

Crash Testing an Engineering Framework in Neuroscience: Does the Idea of Robustness Break Down?¹

ABSTRACT

In this paper I discuss the concept of *robustness* in neuroscience. Various mechanisms for making systems robust have been discussed across biology and neuroscience (e.g. redundancy and fail-safes). Many of these notions originate from engineering. I argue that concepts borrowed from engineering aid neuroscientists in (1) operationalizing robustness; (2) formulating hypotheses about mechanisms for robustness; and (3) quantifying robustness. Furthermore, I argue that the significant disanalogies between brains and engineered artefacts raise important questions about the applicability of the engineering framework. I argue that the use of such concepts should be understood as a kind of simplifying idealization.

“The brain is a physical device that performs specific functions; therefore, its design must obey general principles of engineering.”

Sterling and Laughlin (2015:xv)

1. INTRODUCTION

In this paper I discuss a cluster of issues around the understanding of *robustness* in neuroscience. Systems biologist, Hiroaki Kitano defines

¹ M. Chirimuuta. History & Philosophy of Science, University of Pittsburgh. mac289@pitt.edu. Accepted for presentation at the 2016 Philosophy of Science Association meeting and publication of the proceedings in *Philosophy of Science*.

robustness as, “a property that allows a system to maintain its functions against internal and external perturbations” (Kitano 2004, p.826). According to this definition, in order to determine whether or not a system is robust, one must specify its function, and also specify the kinds of perturbation it faces. Empirically determinable questions then follow about how exactly the system achieves its robustness. Various means for making systems robust have been discussed across biology and neuroscience: copy redundancy, fail-safes, degeneracy, modularity, passive reserve, active compensation, plasticity, decoupling, and feedback (see Figure 1). It is obvious, but still worth emphasising, that most of these notions originate from engineering.

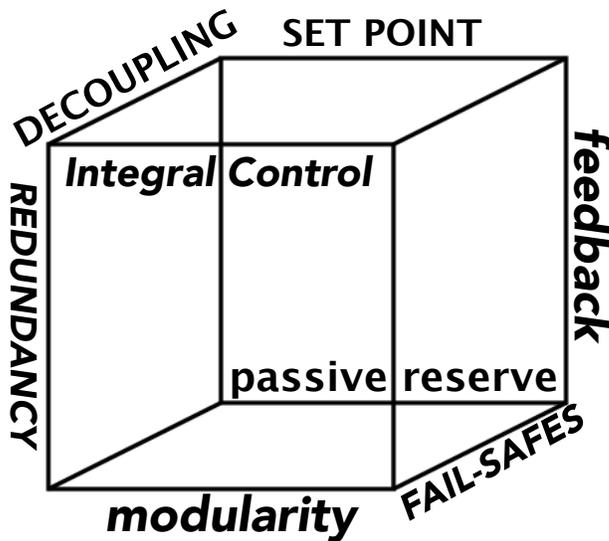


FIGURE 1. The Engineering Framework for Robustness. A set of terms originating from engineering and control theory, which are applied to biological systems to explain how they achieve robust performance.

In Section 2 of this paper I argue that the framework of concepts borrowed from engineering aids neuroscientists in (1) operationalizing robustness by specifying functions of the system and determining possible sources of

perturbation; (2) formulating hypotheses about means for the system to achieve robustness; and (3) showing how robustness may be precisely quantified. This will be shown with examples of neuroscientific research which aims to measure robustness in a retinal circuit (Sterling and Freed 2007), in the motor cortex (Svoboda 2015), and to develop models of homeostatic control (Davis 2006, O’Leary 2014).

In Section 3 I argue that the use of the engineering framework in neuroscience gets stretched, perhaps to breaking point, when applied to systems where (1) there is no principled distinction between processes for robustness and processes which continually maintain the life of the cell; (2) where perturbations are a regular occurrence rather than anomalous events; and (3) where one should not conceive of the system as seeking to maintain a steady state. This point will be illustrated through examination of some recent work from Eve Marder’s laboratory, one of the key centres for research on robustness in neuroscience.

