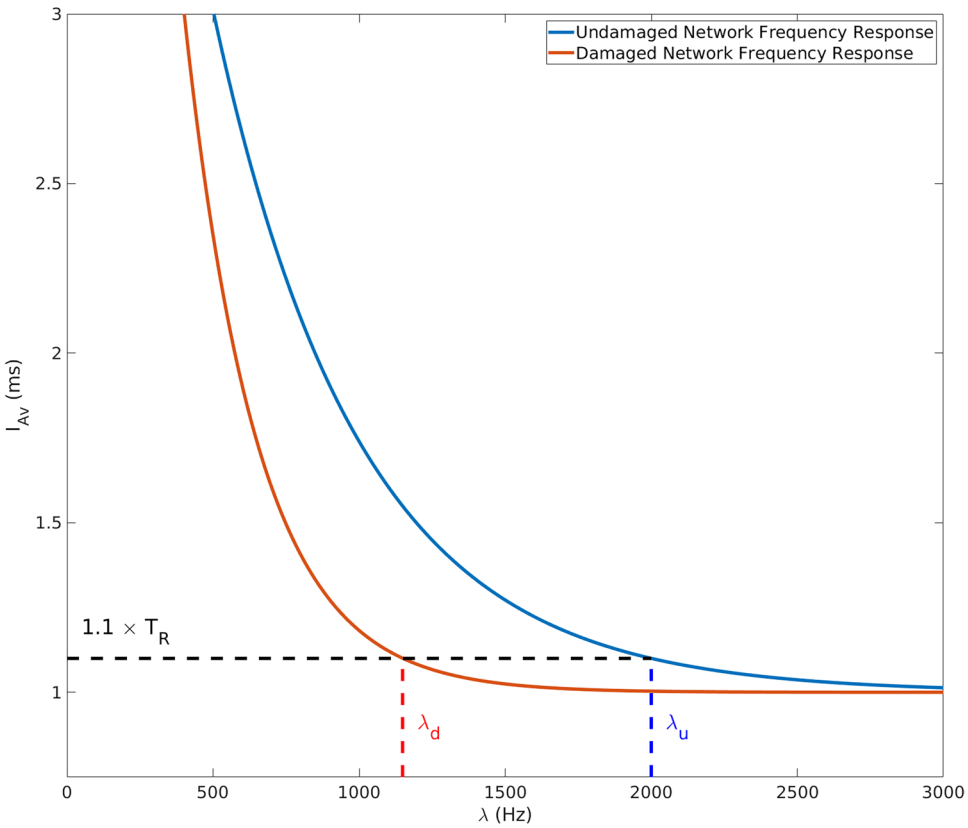

Fig. 1 Prototypical curves for the average interspike interval as a function of frequency for associated undamaged and damaged networks. They decrease monotonically towards an asymptote at T_R. We also have marked λ_u and λ_d, the cutoff frequencies for these networks where the interspike interval reaches $1.1 \times T_R$. Note that while in this illustration $\lambda_d < \lambda_u$, this is not always the case. Our damage metric d_{BW} measures absolute bandwidth difference, so it does not discriminate between the cases where damage incurs a decrease or increase in bandwidth

the specified damage paradigm. That is, it measures how much damage impacts the ability of a network to discriminate between high-frequency stimuli.

It is important to note that there is no inherent value to a higher cutoff frequency over a lower one. The cutoff frequency is simply a property of a network which may have value associated with its function. As an engineering analogue, the value of the cutoff frequency of a low-pass filter is application dependent, and whether raising or lowering it would improve performance is entirely application-dependent.

In this light, the absolute value in Eq. (9) is critical. This metric encodes how significantly damage alters this fundamental network property in either direction, so a more robust network in the d_{BW} sense will better maintain its original associated cutoff frequency in the face of damage.

The choice of globally uniform (or coherent) input stimulus is made out of convenience. In the Appendix, we consider the implications of this assumption by comparing the results of small network experiments with coherent stimulus to those of networks where each neuron receives an independent (incoherent) stimulus. We show that with respect to the d_{BW} metric, there is no characteristic difference in our findings for small networks between use of coherent and incoherent stimuli.

