OPINION ARTICLE

Puzzles in modern biology. V. Why are genomes overwired?

[version 1; peer review: 1 approved, 1 approved with reservations]

Steven A. Frank

Department of Ecology and Evolutionary Biology, University of California, Irvine, CA, 92697-2525, USA

v1 First published: 16 Jun 2017, 6:924
<https://doi.org/10.12688/f1000research.11911.1>
Latest published: 14 Aug 2017, 6:924
<https://doi.org/10.12688/f1000research.11911.2>

Abstract

Many factors affect eukaryotic gene expression. Transcription factors, histone codes, DNA folding, and noncoding RNA modulate expression. Those factors interact in large, broadly connected regulatory control networks. An engineer following classical principles of control theory would design a simpler regulatory network. Why are genomes overwired? Neutrality or enhanced robustness may lead to the accumulation of additional factors that complicate network architecture. Dynamics progresses like a ratchet. New factors get added. Genomes adapt to the additional complexity. The newly added factors can no longer be removed without significant loss of fitness. Alternatively, highly wired genomes may be more malleable. In large networks, most genomic variants tend to have a relatively small effect on gene expression and trait values. Many small effects lead to a smooth gradient, in which traits may change steadily with respect to underlying regulatory changes. A smooth gradient may provide a continuous path from a starting point up to the highest peak of performance. A potential path of increasing performance promotes adaptability and learning. Genomes gain by the inductive process of natural selection, a trial and error learning algorithm that discovers general solutions for adapting to environmental challenge. Similarly, deeply and densely connected computational networks gain by various inductive trial and error learning procedures, in which the networks learn to reduce the errors in sequential trials. Overwiring alters the geometry of induction by smoothing the gradient along the inductive pathways of improving performance. Those overwiring benefits for induction apply to both natural biological networks and artificial deep learning networks.

Keywords

Gene regulation, complex traits, artificial intelligence, deep learning, induction

Open Peer Review

Approval Status

	1	2
version 2		
(revision)		
14 Aug 2017		view
version 1		
16 Jun 2017	view	view

1. **Sean Nee** , The Pennsylvania State University, University Park, USA

2. **David M. McCandlish**, Cold Spring Harbor Laboratory, Cold Spring Harbor, USA

Any reports and responses or comments on the article can be found at the end of the article.

Corresponding author: Steven A. Frank (safrank@uci.edu)

Author roles: Frank SA: Conceptualization, Funding Acquisition, Investigation, Writing – Original Draft Preparation, Writing – Review & Editing

Competing interests: No competing interests were disclosed.

Grant information: National Science Foundation (grant DEB-1251035) supports my research.

The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Copyright: © 2017 Frank SA. This is an open access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

How to cite this article: Frank SA. **Puzzles in modern biology. V. Why are genomes overwired? [version 1; peer review: 1 approved, 1 approved with reservations]** F1000Research 2017, 6:924 <https://doi.org/10.12688/f1000research.11911.1>

First published: 16 Jun 2017, 6:924 <https://doi.org/10.12688/f1000research.11911.1>

Introduction

What determines gene expression? The list keeps growing: transcription factors, methylation, histone codes, DNA folding, intron sequences, RNA splicing, noncoding RNA, and others^{1,2}.

Hundreds of genomic variants affect human traits, such as height³. Consider pathways of influence. Numerous factors affect gene expression. Many genes affect a trait. Vast wiring connectivity links genomic influence to a trait.

An engineer following classic principles of control theory would design a simpler system with fewer connections⁴. Genomes are *overwired*. They have far more nodes and connections than classically engineered systems.

Why are genomes overwired? I discuss possible causes. I then consider wiring density more broadly. What other sorts of systems tend to be overwired?

Computational neural networks in artificial intelligence stand out. Deeply, densely connected computational networks pervade modern life. New computational systems often outperform humans.

The recent computational concepts and methods comprise *deep learning*. The *learning* simply means using data, or past experience, to improve classification of inputs and adjustment of response. The *deep* qualifier refers to the multiple layers of deep and dense network connections^{5,6}.