I will argue that the limitations of the engineering notions are put into stark relief when one examines neural systems through the lens of the process approach to biology (Dupré 2012). The engineering perspective, to the extent that it treats biological systems as pre-specified objects with fixed functions, misses many of the features that make robust biological systems fascinating and which are highlighted by the process view.

In Section 4 I will consider if it is necessary to re-engineer the concepts of robustness to be more in line with the dynamicism of biological systems; or alternatively, if we should accept the engineering perspective as it is, as one amongst many idealizing and simplifying heuristics for understanding complex systems like the brain.

2. PUTTING THE ENGINEERING FRAMEWORK TO USE

The robustness of the brain is one of its many extraordinary attributes. By this I mean the fact that brains can undergo moderately severe external perturbations while still maintaining approximately normal function. Obviously, robustness has its limits and the brain's characteristic patterns of resilience and fragility are an important target of research (Sporns 2010, chap. 10). In order to investigate robustness it is necessary first to specify what sorts of perturbations the system is robust to, and then to quantify how robust it actually is. Explanations of robustness can be developed by testing hypotheses concerning the exact mechanisms by which robust performance is achieved. The engineering framework can be put to effective use in each of these processes.

For example, Sterling and Freed (2007) pose the question of how robust the retinal circuit is. They define robustness as the factor by which intrinsic capacity exceeds normal demand, which is the engineer's notion of margin of safety (p.563). The idea can be illustrated through their comparison with bridge design. An engineer designing a road bridge will consider both the anticipated normal demand (e.g. commuter traffic) as well as the unusual demands that might occasionally be placed on the bridge (e.g. the passage of a 30 ton military vehicle). The unusual demand can be thought of as a "perturbation" in Kitano's terms. A robust design will ensure that the system does not break when pushed beyond normal conditions. For a bridge this can be achieved with passive reserve (using thicker steel than is needed under normal conditions) and redundancy (including additional beams so that there are back-up structures if any parts are compromised).

Sterling and Freed take the bridge case to be analogous to the retinal circuit. Normal demand, for the retina, is the intensity of illumination that the eye will encounter under naturalistic stimulation conditions. The safety factor is calculated by experimental determination of the maximum illumination level under which neurons in the retina can maintain their ability to signal to downstream neurons. Sterling and Freed (2007, p.570) report that,

“across successive stages in this neural circuit, safety factors are on the order of 2–10. Thus, they resemble those in other tissues and systems. Their similarity across stages also accords with the principle of symmorphosis—that efficient design matches capacities across stages that are functionally coupled....”

Sterling and Freed’s explanation of robustness depends on the notion of passive reserve. For photoreceptor neurons, this is calculated as the number of vesicles of neurotransmitter available in their synapse for continuous signalling at high-rates without restocking of the vesicles (p.565-6). In arriving at their conclusion about retinal safety margins, they argue that there are at least twice as many vesicles as needed under normal stimulation conditions. In this case we have seen that a design approach borrowed from civil engineering plays a clear and striking role in these neuroscientist’s definition, operationalization and explanation of robustness in the retina.

Another example comes from Davis’s (2006) review of work on homeostatic regulation² in the nervous system. As he writes:

² Note that Davis makes a conceptual distinction between robust properties and properties under homeostatic control: “In general, robustness describes a system with a reproducible output, whereas homeostasis refers to a system with a constant output” (2006, p.308). I will ignore this difference for the purposes of the paper since homeostatic systems conform to Kitano’s general definition robust systems.

“Homeostatic control systems are best understood in engineering theory, where they are routinely implemented in systems such as aircraft flight control. Recently, biological signaling systems have been analyzed with the tools of engineering theory....” (p.314)

Accordingly, homeostatic control systems have a number of “required features”: 1) a set point which defines the target output of the system; 2) feedback; 3) precision in resetting the output back to the set point, following a perturbation; and (normally) 4) sensors which measure the difference between the actual output and the set point (p.309).