2.4 Network architectures of focus

We consider networks of neurons composed of two-cell feedforward layers in which the layers are also connected by a feedforward structure. This special case of networks is of relevance, as feedforward layered integrate-and-fire models have been used in the past to approach issues such as logic gating and schizophrenia (Vogels, 2005; Vogels et al., 2005; Vogels & Abbott, 2007). The spiking properties of such networks are observed by measuring the output of the last neuron with respect to the order imposed by the feedforward structure within and between layers. To begin, we observe the undamaged network's response. We then compare this to the response of an associated network featuring damage on all axons within layers, but not those between layers. We record from the *same* neuron to obtain the damaged response.

We then consider layer-wise modifications to the network, such as the additions of feedback paths or neurons. The undamaged and damaged responses of this modified network are determined as well. In each such instance, damage is applied only to the axons that were damaged in the unmodified network.

2.4.1 Two-cell networks

The one-layer case is that of the two-cell network. We consider the d_{BW} incurred by damaging the single axon in this network, where the output is observed from the second neuron in the feedforward chain.

We consider two modifications to networks of this type: (1) the addition of an axon in feedback from the second neuron to the first, and (2) the addition of a neuron at the front of the feedforward chain, connected by a new axon to the first cell in the two-cell feedforward chain. The additional cells in modification (2) can be either inhibitory or excitatory. We also consider the case where both modifications (addition of a feedback path and addition of an additional neuron) take place. As such there are five possible modifications for each two-cell feedforward network. Figure 2 provides a graphical representation of these modifications made to a given two-cell network. In determining the d_{BW} for these modified networks, we still consider damage only on the single axon which was present in the two-cell network, and we continue to measure the output at the final cell in the feedforward chain.

2.4.2 Layered networks

In studying layered networks, we focus on the *homogeneous* case, in which all layers are identical. The inter-layer connections are realized as an axon connecting the second neuron in a feedforward chain to the first neuron in the next layer's feedforward chain. With L layers, this network is a $2L$-cell feedforward chain, and we read the output from the last cell in this chain. We make modifications of the same 5 types as we did in the two-cell case layer-wise. Figure 3 provides a graphical representation of these modifications made to a given layered feedforward network.

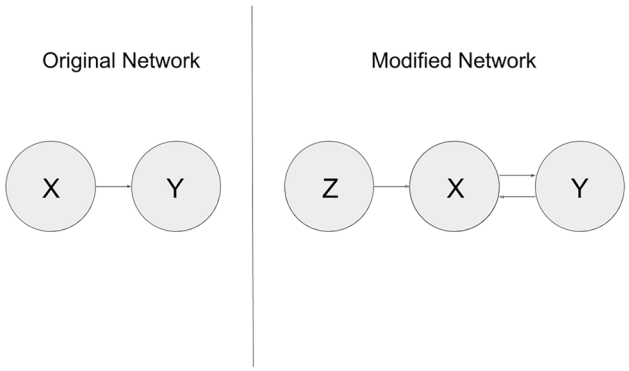

Fig. 2 The left-hand side shows an arbitrary feedforward network consisting of two cells. X and Y represent the type of cell present at that node (inhibitory or excitatory), and X can be the same as Y. The right-hand side shows modifications to the network, where Z is the type of the potentially added neuron. Arrows with white arrowheads represent connections which can be made so as to modify the network – a feedback path can be added, a third neuron can be added at the front of the chain, or both can be added at once

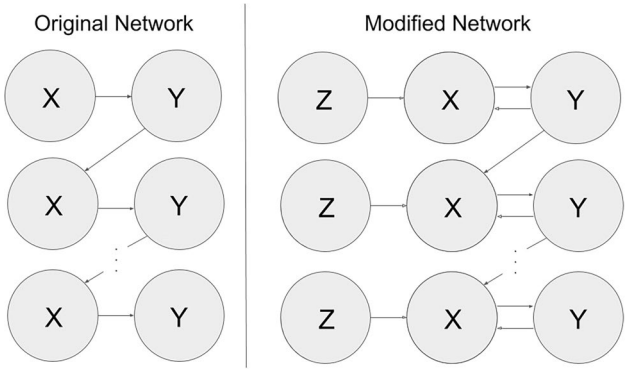

Fig. 3 The left-hand side shows an arbitrary feedforward network consisting of two-cell layers. The right-hand side shows modifications to the network, which are applied uniformly across layers

3 Results and discussion

We present an analysis of the studies described in Section 2. We summarize the results briefly below, and then provide discussion for each network case in detail.

All investigated networks either benefit from or are unaffected by the addition of inhibitory neurons at the front of each layer with respect to d_{BW}-robustness in the face of imposed damage. Due to the exhaustive nature of this study, this provides a plausible, consistent mechanism for the improved robustness of small layered networks, by way of plasticity via the addition of inhibition.