That wiring depth, and the computational techniques to use vast connectivity, triggered the revolutionary advances in performance. I discuss genomic wiring in relation to deep learning. I suggest that the inductive systems of biological adaptation and computation learning gain in similar ways from diffusely and densely wired networks.

Causes

Why do so many factors modulate gene expression? Why is the regulatory network architecture for traits often complex?

Neutral accumulation

A noncoding RNA may, by chance, alter the expression of various genes. Small modulations of expression may have relatively little effect on fitness. If so, a novel noncoding RNA variant may be effectively neutral. Nearly neutral variants accumulate by chance.

Many nearly neutral variants may accumulate over time. As each variant spreads, it changes the genomic environment of gene regulation. When the aggregate effect of many nearly neutral variants becomes significant, natural selection will retune expression to compensate.

After compensation occurs, one cannot remove the layers of accumulated modulating factors without causing deleterious changes in gene expression. What began as neutral accumulation becomes integral to genomic function. Wiring complexity increases irreversibly.

Lynch's neutral theory of genome architecture makes predictions^{7,8}. Smaller population sizes increase chance fluctuations. Greater fluctuations allow larger fitness effects to become nearly neutral. Broader neutrality enhances the rate at which changes accumulate. Smaller populations may tend toward overwiring.

By contrast, large populations more efficiently prune small effects on fitness. Small modulations of gene expression accumulate more slowly. Larger populations may not overwire as readily as smaller populations.

If the fitness effects of modulation tend to be larger, nearly neutral variants will be less common. Prokaryotes may tend to have relatively large deleterious fitness effects of novel modulating factors, because increased genome size and complexity may slow the speed of cellular replication. Eukaryotic genomes may be less sensitive to size and complexity because organismal replication is less strongly coupled to speed of cell division.

Overall, prokaryotes tend to have larger populations and greater sensitivity to genome size and complexity. Such characteristics restrict the scope for neutral accumulation and overwiring. By contrast, eukaryotes tend toward smaller populations and less sensitivity to genome size and complexity. Those characteristics favor neutral accumulation and overwiring. Stronger predictions arise when one can compare closely related organisms that differ in population size and genomic sensitivity.

Robustness and decay

Modulating factors combine to influence traits. The mechanism of combination matters. Consider two alternatives.

First, suppose modulating factors add together to determine a trait. Then, the more modulating factors, the greater the trait's variance. Put another way, the more things that cause fluctuations in gene expression, the more variable the trait. In the classical summation model, the variance contribution of each factor is σ^2 . Summing n components yields a trait variance of $n\sigma^2$, rising with the number of components.

Second, suppose modulating factors average together to determine a trait⁹. When averaging n components, we divide the effect of each component by n . As the number of components rises, the effect of each component declines. Averaging n components yields a trait variance of σ^2/n , declining with the number of components.

One can think about each additional modulating component as perturbing trait expression. Robustness is decreased sensitivity to perturbation. In the averaging model, the greater the number of factors, the weaker the effect of each individual perturbing factor. Thus, averaging reduces sensitivity to each perturbation, enhancing robustness.

If modulating factors average together, the benefits of enhanced robustness can favor an increase in the number of factors⁹. Generally, if the effect of an additional factor causes a sufficient decline in the average contribution of each factor, then natural selection can

favor a tendency for the number of factors to increase. Ultimately, many factors of small effect modulate trait expression.

Under the averaging model, evolutionary dynamics follows an interesting path. An additional modulating factor may be favored because it reduces sensitivity to perturbation. Once the new factor is added and sensitivity is reduced, selective intensity against perturbations weakens. Weaker selection allows the accumulation of additional mutations with larger perturbing effects. That shift in mutation-selection balance causes a decay in the average fitness effect of each factor.

Dynamics progresses like a ratchet^{10,11}. New factors get added for their enhanced robustness. All factors then decay. Taking away a recently added factor exposes the increased deleterious effects of the remaining factors. Exposure of those deleterious effects opposes reversal. One cannot go back.

Gradient smoothing

Hundreds of genomic variants influence traits, such as human height and weight. Most variants have small effects. Many small effects smooth the gradient of trait values.

A smooth gradient means that a trait may potentially change steadily, or monotonically, with respect to underlying genomic changes. We may think of a smoothly increasing path from a starting point up to the highest peak or down to the lowest valley.