Thus control theory offers neuroscientists clear and experimentally testable criteria for determining whether a system undergoes homeostatic regulation, by looking for these required features (e.g. the existence of a set point) in a system. The operating conditions of homeostatic regulation, and the biophysical mechanisms of feedback, sensors, etc., are also open to experimental investigation. Reported examples of properties under homeostatic control are muscle excitation at the neuromuscular junction (p.309) and bursting properties of invertebrate neurons (p.311). More recently, O’Leary et al. (2014, p.818) argue that ion channel expression in their simplified model of invertebrate neurons can be understood as an implementation of *integral control*, a standard control-theoretic architecture.

Figure 2 (if space) schematic for integral control

3. CRASH TESTING THE FRAMEWORK

Before considering the question of whether the engineering framework becomes structurally unsound when applied to some kinds of neural systems, I would like to draw our attention to some of its features. The basic ideas are clearly illustrated in Sterling and Freed's (2007) example of the bridge. When one considers the robustness of an engineered artefact like the bridge, it is presupposed that the system is built up from component parts in such a way as to achieve a specific function. The robustness of the bridge is conceptually distinct from its other designed features or functions, and it can trade off against some of them. For example, the more robust the bridge is to the passage of the occasional heavy vehicle, the more expensive it will be to build (because requiring more steel) (p.563). Moreover, the perturbations against which the system is robust are thought of as atypical events, also conceptually distinct from the normal operations of the system.

There is also the tendency to think of robustness as allowing the system, following a perturbation, to return to its initial stable state. Some experiments specifically involve the operationalization of the robustness of a system as the reversion to a prior state. For example, reporting on an experiment in which mouse premotor cortex in one hemisphere was inhibited using optogenetics during the preparation period for the animal's movement, Svoboda (2015)³ writes, that "[t]his preparatory activity is remarkably robust to large-scale unilateral optogenetic perturbations: detailed dynamics that drive specific future movements are quickly and selectively restored by the network." This notion of robustness as the ability of the system to revert to a

³ To my knowledge, these results have not yet been published in a journal. I have contacted the author to find out if the study is under review or in press.

prior functional states is similar to the idea of *homeostasis* as the ability of a system to stabilize some quantity in spite of external changes.

Figure 3 (if space) After Kitano (2004, Figure 1)

Eve Marder's laboratory has carried out a long term investigation into the ability of neurons to maintain stable electrophysiological properties despite continual turnover of the ion channels embedded in the cell membrane which are responsible for its electrical excitability. This research project is one of the central examples of the study of robustness in neural systems. Marder and her collaborators make ample use of the engineering framework when reviewing other results and reporting their findings. For example, O'Leary et al. (2013, p.E2645) write:

“Both theoretical and experimental studies suggest that maintaining stable intrinsic excitability is accomplished via homeostatic, negative feedback processes that use intracellular Ca^{2+} concentrations as a sensor of activity and then alter[s] the synthesis, insertion, and degradation of membrane conductances to achieve a target activity level.”

What is striking about the characterization of electrophysiological stability in the face of ion channel turnover as a kind of robustness in the face of a perturbation (e.g. p.E2651), is the fact that the turnover is just part of the normal physiology of the cell. There is no functional and stable state of the cell in which this turnover does not occur—a fact which these authors also highlight.⁴ This brings our attention to some strains in the application of the engineering framework to this biological system.

⁴ “neurons in the brains of long-lived animals must maintain reliable function over the animal's lifetime while all of their ion channels and receptors are replaced in the membrane over hours, days, or weeks. Consequently, ongoing turnover of ion channels of various types must occur without compromising

In the basic engineering characterisation of robustness, sketched above, perturbations are different from the normal circumstances in which the system is expected to operate. “Perturbation” carries the everyday connotation of an event which throws the system off balance and is deleterious to its normal functioning. We cannot think of the events of ion channel turnover as perturbations in this sense; they are business as usual for the cell.

Furthermore, it is not in the nature of the system to seek to return to a prior, stable arrangement of its parts. A crucial property of the nervous system is its plasticity: the tendency for its component parts and the connections linking them to be continually sculpted by experience. The homeostatic mechanisms which Marder and colleagues investigate need to be understood as maintaining specific properties (such as a cell’s Ca^{2+} concentration) at a certain point, but not (nor do these researchers claim it) some generalised operation for achieving system-wide internal stability (see §4.4).