For single-layer feedforward networks, there always exist modifications which lower the d_{BW}. In particular, the addition of an inhibitory cell to the front of a feedforward chain will always result in a smaller d_{BW} than in the base case. Thus, for networks of this type, there is a single mechanism by which neuroplasticity can consistently promote recovery.

For larger layered networks of the homogeneous type, the addition of an inhibitory cell to the front of each layer will never increase d_{BW} from its base case value. However, it does not necessarily decrease this d_{BW}, meaning that this modification will at worst have no effect on the bandwidth of the network.

It may be noted that the modifications alone also change the undamaged network's bandwidth, in a manner that is often more significant than the effect of damage. As such,

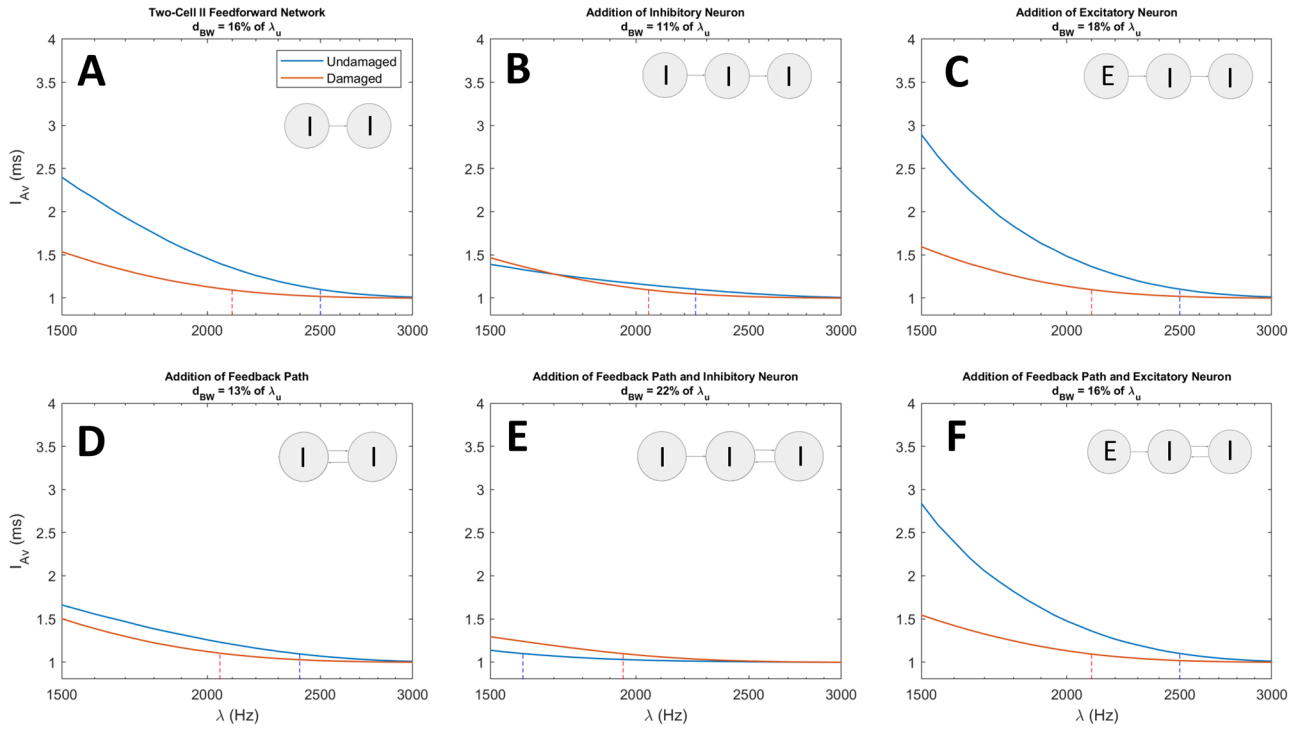

Fig. 4 Damaged and undamaged frequency responses of all tested modifications for a single-layer network consisting of two inhibitory neurons. Dotted lines indicates the cutoff frequencies, λ_u and λ_d. **A** – The frequency responses for the base case: a two-cell feedforward network. **B, C** – The frequency responses for modified networks in which a single neuron is added to the front of the feedforward chain of the base network from that in panel A. It can be seen that addition of an inhibitory cell lowers the d_{BW} value from the base case value, while adding an excitatory cell raises the d_{BW} value. **D** – The frequency response of the modified network in which a feedback path is added from the output neuron to the first neuron in the feedforward chain, resulting in a lower d_{BW}. **E, F** – Frequency responses for modified networks in which both a feedback path and a neuron is added. Addition of feedback and an inhibitory neuron results in a higher d_{BW}, whereas addition of feedback and an excitatory neuron incurs no significant change in d_{BW}