Overwiring leads to many genomic variants of small effect, which in turn smooths the gradient. Thus, we may say that overwiring causes a smooth gradient. What about the converse? Do the benefits of a smooth gradient favor overwiring? Consider three potential benefits.

A smooth gradient enhances adjustability. A densely wired regulatory network has many different connections that can alter traits by a small amount. Such overwired connectivity allows inputs to modulate expression smoothly.

A smooth gradient promotes learning¹². Learning requires adjustment in response to input and measurement of success. A system learns as it steadily climbs the gradient of success by smoothly adjusting expression in response to inputs.

A smooth gradient boosts evolutionary adaptability^{13,14}. Natural selection is essentially a trial and error learning algorithm. The advantages of densely overwired control for learning apply to evolutionary adaptation by natural selection.

The smooth gradient benefits of adjustability, learning, and adaptability can potentially favor overwiring.

Deep learning

Systems can easily adjust, learn, and evolve if they have smooth gradients. Many of the algorithmic tricks and underlying concepts of machine learning and artificial intelligence come down to how one smooths the gradient^{5,6}. A smooth gradient provides a steadily improving path from the starting point to an improved target point.

Some biological networks may be densely wired because of the benefits of gradient smoothing. Ideally, we could analyze how network architecture and connectivity strengths affect gradients. However, we do not yet know enough about the details of biological networks. By contrast, the study of computational networks has advanced greatly in recent years. Those advances in computational studies hint at some principles of networks and gradient smoothing. Those principles provide clues about the design of biological networks by natural selection.

Computational networks are loosely modeled after biological neural networks. A set of nodes takes inputs from the environment. Each input node connects to another set of nodes. Each of those intermediate nodes combines its inputs to produce an output that connects to yet another set of nodes, and so on. The final nodes classify the environmental state, possibly taking action based on that classification.

A network learns by altering its parameters^{5,6}. The parameters set the connection strength between nodes, and how individual nodes combine their many inputs to determine the strength of their output. For example, the input to a network may be an image of a numerical digit. The input nodes are sensors that react to the image. Those sensors initiate activations that pass through all of the connections and layers of the network. The final layer provides a set of ten probabilities, one probability for each of the digits 0, 1, ..., 9.

The network, when presented with an image of the digit 7, classifies the image by returning a set of ten probabilities. The optimal classification is a probability of one for 7 and zero for all other digits. We can calculate an error distance between the optimal classification and the network's guess. An error distance is a function of the differences in the probabilities of the optimal and guessed classification.

The error distance can be used to update the network's parameters. We find a set of small changes in the network parameters that would have yielded a small reduction in the error distance. By following this gradient of improving performance, the network may learn from experience.

That learning approach works as long as there is a smooth path of increasing performance. Improved performance means that the adjustment process truly learns the general features of digit images that enhance future classification. Performance does not improve if adjustments focus on unusual features of the digit images used to train the network. Those unusual features may not be present in many other digit images.

A deep neural network has many layers of nodes between initial inputs and final outputs. Until recently, deep and densely connected computational networks often learned slowly and then got stuck, unable to learn from further information.

Getting stuck often means an unsmooth gradient. Initially, the system learns. It uses past trials to adjust its parameters, yielding a reduction in the error distance for future trials. Then the system gets stuck. Parameter adjustments do not improve future performance.

Put another way, initially the system descended smoothly along the error gradient, improving performance as the error became smaller. Then the gradient flattened out, so that adjustments of the parameters either did not change future error or increased future error.

From that stuck location of parameters, there are no easily discovered altered parameters that follow a smoothly continuing path to a lower point on the error gradient. Other parameter combinations with better performance often exist. But there is no smoothly descending path on the error gradient from the current location to those better combinations.

An improved learning system means a system that smooths the gradient sufficiently, descending on the error gradient to the better locations. The recent revolutionary increase in the performance of deep learning networks arose from a variety of computational adjustments. Many of those adjustments were discovered by trial and error, simply finding that they worked well on real problems^{5,6}.