In the basic engineering conception of robustness, there is a clear conceptual distinction between the features of a system which allow it to carry out its intended function, and those which make the system robust (even if in reality one individual feature can serve both purposes). In the case of the neuron which has continual ion channel turnover and no definite stable state to return to following these “perturbations”, it is not clear that we can make this distinction. A more natural way to think about this and other biological systems is as ones, unlike engineered artefacts, “designed” to keep changing

the essential excitability properties of the neuron” (O’Leary et al. 2013, p.E2645).

and “designed” to maintain functional stability in the midst of this constant change.⁵

The tensions and strains associated with the application of the basic engineering framework to biological systems can be felt more sharply if we appeal to a process metaphysics of biological “things” (Dupré 2012). According to this view, organisms are not substances but *processes*—items whose existence depends on the taking place of certain changes. This highlights the fact that the life of organisms depends on a continual turnover of its component parts, and that the system as a whole, while living, persists longer than its parts. Yet features and functions of the organism remain relatively stable. For example, memories can endure for decades even though the neurons that form them have undergone material change. This stability must be achieved—somehow. And so processes for robustness are not cleanly distinct from the general maintenance processes which keep the organism alive.

The processual nature of neurons is nicely described by Marder and Goaillard (2012, p.563):

“each neuron is constantly rebuilding itself from its constituent proteins, using all of the molecular and biochemical machinery of the cell.”

(and see F n 4)

⁵ This blurring of the lines between mechanisms for robustness and mechanisms for life is highlighted by Edelman & Gally (2001: 13763) in their discussion of the difference between redundancy and degeneracy in biological systems: “the term redundancy somewhat misleadingly suggests a property selected exclusively during evolution, either for excess capacity or for fail-safe security. We take the contrary position that degeneracy is not a property simply selected by evolution, but rather is a prerequisite for and an inescapable product of the process of natural selection itself.” They also discuss another disanalogy between engineered and biological systems—the applicability of “design” talk.

We can contrast this with the substance metaphysics that we usually assume when thinking about engineered artefacts. A bridge or an aeroplane is what it is because of the parts which comprise it. Its existence does not depend on the occurrence of any process. This is not to deny that an expert in the theory of matter might well argue that the steel of the bridge maintains its integrity because of some fundamental processes. The point is that when characterising the robustness of the bridge or the aeroplane we would not resort to such sophistication. Rather, we think of the bridge as a substance and not a process—a steel structure which, in order to maintain its function in the face of perturbation, must resist rather than effect the swapping around of its component parts.

4. EXAMINING REASONS TO RE-ENGINEER

Now that we have noted these disanalogies between biological organisms and engineered things, we ought to worry that the framework borrowed from engineering is misleading when thinking about robustness in the brain and other biological systems. *Is it time to re-engineer our conceptual tools for thinking about robustness to make them more suitable for characterising living things?* In this section I consider four possible answers to this question.

4.1 *No. The terms in the engineering framework are just words that are used to facilitate communication of the neuroscientific results.*⁶

One potential response to the concerns raised in the previous section is that they stem from a superficial fixation on the vocabulary neuroscientists use when writing about their research. Just because the authors discussed above

⁶ A response along these lines was suggested to me by Timothy O’Leary, in conversation.

have employed certain words first introduced by engineers, it does not follow that their understanding of neurophysiology is distorted by comparisons with engineering. For example, I mentioned that the word “perturbation” has a negative connotation which makes it seem inappropriate when describing non-pathological and frequent events like ion channel turnover. It could well be that in the context of this research the term takes on a different meaning—for example, as any event that the system cannot directly control,⁷ such as changes in protein configuration due to thermal noise.

I believe that this response is warranted by what we know of the methodology of some of the investigations discussed above, but not all of them. In the case of Sterling and Freed (2007) I was careful to show that the engineering conceptions directly shaped how these neuroscientists operationalized and quantified robustness, and how they identified mechanisms by which robustness is achieved. There is no indication that they used terms such as “safety factor” to mean something radically different in the context of neuroscience.