Table 3 d_{BW} values for one-layer networks presented in % of λ_u. Highlighted cells show d_{BW} values for modified networks which are lower than the d_{BW} for the respective base case. The addition of an inhibitory cell invariably lowers the d_{BW} from the base case. Similarly, the addition of an excitatory cell in feedforward always raises the d_{BW} from the base case. For an "EI" network, the addition of a feedback path as well as an excitatory neuron in feedforward yields a d_{BW} of 0, meaning that for such a modified network, the damaged and undamaged frequency responses are indistinguishable. Addition of a feedback path alone is restorative only for a network consisting of two inhibitory neurons

Network	Base Case	Add I Cell	Add E Cell	Add FB	Add FB and I	Add FB and E
EE	10	8	24	50	29	67
EI	10	8	24	11	10	0
IE	16	11	18	18	16	18
II	16	11	18	13	22	16

the modifications themselves should not be thought of as mechanisms to recover the original network bandwidth but instead to boost the network bandwidth's reliability or robustness in the face of damage.

3.1 Single-layer feedforward networks

In this section we present the exhaustive numerical study of the effects of damage on single-layer feedforward networks. We compare the d_{BW} values for feedforward networks of two cells with corresponding modified networks of two to three cells.

For example, Fig. 4 shows the λ vs. I_{AV} curves from which d_{BW} is computed for a feedforward network of two inhibitory cells, as well as the effects of the 5 modifications to such a network.

Table 3 contains all such d_{BW} comparisons for the single-layer case. Note that the addition of an inhibitory cell always has a restorative effect on the network's d_{BW} – that is, d_{BW} is always decreased by addition of an inhibitory cell. From an intuitive perspective, the addition of an inhibitory cell at the front of a network decreases the firing rate of the cell to which it is connected. We expect the damage to act as a non-linear low-pass filter, so a decrease in firing rate would make spikes more likely to be faithfully transmitted through a damaged axon. Thus, the restorative nature of the addition of an inhibitory cell in feedforward is consistent with intuition.

In certain networks, other modifications outperform the addition of an inhibitory cell. For example, in the case of EI feedforward networks, the addition of both a feedback connection and an excitatory neuron entirely eliminates changes in the bandwidth relative to the precision of our simulations. The restorative mechanism of such a modification is not particularly intuitive as a result of the competing effects that are present. That is, excitation at the front of a feedforward network should increase the firing rate of the first neuron and make damage more impactful, but the feedback acts so as to lower the firing rate in a manner proportional to the firing rate itself. In this case, these competing effects play out in the favor of the network's recovery. It is clear from this example that for networks with any deviation from the two-cell feedforward structure, it is difficult to intuit which modifications may result in a lowering of d_{BW}. This modification in particular does not have restorative effects for any other two-cell network architecture. No modification other than the addition of a single inhibitory cell is seen to have a consistent restorative effect on networks of this type.

3.2 Homogeneous layered feedforward networks

3.2.1 Two-layer networks

In this section we present the effects of damage on layered networks of the homogeneous type. To begin, we consider two-layered networks, for which Table 4 shows the results of all possible modifications.

Table 4 d_{BW} values for 2-layer networks presented in % of λ_u. Highlighted cells show d_{BW} values for modified networks which are lower than the d_{BW} for the respective base case. The addition of an inhibitory cell at the front of each layer can be seen to provide a restorative effect in three of four networks. In the "EI" network, we see that the base d_{BW} is 0. In this case, the addition of an inhibitory cell does not raise the d_{BW}. Thus, the layer-wise addition of an inhibitory cell in feedforward has, at worst, a neutral effect on two-layer feedforward networks. Unlike in the single layer case, we see here that there is never a better modification than the addition of inhibitory cells. We also see that the addition of both feedback paths and inhibitory neurons in each layer provides a restorative or neutral effect in all four networks

Network	Base Case	Add I Cell	Add E Cell	Add FB	Add FB and I	Add FB and E
EE	33	5	62	82	24	67
EI	0	0	0	0	0	5
IE	18	16	18	18	16	18
II	15	0	16	0	0	16