For example, limiting the connection strength between nodes prevents dominance by a small set of pathways of connectivity. It seems that broad, densely connected networks that retain many pathways of connectivity have greater learning potential. In essence, a deep, densely and broadly connected network provides a robustly smoothed gradient.

Other adjustments include the functions by which individual nodes combine inputs to determine output. No available theory describes exactly how to construct such functions. Again, trial and error has shown certain functions to work well. Most likely, those successful functions enhance the breadth of pathways that can adjust by small amounts in response to new information, again smoothing the gradient.

Network architecture also affects performance. Architecture includes the number of layers of nodes and the manner in which nodes connect. Connections feed forward from inputs to outputs or feed back from later nodes toward earlier nodes. The feature detectors in the sensory input nodes set the initial representation of environmental states. The network generalizes that low-level representation as information passes through the network layers.

Presumably, architecture and representation ultimately contribute to performance through better gradient smoothing. In a sense, better capacity to learn and better gradient smoothing are nearly the same thing. But the emphasis on gradient smoothing can be useful, because it calls attention to the mechanisms by which particular network properties may contribute to better performance.

Over time, we may come to understand the mechanisms that improve performance and smooth gradients in deep learning networks. We can then consider how those advances in computational networks may provide insight into genomic network architecture, sensory representation, and the consequences for gradient smoothing.

We know that densely connected computational and biological neural networks perform spectacularly at learning, and that densely connected genomic networks perform spectacularly in terms of

adjustability and evolvability. We are still trying to understand why.

Geometry of induction

The spectacular performance of large densely wired networks hints at key underlying principles. I conclude by suggesting that large networks are particularly good at smoothing gradients in a way that facilitates induction. Before turning to induction, it is useful to consider deductive principles.

Control theory deduces general principles of wiring to achieve particular design goals⁴. For example, simple feedback often keeps a system near a setpoint. The setpoint may be a fixed temperature or a fixed concentration. Deviation of the output from the setpoint is fed back to the system as an additional input to the controller. If the feedback signal tells the system that it is below its setpoint, the controller triggers increased output.

Many examples of genomic wiring follow simple feedback^{15–17}. Other classic control theory motifs also occur frequently in genomic wiring pathways¹⁸. The deductive theoretical principles of control successfully predict key aspects of genomic wiring.

However, more complex challenges in engineering and in genomes often seem to be solved by deeply, densely wired networks. I call those networks *overwired*, in the sense that their connectivity patterns are much deeper, denser and broader than predicted by classical deductive principles.

Overwired systems may have embedded within them feedback loops and other classic wiring motifs. But those motifs no longer act alone in a simply interpreted manner. Instead, they are enmeshed within such a large web of diffuse connectivity that it is often difficult to trace their particular effects and functions.

Why do some systems wire simply along classical deductive lines and other systems overwire? I have argued that overwired systems smooth gradients to allow adjustability and adaptability. Put another way, such networks can change in response to experience. A sequence of specific events can lead to improvement of future performance. The networks somehow use their specific experience to find general solutions to a challenge. The networks inductively use specific examples to learn general solutions.

Inductive improvement often requires a smooth gradient. Overwiring may be favored because it enhances the scope for small changes in parameters to descend smoothly along a gradient of decreasing error.

The problem is essentially geometric. How do topological changes in network architecture reshape the error gradient? How do particular bounds on connectivity parameters smooth the gradient? How do particular nodal transformations of inputs into outputs alter gradient shape? How do the input sensors and input representations change the error gradient and consequent inductive performance?

Inductive improvement occurs on various timescales. Over short periods of time, an organism may adjust its response to the

environment by changing various parameters within its regulatory network. Over long periods of time, natural selection reshapes the design of the regulatory network. Both short-term adjustments and long-term changes in design arise inductively. Biological systems do not deduce principles. They inductively arrive at abstract representations of environmental challenges. They narrow the error distance along the geometric path of inductive improvement.

Many biological regulatory networks are simple, following closely along classical deductive design principles. In those cases, inductive evolutionary processes discovered those simple deductive principles. Other biological networks are overwired, apparently tuned for inductive potential.

Final questions arise. What sorts of environmental challenges favor classically deductive wiring? What sorts of challenges favor

inductive overwiring? What historical aspects of organismal evolution constrain network design? How can we relate deep learning solutions of engineering problems and genomic wiring solutions of biological problems to a more general geometric theory of induction?