A very explicit statement of the aim to apply engineering principles directly to the understanding of the premotor cortex comes from Svoboda (2015):

“preparatory activity is distributed in a redundant manner across weakly coupled modules. These are the same principles used to build robustness into engineered control systems. Our studies therefore provide an example of consilience between neuroscience and engineering.”

Thus the convergence between a neurophysiological and the engineering perspective on the mouse motor planning system is taken to be an important result of this study. This echoes Sterling and Laughlin’s (2015, pp. xiii-xv) proposal that enquiring to see how engineering principles are implemented in

⁷ I thank Timothy O’Leary for this suggestion.

neural systems, and the attempt thereby to reverse-engineer the brain, leads to insights not otherwise available through routine data collection.

4.2 *No. The inadequacies you point out with the engineering framework are based on a caricature of mechanical engineering, not the actual complex discipline.*⁸

My characterisation of the engineering framework assumes that mechanical engineering (the design of bridges, aeroplanes and such like) is paradigmatic of the engineering approach in general. But of course there are many different kinds of engineering, from mechanical to electronic to communications and chemical. It could well be that the mismatch between understanding the robustness of a highly dynamic entity like the brain, and the rather static conception of robust objects that falls out of the basic engineering framework is just an artefact of only focussing narrowly on the kind of engineering that is actually furthest away from neuroscience.

It would take me beyond the scope of this short article (and well beyond my own knowledge of the subject) to sketch out the various possible frameworks associated with each field of engineering specifically, and to see which conception of robustness is most suitable for biology. However, what I will say is that there is evidence in the studies discussed above that neuroscientists themselves do sometimes draw from the mechanically based caricature. This is particularly true of Sterling and Freed (2007). In contrast, when Davis (2006) and O’Leary (2014) make direct appeal to engineering they refer specifically to models in control theory.⁹ This invites questions, still, about whether the paradigm examples of controlled systems (e.g. a car driven on

⁸ This concern was raised by Arnon Levy and Timothy O’Leary.

⁹ See also Zhang and Chase (2015) on the physical control system perspective on brain computer interfaces for motor rehabilitation.

cruise control, a Watt governor, or an aeroplane flown on autopilot) are dynamical enough capture the processual nature of the nervous system.

4.3 *Yes. The brain is so different from an engineered artefact that the framework is misleading and inappropriate.*

In Sections 4.1 and 4.2 I discussed two reasons for thinking that we should not be concerned about any radical disanalogy between robustness in biological and engineered systems. While I agree that these are important points to keep in mind, I do not think that they diffuse the fundamental concern that when neuroscientists borrow engineers' terms in order to study robustness, they risk mischaracterising the brain as more like an engineered artefact than it actually is. Is the appropriate conclusion, then, that a neural circuit is so different from a bridge or an aeroplane that the engineering framework is simply misleading and should be discarded?

The best way to make this strong negative case is to consider some historical examples in which reasoning by analogy with engineered systems seems to have lead neuroscientists and theorists astray. One example comes from von Békésy, a physicist and communications engineer who turned his attention to inhibition in the nervous system. In his book *Sensory Inhibition* he notes that there are feedback loops everywhere in nervous system and he asks how it is that system manages to avoid ending up in a dysfunctional oscillatory state (1967, p.25). It seems that von Békésy is importing his understanding of systems containing feedback from engineering, and in that context oscillations are normally problematic and efforts must be made to dampen them. These days neuroscientists seek to understand how oscillations in the healthy brain (i.e. its characteristic patterns of endogenous activity) are actually responsible for cognitive functions, and how these oscillations differ

from the ones associated with pathologies such as epilepsy and Parkinson's disease.¹⁰

Another example is the comparison of the effects of “noise” in brains and artificial signalling systems..... GET EXAMPLE

This is very different from how neuroscientists understand noise today, which begins with the idea that brains evolved under constraints imposed by noisy “components”, which has therefore shaped all aspects of neural computation (Faisal et al. 2008). It would be a mistake to think of the brain processing information in the same way as an electronic computer, but with added redundancy to offset the noisiness of individual processing streams.