Table 5 d_{BW} values for 10-layer networks presented in % of λ_u. Highlighted cells show d_{BW} values for modified networks which are lower than the d_{BW} for the respective base case. Layer-wise addition of inhibitory neurons has a restorative effect in three of four networks, only having a neutral effect on a network with a base d_{BW} of 0. Furthermore, in all cases there are no modifications which provide a more restorative effect than the addition of inhibitory cells layer-wise

Network	Base Case	Add I Cell	Add E Cell	Add FB	Add FB and I	Add FB and E
EE	43	24	67	82	67	82
EI	0	0	5	0	0	0
IE	18	16	18	18	18	18
II	15	0	18	0	0	18

Note that there are networks which show an increase in d_{BW} in response to modifications that have restorative effects in their respective two-cell analogues. For example, we saw that the addition of a feedback path and an excitatory cell had restorative effects on the EI feedforward network of the two-cell type, yet the analogous modification to the two-layer network creates the largest observed increase to d_{BW} across all modifications. This indicates that layered networks do not inherit the characteristics of their single-layer analogues.

The addition of an inhibitory cell to each layer of a two-layered homogeneous network never increases d_{BW}, but it also does not always decrease it. This is weaker than the result seen for two-cell feedforward networks, but still allows for a mechanism by which neuroplasticity can attempt to restore a network of this type without risk of detrimental effects. This provides an instance in which two-layer networks inherit characteristics from their one-layer analogues, despite this not being the case in general.

We note also that there are two-layer networks which are very robust to damage with respect to this metric. For example, the network which is composed of EI layers experiences no change in bandwidth due to the addition of damage.

3.2.2 Networks containing more layers

We have recognized a pattern in which the layer-wise addition of inhibitory drive enhances a network's robustness to damage in the one- and two-layer cases. It is natural to ask whether this pattern will continue as the number of layers is increased. Since the complexity of networks increases with the number of layers, we illustrate the potential effects of modification for larger layered networks by focusing on case studies of ten and fifty layers.

Table 5 shows the d_{BW} for ten-layer networks. At this size, it is still the case that layer-wise addition of inhibitory drive never increases a network's d_{BW}. Fifty-layer networks exhibit this property as well, as seen in Table 6. That is, for *all* tested sizes of homogeneous feedforward networks, we observe a consistent restorative effect of additional layer-wise inhibitory drive.

Table 6 d_{BW} Values for 50-layer networks in which only the addition of an inhibitory cell modification is made, presented in % of λ_u. These networks see at worst neutral effects from the addition of inhibitory neurons layer-wise, where this modification is only neutral when the base network has a d_{BW} of 0. Two of four base networks show d_{BW} values of 0, while in the two networks which present non-zero d_{BW}, we observe d_{BW} values of a similar magnitude to those observed in smaller networks

Network	Base Case	Add I Cell
EE	43	24
EI	0	0
IE	18	16
II	0	0

4 Conclusions

For homogeneous layered networks of varying sizes, we find that there is always a modification that can diminish the effect of damage as captured by the d_{BW} metric. Biophysically, these modifications can be interpreted as the rerouting of both extra- and intra-network connections via neuroplasticity. Thus, neuroplasticity may always provide a mechanism for the mitigation of the effects of damage incurred by focal axonal swelling.

In particular, we see that there exists a single modification – the layer-wise addition of inhibition – that always increases the robustness of networks of this type. As this modification's effect is present across networks of all tested sizes, it provides a simple mechanism through which neuroplasticity may aid any homogeneous feedforward network. Other restorative modifications are present in specific network architectures and at specific network sizes. For example, in the single-layer "EI" case, the addition of a feedback connection and an excitatory drive can reduce d_{BW} to 0. Meanwhile, this modification exacerbates damage in all other networks of this size, and in all tested larger networks. Similarly, while the addition of feedback and inhibition has at worst neutral effects on two-layer networks, it can exacerbate damage in networks of other sizes. Thus, the d_{BW} metric is sensitive to perturbations of both network size and network architecture. Besides layer-wise addition of inhibition, there is no tested

modification that acts in a restorative fashion on networks of all architectures and sizes.

Network size, in the sense of number of layers, also plays a restorative role for certain architectures. For example, the "EI" network presents a base d_{BW} of 0 at two layers, ten layers and fifty layers. However, network size does not have a restorative effect in general, as "EE" and "IE" networks present base d_{BW} values that never decrease as network size is increased.