Competing interests

No competing interests were disclosed.

Grant information

National Science Foundation (grant DEB-1251035) supports my research.

The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

References

- Alberts B, Johnson A, Lewis J, *et al.*: **Molecular Biology of the Cell**. Garland Science, New York, 6th edition, 2014.
[Reference Source](#)
- Pollard TD, Earnshaw WC, Lippincott-Schwartz J, *et al.*: **Cell Biology**. Elsevier, San Diego, 3rd edition, 2017.
[Reference Source](#)
- Wood AR, Esko T, Yang J, *et al.*: **Defining the role of common variation in the genomic and biological architecture of adult human height**. *Nat Genet*. 2014; **46**(11): 1173–1186.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Ogata K: **Modern Control Engineering**. Prentice Hall, New York, 5th edition, 2009.
[Reference Source](#)
- Nielsen MA: **Neural Networks and Deep Learning**. Determination Press, 2015.
[Reference Source](#)
- Goodfellow I, Bengio Y, Courville A: **Deep Learning**. MIT Press, Cambridge, MA, 2016.
[Reference Source](#)
- Lynch M: **The Origins of Genome Architecture**. Sinauer Associates, Sunderland, MA, 2007.
[Reference Source](#)
- Fernández A, Lynch M: **Non-adaptive origins of interactome complexity**. *Nature*. 2011; **474**(7352): 502–505.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Frank SA: **Genetic variation of polygenic characters and the evolution of genetic degeneracy**. *J Evol Biol*. 2003; **16**(1): 138–142.
[PubMed Abstract](#) | [Publisher Full Text](#)
- Frank SA: **Maladaptation and the paradox of robustness in evolution**. *PLoS One*. 2007; **2**(10): e1021.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Frank SA: **Evolution of robustness and cellular stochasticity of gene expression**. *PLoS Biol*. 2013; **11**(6): e1001578.
[PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
- Bengio Y: **Learning deep architectures for AI**. *Foundations and Trends in Machine Learning*. 2009; **2**(1): 1–127.
[Publisher Full Text](#)
- Gavrilets S: **Fitness Landscapes and the Origin of Species**. Princeton University Press, Princeton, NJ, 2004.
[Reference Source](#)
- Frank SA: **Natural selection. II. Developmental variability and evolutionary rate**. *J Evol Biol*. 2011; **24**(11): 2310–2320.
[PubMed Abstract](#) | [Publisher Full Text](#)
- Alon U: **An Introduction to Systems Biology: Design Principles of Biological Circuits**. CRC press, Boca Raton, Florida, 2007.
[Reference Source](#)
- Iglesias PA, Ingalls BP: **Control Theory and Systems Biology**. MIT Press, Cambridge, MA, 2009.
[Reference Source](#)
- Cosentino C, Bates DG: **Feedback Control in Systems Biology**. CRC Press, Boca Raton, Florida, 2011.
[Reference Source](#)
- Alon U: **Network motifs: theory and experimental approaches**. *Nat Rev Genet*. 2007; **8**(6): 450–461.
[PubMed Abstract](#) | [Publisher Full Text](#)

Open Peer Review

Current Peer Review Status:

Version 1

Reviewer Report 24 July 2017

<https://doi.org/10.5256/f1000research.12874.r23733>

© 2017 McCandlish D. This is an open access peer review report distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

David M. McCandlish

Simons Center for Quantitative Biology, Cold Spring Harbor Laboratory, Cold Spring Harbor, NY, USA

The outsides of organisms are most often exquisitely, even ruthlessly, adaptive. Inside the organism's body, the situation is more heterogeneous. Both physiology and the function of macromolecular complexes are in many instances technically stunning. On the other hand, the wiring diagram of the cell is bedlam. Even based on our current — and likely quite incomplete — state of knowledge, regulatory networks appear to be both more densely and more broadly interconnected than would seem necessary. This is surely a puzzle of modern biology, and Frank has done us a service by cataloging the live hypotheses and pointing us towards the possibility of a resolution wherein this “over-wiring” simply reflects general principles of inductive inference.