The cautionary tales just told give some concrete indications of how imposition of the engineering framework on to neural systems can lead to conclusions which in retrospect appear false and misguided. But it would be too hasty infer from these two examples that current work on robustness in neuroscience is of dubious standing whenever it appeals to the concepts of engineering. A more general argument is the following: *the brain is not like a bridge (or a computer, or an aeroplane on autopilot...); therefore whenever neuroscientists appeal to terms borrowed from the analysis of such systems, they risk saying things that are simply false because they fail to notice relevant disanalogies*. This lays all the sceptical cards on the table. In the last part of the paper I attempt to mitigate these worries.

¹⁰ For a scientific overview see Buzsáki (2006). For discussion of philosophical implications, see Bechtel and Abrahamsen (2013). See also Knuuttila and Loettgers (2013, p.160) on a parallel difference across engineering and cell biology, where oscillations are found to have a functional role.

4.4 *No. Use of the engineering framework should be thought of as a simplifying strategy.*

Neuroscientist Steven Rose (2012:61) writes that:

“one of the most common but misleading terms in the biology student’s lexicon is homeostasis....[the] concept of the stability of the body’s internal environment. But such stability is achieved by dynamic responses; stasis is death, and homeodynamics needs to replace homeostasis as the relevant concept”¹¹

This seems to capture the problem that was first noted in Section 3, that we should not be misled by the engineering framework into thinking of neural systems as seeking to maintain an initial stable state. But we also noted that the neuroscientists employing control-theoretic models of homeostatic mechanisms are not thinking of their systems as seeking stability in this very general way. Instead, they are modelling the stability of a specific variable—in the case of O’Leary et al. (2014), the concentration of Ca^{2+} —and investigating the mechanisms by which it is controlled. To this end, it is reasonable to interpret the system as an integral controller (p.818).¹² Thus it is still useful to talk about homeostasis with respect to Ca^{2+} concentration, even while thinking of the system as a whole, and in reality, as a “homeodynamic” one.

¹¹ Compare Sterling (2012) on the concept of *allostasis* – stability through change with an emphasis on predictive regulation. Day (2005) and O’Leary and Wyllie (2011), in contrast, argue that the concept of homeostasis easily accommodates these dynamic and predictive aspects, and that the term *allostasis* is therefore superfluous. It is an interesting question (but beyond the scope of this paper) whether the narrow or wide definition of *homeostasis* is currently more prevalent amongst biologists and neuroscientists.

¹² Note that O’Leary et al. (2014) study of homeostasis is via a *model* of a neuron. But the model is realistic enough that it is expected to shed light on actual biophysical mechanisms.

I think of neuroscientists whose investigation of robustness in the brain is scaffolded by the engineering framework as providing *idealized mechanistic explanations*. Their explanatory target is, for example, the process by which overall neuronal activity level is controlled via regulation of ion channel gene transcription through a Ca^{2+} sensitive feedback loop. This is standard fodder for mechanistic explanation. At the same time, the framework of engineering—in this case the schematic of the integral controller—serves to direct attention to specific parts and processes in the extremely complex cellular machinery and to interpret them in control theoretic terms (sensors, feedback loops, etc.), while bracketing other aspects not immediately relevant to the explanation of robustness.

Bechtel (2015, p.92) has presented the case that:

“mechanisms are [to be] viewed not as entities in the world, but as posits in mechanistic explanations that provide idealized accounts of what is in the world.”

His example is the idealization (understood as “falsehood”) that scientists introduce by putting boundaries around putative mechanisms which in nature do not exist. In the cases explored in this paper, the idealization comes in through the analogical reasoning of treating a neuronal system *as if* it is an engineered artefact. This, like the positing of boundaries, is a useful way to simplify the explanandum. It enables neuroscientists to bracket some of the known facts about the brain’s messy, Heraclitean nature. But it means, perhaps, that there is a stark difference between the brain viewed *sub specie aeternatis* (what some neuroscientists call the “ground truth” of the brain) and viewed *sub specie mechinae* (in the guise of a machine).