Our results show mechanisms by which networks of a specific architecture can be made more robust to focal axonal swelling. It is natural to consider similar modifications made to networks of similar architectures. For example, one could consider layered networks that do not have a feedforward structure, or heterogeneous layered networks. With regards to possible future work in this area, it would be appropriate to review the recent attention given to *recurrent network structure* and its likely central role in training / machine learning, especially in the context of neuronal networks capable of spiketrain signal encoding as relevant for the mamalian nervous system Bellec et al. (2020). We see the application of our signal filtering scheme as well as the network bandwidth analysis as promising tools for lending insight into the interplay of this area of research with the effects of traumatic brain injuries and possibilities for recovery and robustness.

Our results on homogeneous networks imply that architecture plays a critical role in the base resilience of a network to damage. Naturally, one could consider the set of all small perturbations to a given network and study their effects systematically and exhaustively. Such a study could potentially isolate aspects of architectures which make networks more or less susceptible to damage. With the use of DAPA, such a study is computationally tractable even for large networks – even ones that feature more complicated neuron dynamics models at the cell level. This line of work can also be interpreted in the context of network degeneracy (Neuberger et al., 2017; Kamaleddin, 2021; Stöber et al., 2023), thus drawing connections with a large body of current research in neuroscience.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s10827-023-00860-0.

Author contributions Both authors contributed equally to this work.

Funding Brian L. Frost is funded by an F-31 grant from the National Institute of Deafness and other Communication Disorders (NIDCD). Grant Number: 1F31DC020621-01. Dr. Stanislav M. Mintchev has no source of funding to report.

Availability of data and materials Simulation software is available at https://github.com/brian-frost-laplante/axon.

Declarations

Ethical approval This is not applicable, as this study did not involve experimentation on any human or animal subjects.

Conflict of Interest The authors declare no conflict of interest.

Competing interests Brian L. Frost is a PhD student at Columbia University, and is funded by the NIDCD. This study does not relate to his source of income or any outside interests. Dr. Stanislav M. Mintchev is a professor of mathematics at the Cooper Union. This study does not relate to his source of income or any outside interests.

References

Abeles, M. (1994). *Firing rates and weil-timed events in the cerebral cortex*, pp. 121–140. Springer. https://doi.org/10.1007/978-1-4612-4320-5_3

Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., & Maass, W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications, 11*(1). https://doi.org/10.1038/s41467-020-17236-y

Bialek, W., Rieke, F., de Ruyter van Steveninck, R. R., & Warland, D. (1991). Reading a neural code. *Science, 252*(5014), 1854–1857. https://doi.org/10.1126/science.2063199

Brette, R., Rudolph, M., Carnevale, T., Hines, M., Beeman, D., Bower, J. M., Diesmann, M., Morrison, A., Goodman, P. H., Harris, F. C., Zirpe, M., Natschläger, T., Pecevski, D., Ermentrout, B., Djurfeldt, M., Lansner, A., Rochel, O., Vieville, T., Muller, E., … & Destexhe, A. (2007). Simulation of networks of spiking neurons: A review of tools and strategies. *Journal of Computational Neuroscience, 23*(3), 349–398. https://doi.org/10.1007/s10827-007-0038-6

Crodelle, J., & Maia, P. D. (2021). A computational model for pain processing in the dorsal horn following axonal damage to receptor fibers. *Brain Sciences, 11*(4), 505. https://doi.org/10.3390/brainsci11040505

Dayan, P., & Abbott, L. F. (2001). *Theoretical Neuroscience : Computational and Mathematical Modeling of Neural Systems*. Cambridge, Mass: MIT Press.

Debanne, D., Campanac, E., Bialowas, A., Carlier, E., & Alcaraz, G. (2011). Axon physiology. *Physiological Reviews, 91*(2), 555–602. https://doi.org/10.1152/physrev.00048.2009

Delahunt, C. B., Maia, P. D., & Kutz, J. N. (2021). Built to last: Functional and structural mechanisms in the moth olfactory network mitigate effects of neural injury. *Brain Sciences, 11*(4), 462. https://doi.org/10.3390/brainsci11040462

Ermentrout, G. B., & Terman, D. H. (2010). *Mathematical Foundations of Neuroscience*. Springer New York, NY. https://doi.org/10.1007/978-0-387-87708-2

Fornberg, B., & Sloan, D. M. (1994). A review of pseudospectral methods for solving partial differential equations. *Acta Numerica, 3*, 203–267. https://doi.org/10.1017/s0962492900002440