My main criticism of this article is that it does not engage sufficiently strongly with either the contemporary or historical literature. On the contemporary side, it would seem appropriate to directly address the ongoing efforts by several groups to formally link population genetics to general principles of inference (of course, Frank has contributed substantially in this area himself by clarifying the relationship between natural selection and information geometry, see e.g. Frank 2012 “Natural selection V. How to read the fundamental equations of evolutionary change in terms of information theory”). These efforts have been recently reviewed by Watson and Szathmari 2016 in a TREE piece “How can evolution learn?”, which hits on many of the same themes as the latter half current manuscript. Watson's work in this area seems particularly relevant, and indeed he calls his theory “Evolutionary connectionism” (Watson *et al.* 2015). An important insight from this series of papers is a possible relationship between the evolutionary problem of evolvability and the statistical problem of overfitting. In particular, they suggest that pressure for developmental simplicity can improve the ability of evolutionary systems to generalize in a manner similar to how regularization, drop outs, or early stopping can prevent over-fitting in machine learning (e.g. Kouvaris *et al.* 2017 “How evolution learns to generalize, using the principles of learning theory to understand the evolution of developmental organisation”). The idea that the topology of regulatory networks is a generic consequence of evolution by gene duplication (as in, e.g. the work of Ricard Solé), and more generally by the expansion of gene families, also seems like it deserves a mention as at least a possible proximal

cause of over-wiring.

On the historical side, I think more could be done to link the current discussion with historical themes in evolutionary thought. For instance, the discussion about many possible genomic changes with small effects smoothing the gradient and allowing evolutionary optimization could be put in the context of Fisher and Wright's disagreements over the structure of fitness landscapes. Wright thought that the reality of building a functional physiology would produce fitness landscapes with many local maxima, so that the key question in evolution was to identify the population-genetic regimes where progress on such a landscape is possible (Wright 1931, 1932). Fisher thought that in high dimensions, these local maxima would largely turn into saddle points, and that in any case, environments were generally changing fast enough that populations were usually chasing a moving optimum rather than adapting on a fixed fitness landscape (Fisher 1930). Frank's discussion of "getting stuck" in the current manuscript provides additional nuance to this classical disagreement by emphasizing the possibility of extended, high-dimensional plateaus that while strictly speaking are saddle points function in an evolutionary sense more like local optima. The reader interested in resolving this puzzle should also be directed to some of the stone-cold classics in this area such as Wagner and Altenberg 1996 and Stoltzfus 1999 ("On the possibility of constructive neutral evolution").

Is the topic of the opinion article discussed accurately in the context of the current literature?

Partly

Are all factual statements correct and adequately supported by citations?

Yes

Are arguments sufficiently supported by evidence from the published literature?

Yes

Are the conclusions drawn balanced and justified on the basis of the presented arguments?

Yes

Competing Interests: No competing interests were disclosed.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response (F1000Research Advisory Board Member) 05 Aug 2017

Steven Frank, Department of Ecology & Evolutionary Biology, University of California, Irvine, USA

I thank David McCandlish for his thoughtful comments and his broad knowledge of the literature. I have added an Appendix that includes the suggested references plus a few others. I agree that there is much prior work and relevant literature. In this series of articles, I am writing to a strict word limit and in a way that can be read by the beginning student as

well as the advanced professional. In this case, I perhaps had limited my references too much. The Appendix should provide a start for those who wish to follow up in the literature.

Competing Interests: No competing interests were disclosed.

Reviewer Report 03 July 2017

<https://doi.org/10.5256/f1000research.12874.r23536>

© 2017 Nee S. This is an open access peer review report distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Sean Nee

Department of Ecosystem Science & Management, The Pennsylvania State University, University Park, PA, USA

This is a stimulating, original and thought-provoking and, so, I recommend publication. I see nothing incorrect, so no changes are requested by me as a referee.

Some thoughts provoked in my mind are as follows.