ACKNOWLEDGEMENTS

I am greatly indebted to Timothy O’Leary, Nancy Nersessian and Peter Sterling for their feedback on this work. I would also like to thank the

participants of the Fall 2015 workshop on Robustness in Neuroscience for discussion of the ideas behind this paper, and the audience at the Spring 2016 Re-Engineering Biology conference for their questions and comments on it. Both of these events were hosted by the Philosophy of Science Center at the University of Pittsburgh.

REFERENCES

Bechtel, W. and Abrahamsen, A. (2013). Thinking dynamically about biological mechanisms: Networks of coupled oscillators. *Foundations of Science*, 18:707–723

Bechtel, W. (2015). Can mechanistic explanation be reconciled with scale-free constitution and dynamics? *Studies in History and Philosophy of Science Part C*, 53: 84–93.

von Békésy, G. (1967). *Sensory Inhibition*. Princeton, NJ: Princeton University Press.

Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford: Oxford University Press

Davis, G.W. (2006). Homeostatic control of neural activity: from phenomenology to molecular design. *Annu. Rev. Neurosci.* 29, 307–323

Day TA (2005). Defining stress as a prelude to mapping its neurocircuitry: no help from allostasis. *Prog Neuropsychopharmacol Biol Psychiatry* 29, 1195–1200

Dupré, J. (2012) *Processes of Life*. Oxford: Oxford University Press

Edelman GM, Gally JA (2001) Degeneracy and complexity in biological systems. *Proc Natl Acad Sci USA* 98(24):13763–13768.

Faisal, A., L. P. J. Selen and D. M. Wolpert (2008) Noise in the Nervous System. *Nature Reviews Neuroscience*. 9:292-303.

Kitano, H. (2004) Biological robustness. *Nature Reviews Genetics*. 5: 826-837.

Knuuttila, T. and A. Loettgers (2013). Basic science through engineering? Synthetic modeling and the idea of biology-inspired engineering. *Studies in History and Philosophy of Science, Part C* 48, 158–169.

Marder, E. and Goaillard, J.-M. (2012) Variability, compensation and homeostasis in neuron and network function. *Nature Reviews Neuroscience*. 7:563-574

von Neumann, J. (2000). *The Computer and the Brain*. New Haven: Yale University Press.

O’Leary T, Williams AH, Caplan JC, Marder E (2013) Correlations in ion channel expression emerge from homeostatic tuning rules. *PNAS*. 110(28): 809–821

O’Leary T, Williams AH, Franci A, Marder E (2014) Cell types, network homeostasis and pathological compensation from a biologically plausible ion channel expression model. *Neuron* 82(4): E2645–E2654.

O’Leary, T. and D. J. A. Wyllie (2011) Neuronal homeostasis: time for a change? *J Physiol* 589.20:4811–4826

Rose, S. (2012). The need for a critical neuroscience. In *Critical Neuroscience: A Handbook of the Social and Cultural Contexts of Neuroscience*, S. Choudhury and J. Slaby (eds.) Hoboken, NJ: Wiley-Blackwell.

Sporns, O. (2010). *Networks of the Brain*. Cambridge, MA: MIT Press.

Sterling, P. and M. Freed (2007). How robust is a neural circuit? *Visual Neuroscience*, **24**, 563–571.

Sterling, P. (2012). Allostasis: A model of predictive regulation. *Physiology & Behavior* 106:5–15

Sterling, P. and S. B. Laughlin (2015). *Principles of Neural Design*. Cambridge, MA: MIT Press.

Svoboda, K. (2015). Probing Frontal Cortical Networks during Motor Planning. Abstract, Center for the Neural Basis of Cognition, 10 November 2015. <http://www.braininstitute.pitt.edu/event/probing-frontal-cortical-networks-during-motor-planning>

Zhang, Y. and S. M. Chase (2015). Recasting brain-machine interface design from a physical control system perspective. *J Comput Neurosci* 39:107–118