Gerstner, W., & Kistler, W. M. (2002). *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press, Cambridge, U.K. New York. https://books.google.com/books?id=Rs4oc7HfxIUC

Hansel, D., Mato, G., Meunier, C., & Neltner, L. (1998). On numerical simulations of integrate-and-fire neural networks. *Neural Computation, 10*(2), 467–483. https://doi.org/10.1162/089976698300017845

Haslinger, R., Klinkner, K. L., & Shalizi, C. R. (2010). The computational structure of spike trains. *Neural Computation, 22*(1), 121–157. https://doi.org/10.1162/neco.2009.12-07-678

Heeger, D., et al. (2000). Poisson model of spike generation. *Handout, University of Standford, 5*(76), 1–13.

Herculano-Houzel, S. (2009). The human brain in numbers: a linearly scaled-up primate brain. *Frontiers in Human Neuroscience, 3*. https://doi.org/10.3389/neuro.09.031.2009

Hodgkin, A. L., & Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology, 117*(4), 500–544. https://doi.org/10.1113/jphysiol.1952.sp004764

Izhikevich, E. M. (2004). Which model to use for cortical spiking neurons? *IEEE Transactions on Neural Networks, 15*(5), 1063–1070. https://doi.org/10.1109/tnn.2004.832719

Johnson, V. E., Stewart, W., & Smith, D. H. (2013). Axonal pathology in traumatic brain injury. *Experimental Neurology, 246*, 35–43. https://doi.org/10.1016/j.expneurol.2012.01.013

Kamaleddin, M. A. (2021). Degeneracy in the nervous system: from neuronal excitability to neural coding. *BioEssays, 44*(1), 2100148. https://doi.org/10.1002/bies.202100148

Lestienne, R. (1996). Determination of the precision of spike timing in the visual cortex of anaesthetised cats. *Biological Cybernetics, 74*(1), 55–61. https://doi.org/10.1007/bf00199137

Lusch, B., Weholt, J., Maia, P. D., & Kutz, J. N. (2018). Modeling cognitive deficits following neurodegenerative diseases and traumatic brain injuries with deep convolutional neural networks. *Brain and Cognition, 123*, 154–164. https://doi.org/10.1016/j.bandc.2018.02.012

Maia, P. D., Hemphill, M. A., Zehnder, B., Zhang, C., Parker, K. K., & Kutz, J. N. (2015). Diagnostic tools for evaluating the impact of focal axonal swellings arising in neurodegenerative diseases and/or traumatic brain injury. *Journal of Neuroscience Methods, 253*, 233–243. https://doi.org/10.1016/j.jneumeth.2015.06.022

Maia, P. D., & Kutz, J. N. (2013). Identifying critical regions for spike propagation in axon segments. *Journal of Computational Neuroscience, 36*(2), 141–155. https://doi.org/10.1007/s10827-013-0459-3

Maia, P. D., & Kutz, J. N. (2014). Compromised axonal functionality after neurodegeneration, concussion and/or traumatic brain injury. *Journal of Computational Neuroscience, 37*(2), 317–332. https://doi.org/10.1007/s10827-014-0504-x

Maia, P. D., & Kutz, J. N. (2017). Reaction time impairments in decision-making networks as a diagnostic marker for traumatic brain injuries and neurological diseases. *Journal of Computational Neuroscience, 42*(3), 323–347. https://doi.org/10.1007/s10827-017-0643-y

Maia, P. D., Raj, A., & Kutz, J. N. (2019). Slow-gamma frequencies are optimally guarded against effects of neurodegenerative diseases and traumatic brain injuries. *Journal of Computational Neuroscience, 47*, 1–16.

Manor, Y., Koch, C., & Segev, I. (1991). Effect of geometrical irregularities on propagation delay in axonal trees. *Biophysical Journal, 60*(6), 1424–1437. https://doi.org/10.1016/s0006-3495(91)82179-8

Maxwell, W. L., Povlishock, J. T., & Graham, D. L. (1997). A mechanistic analysis of nondisruptive axonal injury: A review. *Journal of Neurotrauma, 14*(7), 419–440. https://doi.org/10.1089/neu.1997.14.419

Neuberger, E. J., Gupta, A., Subramanian, D., Korgaonkar, A. A., & Santhakumar, V. (2017). Converging early responses to brain injury pave the road to epileptogenesis. *Journal of Neuroscience Research, 97*(11), 1335–1344. https://doi.org/10.1002/jnr.24202