First: is engineering really as cut-and-dried as we suppose? A cursory reading of fly-by-wire disasters suggests that the elegant theorems of classical control theory may not be as powerful as one would wish. My impression is that engineers are acutely aware that new technologies such as "cyber physical" systems¹, which are most akin to biological systems, are necessitating a complete rethink of the conceptual foundations of their subject matter. Even for more traditional technologies, it is not obvious to me that engineering is as purely "deductive" a subject as we might like to think as we board an aircraft, using the word "deductive" in a way that I may be misconstruing as Frank's usage.

Second: to use Dawkin's convenient metaphor, could a sighted watchmaker really design a "simpler" immune system, for example, than a blind one? If so, is that to do with historical aspects of evolution/population sizes/mutational spectra and so on? If so, what sort of science are we as biologists looking to create: one that says, for example: genome duplication events and large population sizes are responsible for ... what? This would be restricting our thinking about evolution to providing explanations of the contrasting failings of different groups of creatures.

More interesting to me is what I believe Frank is suggesting: the blind watchmaker may have much to teach the sighted ones. This is particularly so in the case of Artificial Intelligences. These are of great interest as both biologists and engineers are only at the starting gate of understanding, and we are all dealing with the question of the design of systems which have a tiny number of component types – "neurons". I expect to see a unification of psychology and AI engineering in the near future.

Time will tell whether any notions that may have floated around in classical thinking about evolutionary genetics will advance this program.

References

1. Lee E.A.: Cyber physical systems: Design challenges. *In Object oriented real-time distributed computing (isorc), 2008 11th ieee international symposium on, IEEE*. 2008. 363-369 [Publisher Full Text](#)

Is the topic of the opinion article discussed accurately in the context of the current literature?

Yes

Are all factual statements correct and adequately supported by citations?

Yes

Are arguments sufficiently supported by evidence from the published literature?

Yes

Are the conclusions drawn balanced and justified on the basis of the presented arguments?

Yes

Competing Interests: No competing interests were disclosed.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Author Response (F1000Research Advisory Board Member) 05 Aug 2017

Steven Frank, Department of Ecology & Evolutionary Biology, University of California, Irvine, USA

I thank Sean Nee for his thoughtful comments that extend the scope of the discussion. This exchange is included as part of the final publication of my article, so I confine my response to these comments.

Italics quote from Nee's review.

First: is engineering really as cut-and-dried as we suppose? ... it is not obvious to me that engineering is as purely "deductive" a subject as we might like to think as we board an aircraft, using the word "deductive" in a way that I may be misconstruing as Frank's usage.

I certainly agree that engineering in practice works as much by trial and error as it does by an abstractly pure deduction. I had intended the labeling of engineering design as "deductive" primarily as a comparison with how evolutionary design by natural selection is relatively more "inductive" in character.

Engineers do use deductive principles of control theory to aid in the design of control systems. When novel deductive understanding of control principles arises, engineers readily

alter their approach to design, benefitting from the improved general insight. By contrast, natural selection is essentially a purely inductive process. That inductive process cannot throw out a past design and start over, but must improve only by layering small inductive gain upon small inductive gain.

Second: to use Dawkin's convenient metaphor, could a sighted watchmaker really design a "simpler" immune system, for example, than a blind one? ... More interesting to me is what I believe Frank is suggesting: the blind watchmaker may have much to teach the sighted ones. This is particularly so in the case of Artificial Intelligences.

I meant the comparison in both of the ways that Nee discusses. A sighted watchmaker would make a different immune system from a blind watchmaker. Whether that different immune system of the sighted watchmaker would be simpler or better is hard to say. However, I suspect that it would be simpler because humans tend to design systems that they can analyze and understand, whereas blind induction does not care about the logic or the complexity of the mechanism.

I agree with the latter aspect in Nee's comments: that the blind watchmaker provides new insights about design that the sighted watchmaker may consider. We are seeing this now in the great advances in artificial intelligence: the goal of the sighted watchmaker has become to improve the ways in which the blind watchmaker's trial and error induction proceeds.

Competing Interests: No competing interests were disclosed.

The benefits of publishing with F1000Research:

- Your article is published within days, with no editorial bias
- You can publish traditional articles, null/negative results, case reports, data notes and more
- The peer review process is transparent and collaborative
- Your article is indexed in PubMed after passing peer review
- Dedicated customer support at every stage

For pre-submission enquiries, contact research@f1000.com

F1000Research