Ofer, N., & Shefi, O. (2016). Axonal geometry as a tool for modulating firing patterns. *Applied Mathematical Modelling, 40*(4), 3175–3184. https://doi.org/10.1016/j.apm.2015.10.017

Ramón, F., Joyner, R. W., & Moore, J. W. (1975). Propagation of action potentials in inhomogeneous axon regions, pp. 85–100. Springer. https://doi.org/10.1007/978-1-4684-2637-3_8

Rudolph, M., & Destexhe, A. (2006). Analytical integrate-and-fire neuron models with conductance-based dynamics for event-driven simulation strategies. *Neural Computation, 18*(9), 2146–2210. https://doi.org/10.1162/neco.2006.18.9.2146

Rudy, S., Maia, P. D., & Kutz, J. N. (2016). Cognitive and behavioral deficits arising from neurodegeneration and traumatic brain injury: a model for the underlying role of focal axonal swellings in neuronal networks with plasticity. *Journal of Systems and Integrative Neuroscience, 2*(2), 114–121. https://doi.org/10.15761/jsin.1000120

Sharp, D. J., Scott, G., & Leech, R. (2014). Network dysfunction after traumatic brain injury. *Nature Reviews Neurology, 10*(3), 156–166. https://doi.org/10.1038/nrneurol.2014.15

Stöber, T. M., Batulin, D., Triesch, J., Narayanan, R., & Jedlicka, P. (2023). Degeneracy in epilepsy: multiple routes to hyperexcitable brain circuits and their repair. *Communications Biology, 6*(1). https://doi.org/10.1038/s42003-023-04823-0

Sussillo, D., & Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron, 63*(4), 544–557. https://doi.org/10.1016/j.neuron.2009.07.018

Tagge, C. A., Fisher, A. M., Minaeva, O. V., Gaudreau-Balderrama, A., Moncaster, J. A., Zhang, X.-L., Wojnarowicz, M. W., Casey, N., Lu, H., Kokiko-Cochran, O. N., Saman, S., Ericsson, M., Onos, K. D., Veksler, R., Senatorov, V. V., Kondo, A., Zhou, X. Z., Miry, O., Vose, L. R., … & Goldstein, L. E. (2018). Concussion, microvascular injury, and early tauopathy in young athletes after impact head injury and an impact concussion mouse model. *Brain, 141*(2), 422–458. https://doi.org/10.1093/brain/awx350

Tang-Schomer, M. D., Johnson, V. E., Baas, P. W., Stewart, W., & Smith, D. H. (2012). Partial interruption of axonal transport due to microtubule breakage accounts for the formation of periodic varicosities after traumatic axonal injury. *Experimental Neurology, 233*(1), 364–372. https://doi.org/10.1016/j.expneurol.2011.10.030

Thapa, N., & Gudejko, M. (2014). Numerical solution of heat equation by spectral method. *Applied Mathematical Sciences, 8*, 397–404. https://doi.org/10.12988/ams.2014.39502

Vogels, T. P., & Abbott, L. (2007). Gating deficits in model networks: a path to schizophrenia? *Pharmacopsychiatry, 40*(S1), 73–77.

Vogels, T. P. (2005). Signal propagation and logic gating in networks of integrate-and-fire neurons. *Journal of Neuroscience, 25*(46), 10786–10795. https://doi.org/10.1523/jneurosci.3508-05.2005

Vogels, T. P., Rajan, K., Abbott, L. F., et al. (2005). Neural network dynamics. *Annual review of neuroscience, 28*, 357.

Wang, J., Hamm, R. J., & Povlishock, J. T. (2011). Traumatic axonal injury in the optic nerve: Evidence for axonal swelling, disconnection, dieback, and reorganization. *Journal of Neurotrauma, 28*(7), 1185–1198. https://doi.org/10.1089/neu.2011.1756

Wu, Y.-T., Gilpin, K., & Adnan, A. (2020). Effects of focal axonal swelling level on the action potential signal transmission. *Journal of Computational Neuroscience, 48*(3), 253–263. https://doi.org/10.1007/s10827-020-00750-9

Wu, Y.-H., Rosset, S., Lee, T.-R., Dragunow, M., Park, T., & Shim, V. (2021). *In vitro* models of traumatic brain injury: A systematic review. *Journal of Neurotrauma, 38*(17), 2336–2372. https://doi.org/10.1089/neu.2020.7402

Young, L.-S. (2022). The brain is a dynamical system. *SIAM News, 55*(6), 1–2.